ELSEVIER

Discrete Applied Mathematics 93 (1999) 205–232

DISCRETE APPLIED MATHEMATICS

Compound constructions of broadcast networks

Michael J. Dinneen[a], Jose A. Ventura[b], Mark C. Wilson[c,*], Golbon Zakeri[d]

[a] *Department of Computer Science, University of Auckland, Private Bag 92019 Auckland, New Zealand*
[b] *Department of Industrial and Manufacturing Engineering, The Pennsylvania State University, University Park, PA 16802, USA*
[c] *Department of Mathematics, Northern Illinois University, DeKalb, IL 60115, USA*
[d] *MCS Division, Argonne National Laboratory, Argonne, IL 60439, USA*

Received 4 January 1997; revised 1 October 1998; accepted 16 November 1998

Abstract

Compound methods have been shown to be very effective in the construction of broadcast graphs. Compound methods generate a large broadcast graph by combining multiple copies of a broadcast graph G using the structure of another broadcast graph H. Node deletion is also allowed in some of these methods. The subset of connecting nodes of G has been defined as solid h-cover by Bermond, Fraigniaud and Peters, and center node set by Weng and Ventura. This article shows that the two concepts are equivalent. We also provide new properties for center node sets, including bounds on the minimum size of a center node set, show how to reduce the number of center nodes of a broadcast graph generated by a compound method, and propose an iterative compounding algorithm that generates the sparsest known broadcast graphs in many cases.

MSC: primary 90B12; secondary 05C90; 68R10; 94A05

Keywords: Broadcasting; Communication network; Minimum broadcast graph; Graph compound

1. Introduction

Communication in networks is a process whereby a set of messages, generated by a set of originators, is transferred to a set of receivers. The nodes of the network are the possible originators and receivers of the messages, and the edges are the communication lines which allow the direct transmission of messages between certain pairs of nodes. There is a wide range of network design problems in communication networks which are differentiated by placing constraints upon the set of messages, the originators, the

* Corresponding author.
E-mail addresses: mjd@cs.auckland.ac.nz (M.J. Dinneen), jav1@psu.edu (J.A. Ventura), wilson@math.niu.edu (M.C. Wilson), zakeri@mcs.anl.gov (G. Zakeri).

PII: S0166-218X(99)00043-8

receivers, the edges, the network topology, and the transmission characteristics of the network [1,10,13,14,20,23]. A communication network can be modeled as a connected graph $G=(V(G),E(G))$ without loops or parallel edges, consisting of a set of nodes $V(G)$ with cardinality $v(G)$, and a set of undirected edges $E(G)$ with cardinality $e(G)$. We shall freely interchange the network and graph-theoretic terminology throughout.

Broadcasting is a special type of network communication in which a single message, originated at any node, is transmitted to all the other nodes of the network. Broadcasting is usually required to be completed as rapidly as possible by a sequence of transmissions through the communication lines. It is assumed that broadcasting is carried out under the following three constraints [8,9]: (i) each transmission requires one unit of time, (ii) a node can make at most one transmission in one time unit, and (iii) a node can only transmit the message to its neighbors (two nodes are called neighbors if they are connected by an edge). Thus, in one time step, the number of informed nodes can at most be doubled. This implies that after m time steps the number of nodes that have received the message, including the originator, is at most 2^m. The *broadcast time* $b(G)$ of a graph G is the minimum number of time steps in which broadcasting can be achieved in G regardless of the originator of the message. From the above it is clear that $b(G)\geqslant\lceil\log_2 v(G)\rceil$. A *broadcast graph* is usually defined to be a graph G satisfying $b(G)=\lceil\log_2 v(G)\rceil$ (this terminology, though somewhat illogical, is well established). Complete graphs are obviously broadcast graphs.

Suppose that a node u in a network G is the originator of the message. A *broadcast protocol* (or broadcast tree) $P(G,u)$ is a rooted spanning tree in which the originator u is the root and all the nodes are labeled by their receiving times. In a broadcast protocol, each edge is used exactly once and the message is always transmitted from parent to child. In order that a network G be a broadcast graph, each node in the network must have a broadcast protocol that can be completed in $\lceil\log_2 v(G)\rceil$ time steps.

The problem of recognizing whether an arbitrary network is a broadcast graph is $\mathcal{NP}$-complete [8]. A *minimum broadcast graph* (mbg) is a broadcast graph with the minimum possible number of edges for its given number of nodes, and the *broadcast function* $B(n)$ is defined to be the number of edges of every mbg with n nodes. There is no known feasible method for determining $B(n)$ for an arbitrary value of n. Farley et al. [9] showed that hypercubes are mbg's and so $B(2^m)=m2^{m-1}$ for $m\geqslant 0$. Khachatrian and Haroutunian [15] and Dinneen et al. [7] proved independently that $B(2^m-2)=(m-1)(2^{m-1}-1)$ for $m\geqslant 2$. Farley et al. also determined the values of $B(n)$, for $1\leqslant n\leqslant 15$. Bermond et al. [2] and Weng and Ventura [25] published known values of $B(n)$ for $17\leqslant n\leqslant 63$. Recently, Saclé [22] gave lower bounds on $B(2^m-k)$, for $m\geqslant 3$ and $3\leqslant k\leqslant 6$, Fig. 1 shows an example of an mbg with 15 nodes and 24 edges and one of its broadcast protocols.

Direct construction of broadcast graphs is a difficult process, since in the worst case one must check that every node has a broadcast protocol taking time $\lceil\log_2 v(G)\rceil$. Thus most authors have concentrated on constructions which combine several known broadcast graphs to create new ones.

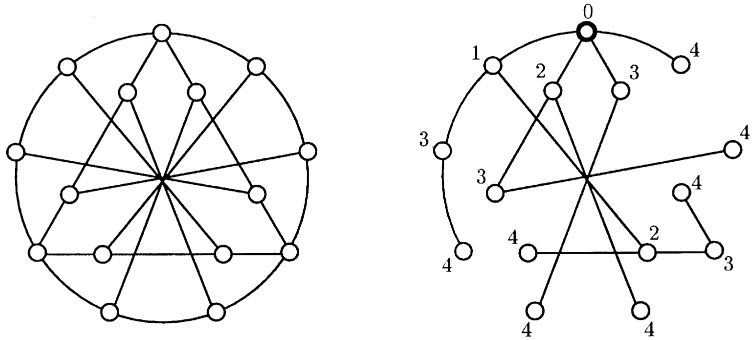

(a) Mbg with 15 nodes. (b) Broadcast protocol.

Fig. 1. An mbg with 15 nodes and a selected protocol.

In 1979 Farley [8] introduced broadcast graphs, proposed a recursive algorithm to construct broadcast graphs with an arbitrary number of nodes n, and showed that the number of edges of the broadcast graphs produced by his algorithm is bounded by $(n/2)\lceil \log_2 n \rceil$. Chau and Liestman [5] developed an algorithm which constructs broadcast graphs by interconnecting 5, 6 and 7 smaller broadcast graphs. They also improved Farley's bound on $B(n)$, for n in the range $(2^{m-1}, 7.2^{m-3})$, $m \geqslant 1$. Gargano and Vaccaro [11] proposed three algorithms based on the interconnection of hypercubes of small dimension to build up larger broadcast graphs. Chen [6] presented a method similar to the second algorithm of Gargano and Vaccaro, and then suggested the recursive application of his first method to construct larger broadcast graphs. In Grigni and Peleg's algorithm [12], hypercubes and generalized Fibonacci numbers are used to construct broadcast graphs with $O(L(n)n)$ edges, where $L(n)$ is the number of leading 1's in the binary representation of $n-1$. In practice, however, the other methods seem to require fewer edges. Bermond et al. [3] proposed four methods to construct broadcast graphs and used them to produce new broadcast graphs for $18 \leqslant n \leqslant 63$. Ventura and Weng [24] developed a method based on the concepts of aggregated nodes and aggregated edges (which are used to replace ordinary nodes and edges, respectively, of known mbgs, for $9 \leqslant n \leqslant 15$) to construct sparse broadcast graphs. The central idea of all the methods discussed so far is to produce larger broadcast graphs by combining small known mbg's or broadcast graphs using as few edges as possible without violating the constraint of being able to broadcast in minimum time from any node.

More recently another class of combination methods using compound graphs has been developed by Bermond et al. [2], generalizing a construction of Khachatrian and Haroutunian [15]. A more general method, which allows for systematic vertex deletion, was proposed by Weng and Ventura [25]. A key ingredient of this last method, called the *doubling procedure*, is a center node set, defined via so-called official broadcasting.

In this article, we investigate the constructions of [2,25], treating official broadcasting and center node sets in more detail. Iterative algorithms based on these constructions are presented and analyzed. Computational results are obtained which improve most of the best known bounds on numbers of edges and size of center node sets achieved by previous methods.

This article is rather more self-contained than might be expected, and includes several known results. Our purpose in doing this is threefold. The experience of two of the authors made it clear that a full account would be valuable to those entering the field. In addition, in the course of our literature search we encountered many errors and omissions, some rather serious, in the references in the bibliography of this article. Finally, several basic and useful results have been repeatedly rediscovered, because, in our opinion, they have never been explicitly stated and proved in the published literature. It is our hope that our approach will help to remedy these problems.

An outline of the paper is as follows. In Section 2 we recall the definitions of official broadcasting and center node set and demonstrate the equivalence of the latter concept and that of solid h-cover. We derive bounds on the minimal size of a center node set. Section 3 focuses on ways of reducing the center node sets generated by the compounding procedures. The short Section 4 discusses the general framework of the iterative algorithms and shows the limitations of the center node reductions. Section 5 has three subsections. In the first of these we discuss our practical implementation of the iterative algorithms. In the second part the initial input used by the algorithms is verified. An important feature here is an efficient and accurate calculation of good center node sets for known broadcast graphs. In the third part, we present computational results that compare our refined algorithm with previously described methods. A table containing the known values of $B(n)$ and the best upper bounds on $B(n)$, for $17 \leqslant n \leqslant 127$, is also presented. Finally, some open problems and directions for future research are discussed in Section 6.

2. Preliminaries

This section lays out the basic notation and definitions concerning center node sets which will be used in the rest of the paper. The idea of official broadcasting, described below, leads naturally to the definition of a center node set. We establish the equivalence of the center node sets of Weng and Ventura [25] and the solid h-covers of Bermond et al. [2]. In so doing we recall the compounding methods introduced in these papers. We obtain elementary bounds for the minimum size of center node sets which will prove useful later.

The following definition will be used throughout. Let (n,k,i) be a triple of integers with $n > 0$ and $0 \leqslant i < k$. Then

$$\lceil \log_2(nk-i) \rceil \leqslant \lceil \log_2(nk) \rceil \leqslant \lceil \log_2 n \rceil + \lceil \log_2 k \rceil.$$

We say that the triple (n,k,i) *satisfies the broadcast condition* if equality holds in both inequalities. If $i = 0$ we say simply that (n,k) satisfies the broadcast condition.

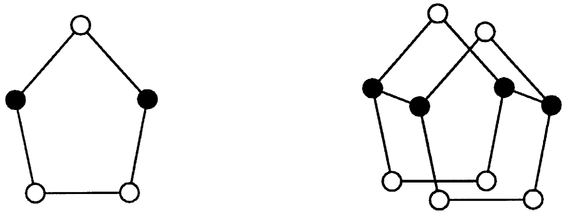

(a) Mbg with 5 nodes. (b) Mbg with 10 nodes.

Fig. 2. Constructing an mbg with 10 nodes by compounding.

We recall that a *vertex cover* of a graph G is a subset S of $V(G)$ such that every edge of G is incident to at least one vertex of S.

2.1. Official broadcasting and center node sets

The *compound* of a graph G into another graph H relative to a set $S \subseteq V(G)$, denoted by $G_S[H]$, is the graph obtained by replacing each node of H with a copy of G, and each edge of H by a matching between the corresponding copies of S. Thus, $V(G_S[H]) = V(G) \times V(H)$ in a natural way.

In [15] a compounding algorithm was developed in which G is a broadcast graph restricted to have a maximum degree bounded by $\lceil \log_2 v(G) \rceil - 1$, S is a vertex cover of G and $H = K_2$.

In [2] a more general compounding procedure was presented, in which G and H are general broadcast graphs satisfying the broadcast condition, and S was what was called a solid h-cover of G (see Section 2.2 below). Fig. 2 illustrates this method by generating an mbg with 10 nodes from two copies of an mbg with 5 nodes. In this example S is a solid 2-cover defined by the black nodes in Fig. 2(a) and $H = K_2$.

It is not possible to generate broadcast graphs with n nodes for every n by this procedure (for example, when n is prime). Weng and Ventura [25] proposed a generalized compounding algorithm, which they called the *doubling procedure*, which includes the method of [2] as a special case. The doubling procedure's extra generality mainly arises from the fact that nodes may be deleted in certain copies of G. The method constructs a broadcast graph from given broadcast graphs G and H, and an integer i with $0 \leqslant i \leqslant v(H) - 1$, such that $(v(G), v(H), i)$ satisfies the broadcast condition. The subset S of nodes of G used for connecting the copies of G was called a center node set in [25].

Weng and Ventura introduced the concepts of official broadcasting and center node sets in order to describe the doubling procedure. In *official broadcasting*, certain nodes of the network, called *center nodes*, are given the authority to make a message official. A message, originated at any node, must be made official during the broadcast protocol. The official message must then be transmitted to all the nodes of the network. During the broadcast protocol a message is official if it has been officialized by a center node; otherwise, it is unofficial. It is assumed that an unofficial message becomes official immediately after it arrives at a center node. In official broadcasting, where all the nodes must receive an official message, a center node will only receive one message, so that if the incoming message is unofficial, it will be officialized immediately after its arrival. A non-center node may receive one or two messages. In the first case, the message must be official. In the second case, the first message must be unofficial and the second one official. In addition, in official broadcasting, it is possible for a non-center node to send an unofficial message to a neighbor and receive an official message during the same time step.

A more formal description is as follows. An *official broadcast protocol* for a node u with respect to a set $S \subseteq V(G)$ in a graph G, denoted by $P(G,u;S)$, is a spanning subgraph of G containing all of its nodes, in which the nodes are labeled by one or two receiving times, all of which are at most $b(G)$. If a node has two receiving times, it must not belong to S, the first receiving time is for the unofficial message and the second one for the official message. If $u \in S$, each node is labeled by a single receiving time, and the official broadcast protocol is a spanning tree rooted at u, that is, just an ordinary broadcast protocol.

Given an official protocol $P = P(G,u;S)$ we define the *unofficial part* P_u of P to be the tree rooted at u induced by all edges that transmit an unofficial message and the corresponding nodes. The *official part* P_o of P is a forest of rooted trees induced by all edges which transmit an official message. These trees are rooted at nodes in S which receive an unofficial message. Denote the set of all such nodes by $V_{cu}(P)$. The forest P spans G.

In particular, if every node of a graph G has an official protocol of time $b(G)$ with respect to S, then S is called a (minimal) *center node set* of G. Clearly if S is a center node set and $S \subseteq T \subseteq V(G)$ then T is also a center node set. In this paper we shall be interested only in the case where G is a broadcast graph. In this case, unless otherwise stated all protocols are assumed to take $\lceil \log_2 v(G) \rceil$ time steps.

If P is an ordinary broadcast protocol, then a node is *idle at time t* if it is aware of the message at time $t-1$ and does not transmit the message at time t. It is important to note that a node u can be a non-center node only if in some broadcast protocol for u, some node is idle and can send the official message back to u. In practice we construct official protocols by using ordinary protocols and showing that there are enough idle nodes to inform all non-center nodes that receive an unofficial message.

The notion of official broadcasting was introduced purely as a way to describe the set of connecting nodes in a compound broadcast graph. However it may have other applications, such as minimizing costs in the design of networks in which message authentication is required. Hence obtaining the smallest possible center node sets of a given graph is a problem of some interest.

Given a graph G, define the *center node number* $cn(G)$ to be the minimum size of all center node sets for G. A center node set of size $cn(G)$ is called an *optimal center node set* (*ocns*) of G.

There may exist multiple ocns's for a given graph G, and the cardinality of ocns's of non-isomorphic graphs with n nodes and m edges may be different. There is no known

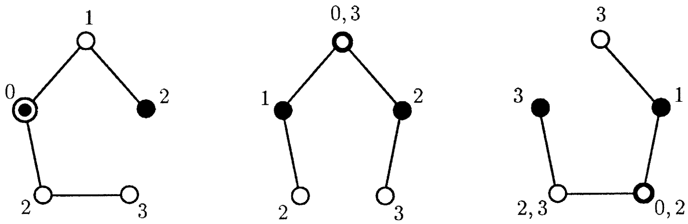

Fig. 3. The official broadcast protocols for the mbg in Fig. 2(a).

polynomial-time algorithm for computing $cn(G)$ for an arbitrary graph G. In the mbg with 5 nodes presented in Fig. 2(a), the black nodes are center nodes and the white nodes are non-center nodes. Actually, the two center nodes define an ocns for the mbg. Fig. 3 shows the three different official broadcast protocols for this mbg with 5 nodes.

2.2. Compound methods

We can now describe the *doubling procedure* of [25]. Fix a broadcast graph G with center node set S. If $S \neq V(G)$ then choose a non-center node v of G (usually of minimal degree). Otherwise let v be any (center) node of S. Construct a new network G^v by deleting v and all its incident edges from G, and adding the required edges to form a clique among the neighbors of v. Let $S^v = S \setminus \{v\}$.

Let H be a broadcast graph. For a fixed integer i with $0 \leqslant i \leqslant v(H) - 1$, we can construct a network $\mathcal{G}$ by connecting $v(H) - i$ copies of G and i copies of G^v as follows. For each fixed $s \in S^v$, connect all $v(H)$ copies of s to form a copy H_s of H. If v is a center node of G, then choose any broadcast graph with $v(H) - i$ nodes and connect the $v(H)-i$ copies of v to form a graph H^* isomorphic to this broadcast graph. If v is a non-center node then let H^* be the empty graph. The procedure is illustrated in Fig. 4. We shall at times refer to *vertical* (within G or G^v) or *horizontal* (within H or H^*) broadcasting; these terms should be interpreted as illustrated in that figure.

Theorem 2.1 (The Doubling Procedure [25]). *Let G and H be broadcast graphs and let S be a center node set for G. Suppose that $(v(G), v(H), i)$ satisfies the broadcast condition.*

(i) *The graph $\mathcal{G}$ constructed by the doubling procedure above is a broadcast graph.*
(ii) *The set $\bigcup_{s \in S^v} v(H_s) \cup v(H^*)$ is a center node set for $\mathcal{G}$.*

We point out here that the example on p. 292 of [25] is invalid, since $(5, 17)$ does not satisfy the broadcast condition.

The doubling procedure has the following important special case.

Theorem 2.2 (The Ordinary Compounding Method (cf. [2])). *Let G and H be broadcast graphs and let S be a center node set for G. Suppose that $(v(G), v(H))$ satisfies the broadcast condition.*

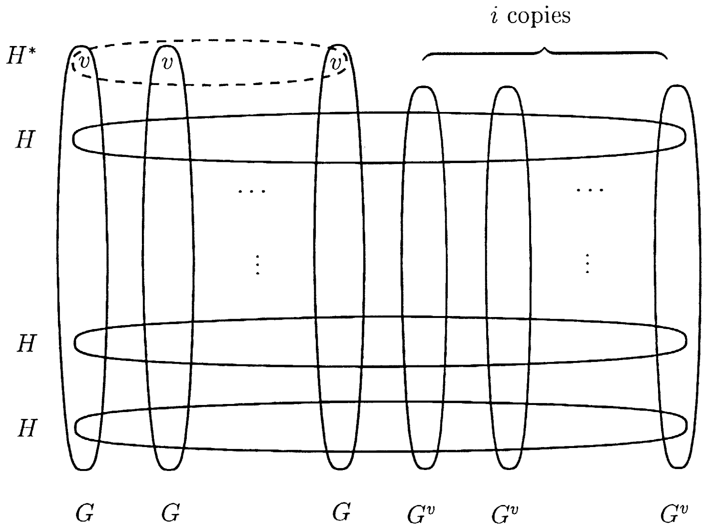

Fig. 4. The doubling procedure of Weng and Ventura (see [25]).

(i) *The compound graph* $\mathscr{G} = G_S[H]$ *is a broadcast graph.*
(ii) *The product* $S \times V(H)$ *is a center node set for* $\mathscr{G}$.

In [2] similar results to Theorem 2.2(i) are presented. The only difference is that S is required to be a solid 1-cover or solid 2-cover as defined in that paper. We now establish the connection between the two papers.

Recall from [2] that for a given broadcast graph G, a subset S of nodes is a *solid* 1-*cover* if S is a vertex cover of G and for every $u \notin S$, there is a broadcast protocol for u such that at least one neighbor of u is idle at some time during the broadcast.

Proposition 2.3. *A subset S of $V(G)$ is a solid* 1-*cover if and only if S is both a vertex cover and a center node set.*

Proof. Let S be a solid 1-cover for G. If $u \in S$ then there is an official protocol for u by definition. If $u \notin S$ then every neighbor of u belongs to S. Thus in the protocol guaranteed by the definition of solid 1-cover, for every node v except u, each message is official on arrival at v. It remains only to inform u officially, and this is done by the idle neighbor supplied by the definition.

For the reverse direction, if S is a center node set and vertex cover then for each $u \notin S$ there is an official protocol for u with respect to S. The node which sends the official message to u is necessarily in S and must be idle in some broadcast protocol for u. $\square$

In [2] it is stated that one can define solid h-covers for every $h \geqslant 1$ and the conclusion of Theorem 2.2(i) will hold with S assumed to be a solid h-cover. An inspection of the proof shows that it is necessary to *define* a solid h-cover to be a center node set which is an h-cover, that is, a path cover for all paths of length h (note that a vertex cover is a 1-cover). Since every center node set is an h-cover for some h, the two concepts then coincide. For brevity, we will refer to center node sets that are also vertex covers as solid 1-covers.

It is possible that one might obtain a broadcast graph by compounding with respect to a set of nodes that is not a center node set (this would require a different method of proof to that in [25] and [2]). We believe that this is not the case and that center node sets are crucial for such a result.

Conjecture 2.4. Let G and H be broadcast graphs and let $S \subseteq V(G)$. If $G_S[H]$ is a broadcast graph then S is a center node set for G.

Even when $H = K_2$ the above conjecture seems rather difficult.

2.3. General bounds on center node sets

We now present some elementary results which are useful in obtaining bounds on $cn(G)$.

Perhaps the easiest way to obtain a good center node set is the following widely applicable result which follows readily from Proposition 2.3 and [2].

Proposition 2.5. *Let G be a broadcast graph of broadcast time $m = \lceil \log_2 v(G) \rceil$ and let $S \subseteq V(G)$. If S is a vertex cover for G such that every neighbor of every node not in S has degree at most $m-1$, then S is a center node set for G.*

Proof. Let v be a node not in S. In every broadcast protocol P for v the first neighbor receiving the unofficial message must be idle at time no later than m and can therefore send the official message back to v. All other nodes already have an official message and so this yields an official protocol for v. $\square$

We now proceed to derive a lower bound for $cn(G)$ that is valid under special conditions on $v(G)$. Observations similar to the following two results have been made by many authors (see for example [4,9,11,12,17,18]). The next result also holds in the case when the node v is never idle.

Proposition 2.6. *Let G be a broadcast graph with broadcast time m. Let v be a node of G. Suppose that in a broadcast protocol for v, v itself is idle at times $t_1 < t_2 < \cdots < t_l$.*

(i) $v(G) \leqslant 2^m + 1 - 2^l$.

(ii) *Equality occurs if and only if v is idle for the last l time steps, and no other node is idle.*

Proof. Denote by $I(i)$ the number of nodes informed by time i. Then $I(t_1-1)\leqslant 2^{t_1-1}$. Since v is idle at time t_1 and the size of the forest formed by the informed nodes other than v can at most double at time t_1, we have $I(t_1)\leqslant 2(2^{t_1-1}-1)+1=2^{t_1}(1-2^{-t_1})$. An easy induction yields $I(t_l)\leqslant 2^{t_l}(1-(2^{-t_1}+\cdots+2^{-t_l}))$. Thus $I(m)\leqslant 2^{m-t_l}I(t_l)\leqslant 2^m(1-(2^{-t_1}+2^{-t_2}+\cdots+2^{-t_l}))$. This last expression is clearly maximized when and only when $t_1=m+1-l,\ldots,t_l=m$ and the corresponding maximal value is 2^m+1-2^l. □

Let $\delta(G)$ denote the minimum degree of a node of G.

Corollary 2.7. *Let G be a broadcast graph with broadcast time m. Then*

$$\delta(G)\geqslant m-\lfloor\log_2(2^m+1-v(G))\rfloor.$$

Proof. If an originating node of G has degree $k\leqslant m$ then it must be idle at least $m-k$ times. Thus $v(G)\leqslant 2^m+1-2^{m-k}$ by Proposition 2.6. Solving for k yields $k\geqslant m-\log_2(2^m+1-v(G))$ and the result follows since $\delta(G)\in\mathbb{Z}$. □

A weaker bound, with the floors replaced by ceilings, is contained in Theorem 2.2 of [12]. The latter bound has a simple interpretation in terms of the binary expansion of $v(G)-1$, but it seems that the bound above does not.

The corollary yields the important special cases that broadcast graphs on 2^m, 2^m-1 and 2^m-2 nodes have minimum possible degrees m, $m-1$ and $m-1$, respectively.

For each $n\in\mathbb{Z}$ let $V_n(G)$ denote the set of nodes of G of degree n.

Proposition 2.8. *Let G be a broadcast graph with* $2^m+1-2^{m-\delta}$ *nodes for some* δ *with* $0\leqslant\delta\leqslant m$.

(i) *Let v be a node of degree* δ. *Then v is contained in every center node set of G. Hence* $cn(G)\geqslant|V_\delta(G)|$.

(ii) *If* $\delta<m$ *and* V_δ *is a vertex cover then* $cn(G)=|V_\delta(G)|$.

Proof. By Proposition 2.6(ii), in an official broadcast protocol for v with respect to any subset S of $V(G)$, the only node that can possibly be idle is v itself (for the last $m-\delta$ time steps). Thus no node can send an official message back to v and so necessarily $v\in S$. This proves (i), and (ii) follows directly from Proposition 2.5. □

Taking $\delta=m$ and $\delta=m-1$, respectively, we see that (i) above generalizes Theorems 2.1 and 2.3 of [25].

Finally, general bounds on $cn(G)$ are given by the next result. The lower bound is attained by all complete graphs. We believe that the degree hypothesis in (ii) is redundant, but do not have a proof of this.

Proposition 2.9. *Let G be a broadcast graph with broadcast time* $m=\lceil\log_2 v(G)\rceil$.

(i) *If* $v(G)=2^m$ *then* $cn(G)=v(G)$, *whereas if* $v(G)<2^m$ *then*

$$m-\lfloor \log_2(2^m-v(G))\rfloor \leqslant cn(G).$$

(ii) *If there is a subset* R *of nonadjacent nodes of* $V(G)$ *all of whose neighbors have degree at most* $m-1$, *then*

$$cn(G)\leqslant v(G)-|R|.$$

Proof. If $v(G)=2^m$ then no node can be idle during a broadcast so all nodes are center nodes. If $v(G)<2^m$ then suppose that k is the size of a center node set for G. If $k>m$ the result follows trivially so we may assume $k\leqslant m$. Let v be a non-center node. Then, in every broadcast originating at v, the maximum number of nodes with the official message at time k is 2^k-1, and this occurs only if all neighbors of v are center nodes. Thus the maximum size of the official forest at time m is $2^{m-k}(2^k-1)$. This implies that $2^m-2^{m-k}\geqslant v(G)$. Solving for k yields $k\geqslant m-\log_2(2^m-v(G))$ and (i) follows since $k\in\mathbb{Z}$. nodes. Part (ii) follows from Proposition 2.5 since $V(G)\setminus R$ is a vertex cover for G. □

Of course for a given $v(G)$, a reduction in the number of edges (while keeping the resulting graph a broadcast graph) may force an increase in the number of center nodes. For example, Conjectures 1 and 2 in [25] say essentially that for each mbg G of order $v(G)=2^m-1$: (i) all nodes have degree m or $m-1$, (ii) there are $\lceil v(G)/(m+1)\rceil$ nodes of degree m, and (iii) the set of all nodes of degree $m-1$ is a vertex cover of G. Thus in this case $cn(G)$ is asymptotic to $\frac{m}{m+1}(2^m-1)$, whereas the lower bound above is only $m+1$. This playoff between edges and center nodes is of crucial importance in the compounding methods.

3. Reducing center node sets in compound methods

In this section we consider ways of reducing the center node sets generated by the compounding procedures under consideration. The bounds on $cn(\mathcal{G})$ given by Theorems 2.1 and 2.2 are very poor in general and they deteriorate upon further compounding. The main results in this section are Theorem 3.1 and especially Theorem 3.3, which improve these bounds considerably.

We first treat the simpler special case of Theorem 2.2. The next result reduces the center node sets obtained by applying that theorem.

Theorem 3.1. *Let* G *and* H *be broadcast graphs and let* S,T *be center node sets for* G,H *respectively. Suppose that* $(v(G),v(H))$ *satisfies the broadcast condition. Then* $\mathcal{S}=S\times T$ *is a center node set for* $\mathcal{G}=G_S[H]$.

Proof. Let u be a node of $\mathcal{G}$. If u belongs to some copy H_s of H then first broadcast officially in H_s with respect to T_s. After $\lceil \log_2 v(H)\rceil$ time steps all nodes in H_s have

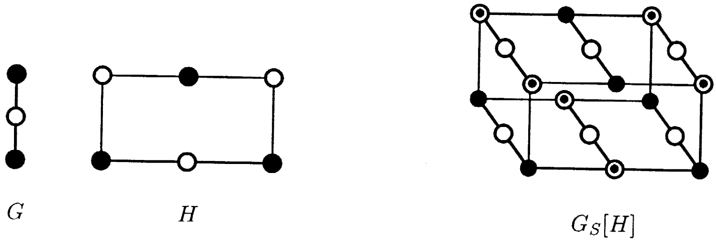

Fig. 5. Compounding with center node reduction.

the official message. Now broadcast vertically in the appropriate copy of G. All nodes of $\mathscr{G}$ are officially informed by time $\lceil \log_2 v(G) \rceil + \lceil \log_2 v(H) \rceil$.

Now suppose that u is a node of some copy G_i of G and does not belong to any H_s, i.e. u does not belong to any copy S_i of S. Let P be an official broadcast protocol for u in G with respect to S. Recall the notation $V_{cu}(P)$ from Section 2. Broadcast in G_i according to P_u, ending at $V_{cu}(P)$. On receiving the message each element of $V_{cu}(P)$ then broadcasts horizontally in its corresponding copy of H, and this broadcast is official and takes $\lceil \log_2 v(H) \rceil$ time steps. After this is completed, each node in each copy of $V_{cu}(P)$ continues according to P_o. Since it takes $\lceil \log_2 v(G) \rceil$ time steps to perform official broadcasting in G according to P_u and P_o, the total time needed to broadcast officially in $\mathscr{G}$ is $\lceil \log_2 v(G) \rceil + \lceil \log_2 v(H) \rceil$. $\square$

Fig. 5 exemplifies this result. Here S is the set consisting of the two black nodes of G. The circled black nodes in $G_S[H]$ are those center nodes produced by Theorem 2.2 which can be removed according to Theorem 3.1.

We now collect some data about the above construction for later reference.

Corollary 3.2. *The following statements hold for the ordinary compounding method.*

$$v(\mathscr{G}) = v(G)v(H)$$
$$e(\mathscr{G}) = v(H)e(G) + |S|e(H)$$
$$|\mathscr{S}| = |S||T|$$

We now describe a slight generalization of the doubling procedure (as described immediately before Theorem 2.1). We do not require all H_s to be isomorphic, only that the H_s are broadcast graphs and have $v(H_s)$ all equal to a common integer which we denote $v(H)$. The $e(H_s)$ need not have the same value. Finally, H^* need not be a broadcast graph; it is only necessary that broadcast time of H^* is at most that of each H_s. We refer to this construction as the *generalized compounding method*. In order to avoid excessively long formulas, the resultant graph will simply be denoted by $\mathscr{G}$.

Under these weakened hypotheses we can prove the result below. Note that part (ii) drastically reduces the size of the center node set obtained in Theorem 2.1. The condition on U in (ii) is always satisfied if H^* is a broadcast graph and U a center node set for H^*.

Theorem 3.3 (The Generalized Compounding Method). *Suppose that* $(v(G), v(H), i)$ *satisfies the broadcast condition.*

(i) *The graph* $\mathscr{G}$ *constructed by the generalized compounding method above is a broadcast graph.*

(ii) *Let* T_s *be a center node set for* H_s. *Choose* $U \subseteq V(H^*)$ *so that official broadcast in* H^* *with respect to* U *can be completed in* $\lceil \log_2 v(H) \rceil$ *time steps. Then* $\mathscr{S} = \bigcup_{s \in S^v} T_s \cup U$ *is a center node set for* $\mathscr{G}$.

Proof. Recall that it was proved in Theorems 3.1 and 3.2 of [25] that official broadcast in G^v can be carried out with respect to S^v in $\lceil \log_2 v(G) \rceil$ time steps. Furthermore, for each (official) protocol P originating at any node of G other than v, there is a modification P^v of P which works for G^v. Only for case 4 below is it necessary to know the details of P^v and these details are mentioned there.

Let u be a node of $\mathscr{G}$. We show that official broadcast with respect to $\mathscr{S}$ is possible from u in $\lceil \log_2 v(G)v(H) \rceil$ time steps. There are 4 cases depending on the location of u.

Case 1: First suppose that u belongs to some H_s. As in the proof for the special case above, official broadcast within H_s with respect to T_s can be completed in $\lceil \log_2 v(H) \rceil$ time steps. At this stage all nodes in H_s have an official message. In the remaining $\lceil \log_2 v(G) \rceil$ time steps, each node in H_s broadcasts vertically inside its copy of G or G^v. By this time, all nodes in $\mathscr{G}$ are officially informed, proving the result in this case.

Case 2: Next suppose that u is a node of H^*. Then by definition of H^*, all nodes of G are center nodes and so all nodes in $\mathscr{G}$ belong to H^* or some H_s. Also u belongs to a unique copy G_i of G. First broadcast vertically in G_i, which takes $\lceil \log_2 v(G) \rceil$ time steps, at the end of which all nodes in G_i have the message. Now broadcast horizontally and officially in the appropriate H_s or H^*. This can be done in $\lceil \log_2 v(H) \rceil$ time steps and so all nodes are officially informed in the correct time.

Case 3: Suppose now that u is a node of some copy G_i of G but not a node of any H_s. As in the proof of Theorem 3.1, broadcast according to the unofficial part P_u of some official protocol P, ending at $V_{cu}(P)$. Each center node in $V_{cu}(P)$ broadcasts horizontally and officially in its corresponding copy of H. Then each copy t of a node in $V_{cu}(P)$ continues as in P_o (resp. $(P^v)_o$) if t belongs to a copy of G (resp. G^v). As in the special case above this can all be completed in the required time.

Case 4: Finally, we consider the most difficult case, where u is a node of some G^v and not a node of any H_s. Note that since $u \notin H_s$, u is not a center node. Thus not all nodes of G are center nodes and hence by construction of G^v, v is not a center node.

It is now necessary to consider the details, given in [25], of the definition of P^v.

Let P be a broadcast protocol originating at u, and let v' broadcast to v in P at time t. Suppose that v calls its neighbors $v_1, \ldots, v_k$ at times $t_1 < \cdots < t_k$, respectively. Then P^v is obtained by eliminating the messages from v and adding a call from v' to v_1 at time t and v_i to v_{i+1} at time t_i for $1 \leqslant i \leqslant k-1$. This is possible since in the construction of G^v, v was deleted and a clique C was formed between all its neighbors in G. Once each v_i is informed, it broadcasts away from C just as it would in P. The exception

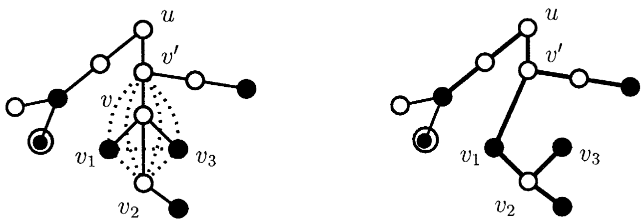

Fig. 6. Example for the proof of Theorem 3.3.

is v_k, which receives the message 1 time unit earlier than it would have in P, so we make v_k wait idle for 1 time unit before proceeding as in P. Then by construction P^v takes the same time as P. Note that outside C, P^v and P are exactly the same.

Now suppose further that P is an official protocol with respect to S. Recall that $V_{cu}(P)$ denotes the subset of center nodes in S which receive an unofficial message in P. We proceed to define a subprotocol Q, which will be used below in concluding the proof of the theorem, as follows.

Let $\overline{P}$ be the tree obtained from P by deleting the subtree rooted at v. Note that protocols P and P^v are identical as far as nodes in $\overline{P}$ are concerned. The restriction of Q to $\overline{P}$ is defined to be the same as P_u. For $P \setminus \overline{P}$ let Q follow P^v exactly to the point (and no further) that all nodes in $V_{cu}(P)$, which have v as an ancestor, are informed. By construction of P^v and Q, the time taken to execute Q is precisely the time taken to execute P_u.

The constructions of the last few paragraphs are illustrated in Fig. 6. The solid nodes are center nodes. In the left picture, we see P. All those center nodes except the circled one lie in $V_{cu}(P)$. The clique C is indicated by the dotted lines. The right picture shows P^v. The thicker edges define the subprotocol Q.

We now conclude the proof of case 4 and of the theorem. Let P be an official protocol for u in G. Broadcast in G^v according to Q. Now broadcast horizontally and officially in the appropriate H_s, which takes $\lceil \log_2 v(H) \rceil$ time steps. Note that every copy of each $w \in V_{cu}(P)$ has the official message in each copy of G and G^v. Each copy of each $w \in V_{cu}(P)$ now continues broadcasting until the end of P or P^v as appropriate. Recall that Q takes the same time as P_u to reach $V_{cu}(P)$. Since P^v takes $\lceil \log_2 v(G) \rceil$ time steps, and Q and P_o together also use this many time steps, all nodes in $\mathscr{G}$ are officially informed by the required time. This concludes the proof of the theorem. □

We collect some data about the above construction. Here we let $d = \deg v$, and for the given center node set, δ_n and δ_c denote the minimal degree of a non-center and center node respectively.

Proposition 3.4. *The following statements hold for the generalized compounding method.*

$$v(\mathscr{G}) = v(G)v(H) - i,$$

$$e(\mathcal{G}) = \begin{cases} v(H)e(G) + i \cdot d(d-3)/2 + \sum_{s \in S} e(H_s) & \text{if } S \neq V(G), \\ v(H)e(G) + i \cdot d(d-3)/2 + \sum_{s \in V(G)\setminus\{v\}} e(H_s) + e(H^*) & \text{if } S = V(G), \end{cases}$$

$$|\mathcal{S}| = \begin{cases} \sum_{s \in S} |T_s| & \text{if } S \neq V(G), \\ \sum_{s \in V(G)\setminus\{v\}} |T_s| + |U| & \text{if } S = V(G), \end{cases}$$

$$\delta_n(\mathcal{G}) \leqslant \begin{cases} \min_{s \in S} \{d, \delta_n(H_s) + \delta_c(G)\} & \text{if } S \neq V(G), \\ \delta_c(G) + \delta_n(H^*) & \text{if } S = V(G) \text{ and } U \neq V(H^*), \\ v(\mathcal{G}) - v(H) + \delta_n(H_s) & \text{if } S = V(G), U = V(H^*) \text{ and } T_s \neq V(H_s), \end{cases}$$

$$\delta_c(\mathcal{G}) \leqslant \begin{cases} \min_{s \in S} \{\delta_c(G) + \delta_c(H_s)\} & \text{if } S \neq V(G), \\ \delta_c(G) + \delta_c(H^*) & \text{if } S = V(G). \end{cases}$$

Proof. Only the entries for δ_n and δ_c require explanation, the first three being exact.

If $S \neq V(G)$ then v is a non-center node of G which remains a non-center node in $\mathcal{G}$. Furthermore v has no new neighbors in $\mathcal{G}$ and so $\delta_n(\mathcal{G}) \leqslant d$. Now fix $s \in S$ so that $\deg_G s = \delta_c(G)$ and let u be a non-center node of H_s. The copy of s at u is a non-center node of $\mathcal{G}$ of degree at most $\delta_n(H_s) + \delta_c(G)$.

If $S = V(G)$ and $U \neq V(H^*)$ then let u be a node of H^* of degree at most $\delta_n(H^*)$ and let s be a node of G of degree at most $\delta_c(G)$. The copy of s at u is a non-center node of $\mathcal{G}$ of degree at most $\delta_c(G) + \delta_n(H^*)$.

If $S = V(G)$ and $U = V(H^*)$ then the only possible non-center nodes belong to some H_s. A non-center node of H_s of degree at most $\delta_n(H_s)$ has degree at most $V(\mathcal{G}) - v(H) + \delta_n(H_s)$ in $\mathcal{G}$.

For δ_c the proofs are similar. If $S \neq V(G)$ then a center node w of G of degree $\delta_c(G)$ is not deleted in $\mathcal{G}$. If u is a center node of some H_s then the copy of t at u has degree at most $\delta_c(G) + \delta_c(H_s)$. On the other hand, if $S = V(G)$ then w may be deleted in the construction of $\mathcal{G}$, since w may in fact be the distinguished node v. In this case, if u is a node of H^* of minimal degree then the copy of w at u is a center node of degree at most $\delta_c(G) + \delta_c(H^*)$. $\square$

In the case where both $v(G) = 2^a$ and $v(H) = 2^b$, Theorem 3.3 yields a center node set $\mathcal{S}$ equal to the entire set of nodes $V(\mathcal{G})$. If $v(\mathcal{G}) \neq 2^{a+b}$ then this is most likely not optimal (see Proposition 2.9), and yet this situation may be inevitable. For example, if $n = 2^m - 1$ then suppose that $n = xy - i$ and (x, y) satisfies the broadcast condition. Let $a = \lceil \log_2 x \rceil, b = \lceil \log_2 y \rceil$. Then $a + b = m$ and $xy - i = 2^m - 1$. The only solution is $x = 2^a, y = 2^b$ and $i = 1$, since if $x \leqslant 2^a - 1$ or $y \leqslant 2^b - 1$ then $xy \leqslant 2^m - 2$. Also, there is no solution with $i = 0$ and so the ordinary compounding method never generates a broadcast graph with $2^m - 1$ nodes.

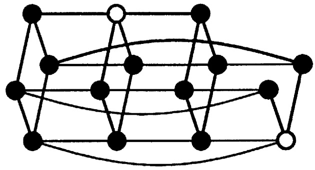

Fig. 7. Example for Proposition 3.5.

There is a slightly better method we can use in this special case, which reduces the center node sets obtained. Fix $a, b > 0$ with $a + b = m$. Proceed exactly as in the doubling procedure with G and H being broadcast graphs with 2^a and 2^b nodes respectively, except that G^v may be replaced by any broadcast graph G' on $2^a - 1$ nodes. Call the resulting graph $\mathscr{G}$. Let U be a center node set for H^* and W a center node set for G'.

Proposition 3.5. *In the notation of the above paragraph, the graph $\mathscr{G}$ is a broadcast graph and $\mathscr{S} = V(\mathscr{G}) \setminus [(V(H^*) \setminus U) \cup (V(G') \setminus W)]$ is a center node set for $\mathscr{G}$.*

Proof. Let $u \in V(\mathscr{G})$. If u is a node of some copy G_i then broadcast within G_i, taking a time steps. Then each node in G_i except possibly u has an official message. This is because u may be in $V(H^*) \setminus U$, but all other nodes of G_i are center nodes. Then broadcast horizontally in each copy of H, ensuring this is done officially in H^* with respect to U. This takes b time steps.

If u is a node of G' then first broadcast horizontally, which is possible since every node of G' also belongs to some copy of H. This takes b time steps. Every node receiving this message is a center node and so all nodes except possibly u (in the case where $u \in V(G') \setminus W$) now have the official message. Now broadcast vertically in the remaining a time steps, ensuring that this is done officially in G' with respect to W. □

Fig. 7 illustrates this procedure. Here G and H are both the hypercube Q_2, while G' and H^* are both the mbg P_2 with 3 nodes. The white nodes are those which can be made non-center nodes according to Proposition 3.5.

Proposition 3.6. *The following statements hold for the construction of Proposition* 3.5.

$$v(\mathscr{G}) = 2^m - i,$$
$$e(\mathscr{G}) = m2^{m-1} + i \cdot e(G') + e(H^*) - (i \cdot a2^{a-1} + b2^{b-1}),$$
$$|\mathscr{S}| = 2^m - i - [(v(H^*) - |U|) + (v(G') - |W|)],$$
$$\delta_n(\mathscr{G}) \leqslant \min\{a + d_n(H^*), b + d_n(G'), m\},$$
$$\delta_c(\mathscr{G}) \leqslant \min\{a + d_c(H^*), b + d_c(G'), m\}.$$

4. Iterating the compounding methods

In order to generate broadcast graphs with arbitrarily large number of nodes using the compound methods discussed above, it is necessary to iterate. Starting with a base table of broadcast graphs and center node sets, one can generate new broadcast graphs, each with a corresponding center node set, add these to the base table and repeat.

There will be gaps in the output for the ordinary compound method. For example, graphs whose order is prime or of the form $2^m - 1$ cannot be generated. However, the doubling procedure generates graphs of all orders as long as the initial list contains the broadcast graphs with 1 and 2 nodes. This last observation follows from the fact that the broadcast condition is always satisfied by $(2,n)$ and $(2,n,1)$ for any positive integer n.

It is natural to suspect that the smaller center node sets given by Theorems 3.1 and 3.3 would lead to the broadcast graphs generated at later stages having fewer edges than if the much larger center node sets given by Theorem 2.2 or Theorem 2.1 were used. Somewhat surprisingly, this is not the case, at least for the ordinary compounding method, as the following "associativity" result shows.

Say that two broadcast graphs generated by the same iterated compounding method are equivalent if they have the same number of nodes and the same number of edges.

Proposition 4.1. *Any broadcast graph eventually output from an initial list by iterating the ordinary compounding method with center node reduction as in Theorem* 3.1 *is equivalent to a compound of some broadcast graph G into some broadcast graph H, where G is in the initial list and H has been produced by the iterative method.*

Proof. We claim first that if G, H and K are broadcast graphs and if S and T are center node sets for G and H, respectively, then $(G_S[H])_{S\times T}[K]$ and $G_S[H_T[K]]$ are equivalent. This is a straightforward computation using the data in Proposition 3.2. The common vertex number is $v(G)v(H)v(K)$ and the common edge value $v(H)v(K)e(G) + |S|e(H)v(K) + |S||T|e(K)$. Note that if $(v(G), v(H))$ and $(v(G)v(H), v(K))$ satisfy the broadcast condition then it follows from the definition that $(v(H), v(K))$ and $(v(G), v(H)v(K))$ also do. A simple induction on the number of factors in a compound now proves the result. $\square$

In the formula in Proposition 3.2 for $e(\mathscr{G})$ the size of a center node set for H does not appear. It follows from Proposition 4.1, by induction on the number of factors in a compound, that center node reduction of Theorem 3.1 does not improve the upper bounds on $B(n)$ generated by the iterated ordinary compounding method. However using center node reduction does result in the generation of many more non-isomorphic equivalent broadcast graphs. This improves the chances of finding broadcast graphs with "good" (for example, symmetric) broadcast protocols.

Empirical results suggest that the same phenomenon occurs with the doubling procedure, though we do not have a general proof of this fact. Proposition 4.1 can be generalized somewhat. We need only assume that both pairs (G,H) and (H,K) are combined as in the ordinary compound method. The general case seems much harder.

5. Computational results

In this section we discuss our implementation of the iterated compounding methods. The first subsection deals with the algorithms themselves, the second with the initial data used and the third with the empirical results obtained.

5.1. The algorithms

We have implemented two iterative algorithms as described in the last section. One uses ordinary compounding with the center node reduction as described in Theorem 3.1. The other combines the generalized compounding method as in Theorem 3.3, the special case center node reduction of Theorem 3.5 and the standard method of vertex deletion. The code for these programs is available from the first author on request.

It is infeasible to store and use all information about each known broadcast graph. We use the data format of Table 1. Each row corresponds to a fixed broadcast graph G with center node set $S(G)$. We allow the possibility of multiple rows for each $v(G)$ to take advantage of the structure of different broadcast graphs and/or different center node sets.

Here $d_n(G)$ and $d_c(G)$ denote upper bounds for, respectively, $\delta_n(G)$ and $\delta_c(G)$, calculated with respect to $S(G)$. The column labeled 'Ref' gives a reference to where the given graph G (and its center node set) can be found. The bold entries in column $e(G)$ indicate when a broadcast graph is an mbg. In the same column, if the given graph $\mathcal{G}$ is generated by the doubling procedure, the relevant data $(v(G), v(H), i)$ or $(v(G), v(H))$ are listed. The G and H referred to also belong to the table.

Of course it is possible that better results might be obtained if more information were stored. For example, keeping the two lowest degrees of non-center nodes would improve some of the bounds in Proposition 3.4.

The first algorithm generates one new row for each ordered pair of input rows for which $(v(G), v(H))$ satisfies the broadcast criterion, using the update formula of Proposition 3.2.

For the second algorithm, rows are generated as follows. For each i with $0 \leqslant i \leqslant v(H) - 1$ the broadcast condition on $(v(G), v(H), i)$ is tested, and if satisfied a new row generated using the formulae of Proposition 3.4. Note that the expressions in those formulae are increasing functions of δ_c, δ_n and so the formulae give valid upper bounds if δ is replaced by d throughout. Note also that H^* can be any broadcast graph with $v(H) - i$ nodes and so for H^* we use each broadcast graph already available in the table with this property. We have already shown that there must already exist at least one such broadcast graph in the table.

Table 1
Standard input data

$v(G)$	$e(G)$	$\|S(G)\|$	$d_n(G)$	$d_c(G)$	Ref.
1	**0**	1		1	K_1
2	**1**	2		1	$Q_1 = K_2$
3	**2**	2	2	1	P_2
4	**4**	4		2	$Q_2 = (2,2)$
5	**5**	2	2	2	C_5
6	**6**	3	2	2	C_6
7	**8**	5	3	2	[25]
8	**12**	8		3	$Q_3 = (2,4)$
9	**10**	3	2	2	[25]
10	**12**	4	2	3	Fig. 2
11	**13**	5	2	2	[2]
12	**15**	6	2	3	(6,2)
13	**18**	7	3	2	[2]
14	**21**	7	3	3	Fig. 11
15	**24**	12	4	3	Fig. 1
16	**32**	16		4	$Q_4 = (2,8)$
17	**22**	5	2	3	[25]
18	**23**	6	2	3	(9,2)
19	**25**	7	2	2	[21]
20	**26**	8	2	3	[21]
21	**28**	11	2	3	Fig. 12
22	**31**	10	2	3	(11,2)
23	34	11	2	3	Fig. 13
24	36	11	3	3	[2]
25	40	15	3	3	[25]
26	**42**	15	3	3	[22]
27	**44**	15	3	3	Fig. 8
28	**48**	15	3	3	[22]
29	**52**	21	4	3	Fig. 8
30	**60**	15	4	4	Fig. 11
31	**65**	25	5	4	[25]
32	**80**	32		5	$Q_5 = (2,16)$
37	56	13	2	4	Fig. 13
39	59	16	2	4	Fig. 13
43	70	21	2	3	Fig. 13
49	94	30	3	4	[25]
57	126	28	4	4	Fig. 12
58	**121**	34	4	4	[22]
59	**124**	33	4	4	Fig. 9
60	**130**	30	4	4	Fig. 9
61	**136**	43	4	4	Fig. 10
62	**155**	31	5	5	[7]
63	**162**	54	6	5	[25]
64	**192**	64		6	$Q_6 = (2,32)$
$2^m - 2$	$(m-1)(2^{m-1}-1)$	$2^{m-1}-1$	m	m	[7]

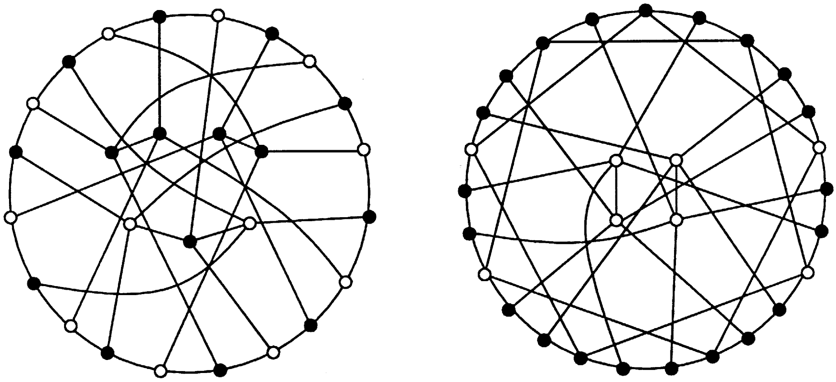

Fig. 8. Mbg's with 27 and 29 nodes (see [22]).

In addition, we generate rows by using Theorem 3.5 and the update formulae of Proposition 3.6 when and only when the resulting broadcast graph has size $2^m - 1$.

Again it is infeasible to store and use all broadcast graphs generated in each iteration. Since we are only concerned with the numbers of edges and center nodes it is reasonable to define a partial ordering $\preccurlyeq$ on the rows so that row $G_1 \preccurlyeq$ row G_2 if and only if $v(G_1) = v(G_2)$, $e(G_1) \leqslant e(G_2)$ and $c(G_1) \leqslant c(G_2)$. In this situation we say that G_1 *dominates* G_2. After a given iteration only the minimal rows with respect to this order are kept for the next iteration.

For a preassigned positive integer M, the iteration in both algorithms stops when all rows with $v(G) \leqslant M$ show no change in two successive iterations. As a final step the second algorithm then tests each entry of the table to see whether vertex deletion improves the number of edges and if so generates a new row in this way.

5.2. The initial data

We used the broadcast graph data listed in Table 1, with the addition of a few other mutually non-dominant broadcast graphs that are not mbg's (available by request). We now establish the validity of the entries in the initial table for which center node sets have not been given by other authors. As mentioned above the calculation of small center node sets is a nontrivial procedure, which has led to errors in the literature. Results such as Proposition 4.1 underline the importance of establishing good bounds on $cn(G)$ for the initial input.

The entries with 26–29 and 58–61 nodes are the mbg's discovered by Saclé [22]. In that paper center node sets are given for the cases 26, 28 and 60. We treat the remaining cases here.

27: A vertex cover of size 15 with all nodes of degree at most 4 is shown in Fig. 8.

29: The graph is displayed in Fig. 8. By Proposition 2.8 all 16 nodes of degree 3 must lie in every center node set. Estimating the size of a broadcast tree shows that every vertex of degree 4 must send to another of degree 4 in every broadcast protocol.

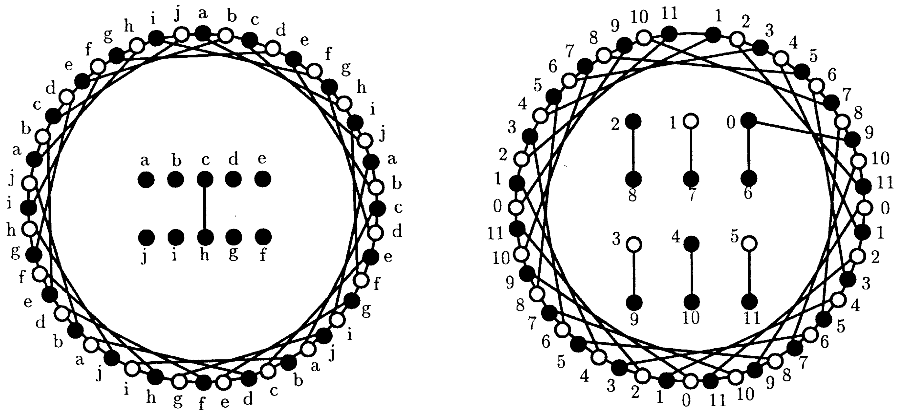

Fig. 9. Mbg's with 58 and 59 nodes (see [22]).

If a vertex of degree 4 has a neighbor of degree 4 which is not a center node then we obtain a contradiction since the maximum size of an official forest obtainable in such a situation is 24. Thus it is necessary to add to our candidate center node set a minimal vertex cover of the subgraph induced by the degree 4 nodes. Since these form a cycle of order 9 such a cover has size 5. We claim that the resulting set of size 21 is an ocns for G. If the originator is a non-center node of degree 4 then since all its neighbors are center nodes of degree at most 4 the result is clear. Now suppose that the originator has degree 5. Each such vertex v has one neighbor w of degree 5 and one of degree 3. In every broadcast protocol for v, either v informs a degree 3 node at time 1, or else v first informs w and then both v and w inform a degree 3 node at time 2. In either case there are enough degree 3 nodes idle to inform both v and w officially. All other nodes receive only the official message and so this case is completed.

58: The maximum degree is 5 and a solid 1-cover of size 34 is shown in Fig. 9. For clarity not all edges have been drawn; each labeled node inside the main circle is connected to all nodes on the circle which have the same label.

59: The maximum degree is 5 and a solid 1-cover of size 33 is shown in Fig. 9. The same convention is followed as in the previous case.

61: This is similar to case 29. The graph is displayed in Fig. 10, with the same convention as in the previous two cases. By Proposition 2.8 all nodes of degree 4 must lie in every center node set. Adding the vertex of degree 6 labeled '*', we obtain a set of size 43 which the following observations show is a center node set.

If the originator v has degree 5 then estimating the size of a broadcast tree shows that its first message must go to a w of node of degree 6 and its second to a node x of degree 4. Then x is idle at time at most 6 and so can inform v officially. If w is the node labeled '*' then there is nothing left to prove. Otherwise w sends to a degree 4 node $y \neq x$ at time 2, and this is then idle at time at most 6 and can inform w officially.

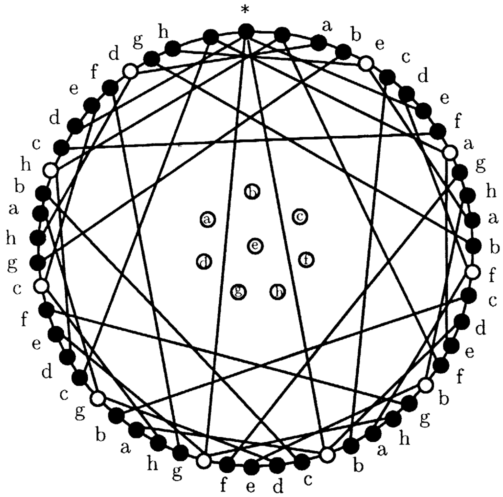

Fig. 10. Mbg with 61 nodes (see [22]).

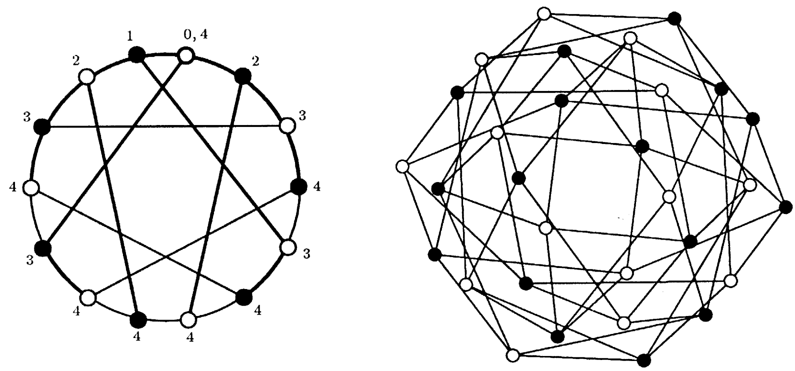

Fig. 11. Bipartite vertex-transitive mbg's with 14 and 30 nodes (see [7]).

If the originator is a non-center node of degree 6 then it must send to a degree 4 node at time at most 2 and the result follows as above.

The mbg's with $2^m - 2$ nodes presented in [7] are regular bipartite and of degree $m - 1$, with either half of the bipartition yielding a vertex cover. Two examples are found in Fig. 11. The mbg with 21 nodes and 28 edges of [18] has maximum degree 4 and so a solid 1-cover for this graph is shown in Fig. 12. The broadcast graph with 57 nodes and 126 edges from [3] has a solid 1-cover with all nodes of degree at most 5 and is also shown in Fig. 12.

The broadcast graphs discovered by Ridwan [21] are shown in Fig. 13 along with the center node sets given in that paper.

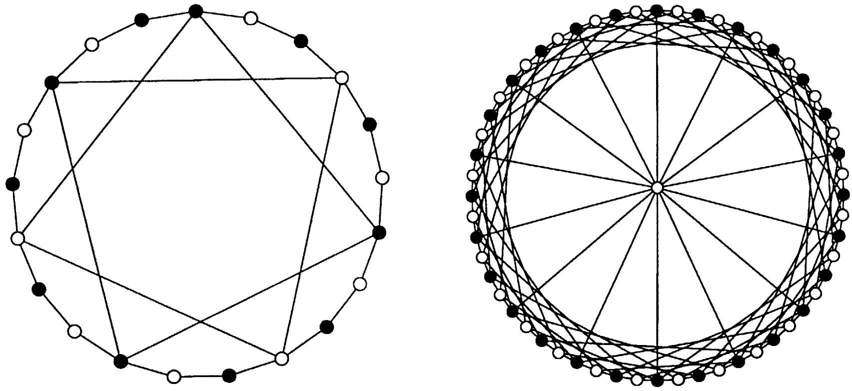

Fig. 12. Broadcast graphs with 21 and 57 nodes (see [3,18]).

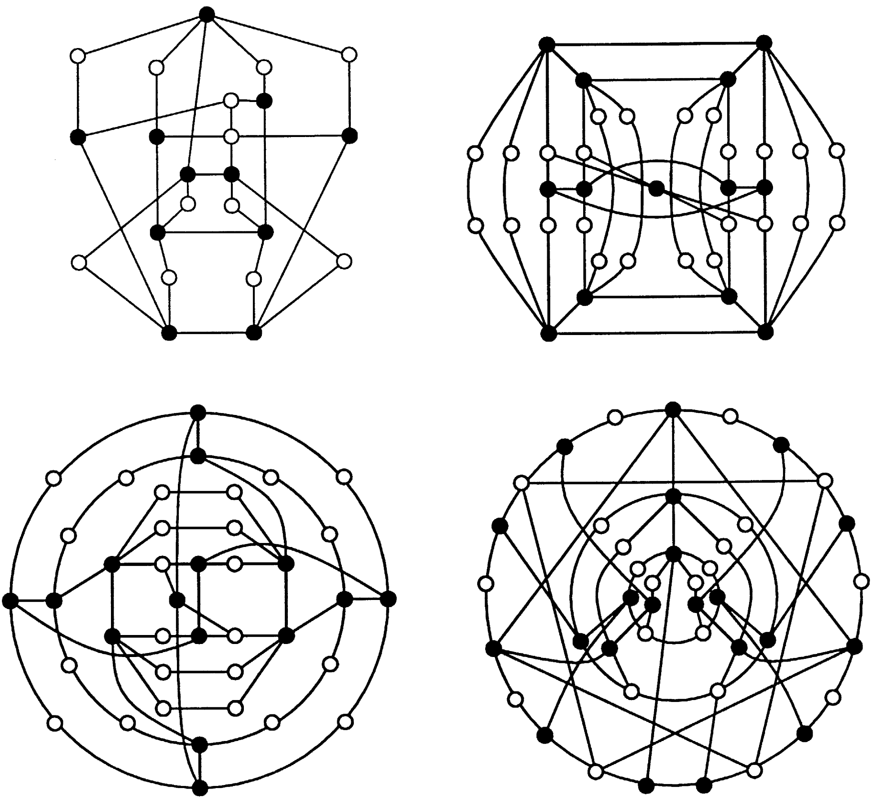

Fig. 13. Broadcast graphs with 23, 37, 39 and 43 nodes (see [21]).

5.3. Empirical results

Programs implementing each algorithm described above were run subject to the added heuristic (for memory reasons) that after each iteration, at most 10 mutually nondominant rows were used in the next iteration, for each fixed number n of nodes. The larger number of combinations which must be considered in the doubling procedure was reflected in the fact the running time for the first algorithm was only a few minutes as opposed to several hours for the second. Computations in this paper were carried out using the University of Auckland's SGI Power Challenge GR computer (with 16 R10000 processors).

In the range $n \leqslant 16\,384$, the first (ordinary compounding) algorithm generated upper bounds for $B(n)$ for 1618 of the values of n. For these values of n the second algorithm lowered those bounds 75.3% of the time, with a mean improvement of 9.42%.

In the same range the center node reduction techniques in this paper achieved a mean reduction of around 50% in the upper bounds for $cn(G)$ formerly produced by the compounding methods.

In [22] lower bounds for $B(n)$ were presented, where n has the form $2^m-3, 2^m-4$, 2^m-5 or 2^m-6. For these values of n less than 16 384 the best upper bounds of this paper exceed the lower bounds by an average of less than 6%. Most notably it follows from Saclé's lower bound and our upper bound that $314 \leqslant B(122) \leqslant 315$.

The best upper bounds for $B(n)$ generated by various methods, for certain ranges of n, are displayed in Table 2. These sample values of n were established by Gargano and Vaccaro [11]. The second column (WV or Ref) indicates the previous known bound before Bermond et al. [2]. The default reference WV is [25]. We note that the formula for the number of edges given, when $n = 2^m - 2^k - j$, by the construction in Theorem 3 of [15] is incorrect. We used the formula $S(n) = n(m-k+1) - \frac{1}{2}(m-k)(3m+k-1)+j$ to generate 5 of the entries labelled [15] in Table 2. The BFP and $DVWZ$ columns indicate the best known bounds obtained in this current paper, by our first (ordinary compounding) and second algorithms, respectively. Only the cases $n = 513, 896, 1008, 16128$ were presented in [2], and the corresponding bounds on $B(n)$ were the same as in this paper. We note that the justification for the BFP entry given in [2] for the $n = 1008$ row was incorrect, being based on the assumption that Labahn's mbg G with 63 nodes has a solid 1-cover of size 36. In fact $cn(G) = 54$ by Proposition 2.8 and this gives a value of 4320 for the particular compound used. However, one may instead take G to be the regular mbg with 126 nodes and H to be the hypercube Q_3. This yields $B(1008) \leqslant 3780$, the entry shown in Table 2.

Many authors have published tables of the best known upper bounds for $B(n)$ for $1 \leqslant n \leqslant 64$. Table 1 shows the best known upper bounds for $B(n)$ for $17 \leqslant n \leqslant 127$. The (exact) bounds for $n \leqslant 16$ (which are listed in Table 3) are credited to Farley et al. [8,9].

Each entry in the column labeled 'Compound' gives a construction that matches the best known bound. In every case except for $n = 114$, the G and H in the entry $(v(G), v(H), i)$ also occur in Table 3. In this special case the algorithm uses the graph

Table 2
Comparison of recent methods for bounding $B(n), n = 2^9$–2^{10} and $n = 2^{13}$–2^{14}

n	VW or Ref	BFP	$DVWZ$
513	1 026	948	934
514	1 028		935
515	1 030		936
516	1 032	999	937
517	1 034		938
520	1 040	976	941
521	1 042		942
528	1 056	1020	959
529	1 058	1156	960
544	1 088	1104	1013
545	1 090		1014
576	1 152	1140	1128
577	1 410		1129
640	1 536	1472	1348
641	1 729		1349
768	2 112 [15,24]	2032	1932
769	2 690		1932
896	2 688 [15]	2688	2678
897	3 264		2676
960	3 120 [15]	3040	3040
961	4 040	3640	3479
992	3 472 [15]	3472	3472
993	4 616		3900
1008	3 780 [15]	3780	3780
1009	4 889		4197
1016	4 064 [15]	4064	4064
1017	5 017		4466
1020	4 335 [15]	4335	4335
1021	5 078		4626
1022	**4 599** [7,15]	—	—
1023	5 106		5082
8193	20 482		17 083
8194	20 484		17 084
8195	20 486		17 085
8196	20 488		17 086
8197	20 490		17 087
8200	20 496		17 090
8201	20 498		17 091
8208	20 512	17 856	17 098
8209	20 514		17 099
8224	20 544		17 114
8225	20 546		17 115
8256	20 608	20 016	17 146
8257	20 610		17 147
8320	20 736	19 200	17 406
8321	20 738		17 407
8448	20 992	20 144	17 968
8449	20 994		17 969

Table 2

n	VW or Ref	BFP	$DVWZ$
8704	21 504	22 784	19 260
8705	21 506		19 261
9216	22 528	22 464	21 176
9217	29 698		21 177
10 240	32 715 [15]	31 744	24 541
10 241	32 717 [15]		24 542
12 288	36 811 [15]	43 008	35 637
12 289	51 125 [15]		35 637
14 336	57 266 [15]	57 344	49 748
14 337	68 610		49 746
15 360	65 280 [15]	64 000	61 584
15 361	89 093 [11]		61 579
15 872	71 424 [15]	71 424	71 424
15 873	99 334 [11]		71 469
16 128	76 608 [15]	76 608	76 608
16 129	108 807 [11]		83 464
16 256	81 280 [15]	81 280	81 280
16 257	111 883		87 939
16 320	85 680 [15]	85 680	85 680
16 321	113 356		92 169
16 352	89 936 [15]	89 936	89 936
16 353	114 060		96 294
16 368	94 116 [15]	94 116	94 116
16 369	114 397		100 379
16 376	98 256 [15]	98 256	98 256
16 377	114 557		104 452
16 380	102 375 [15]	102 375	102 375
16 381	114 648 [8]		106 548
16 382	**106 483** [7,15]	—	—
16 383	114 674 [8]		114 626

G with 57 nodes from Table 1, which does not yield the best bound on $B(57)$. This shows the usefulness of keeping several mutually nondominant broadcast graphs, rather than taking just the one with the fewest edges, and underlines the importance of finding small center node sets.

New bounds for $B(n)$ are indicated in each row of Table 3 where we omit a reference. Again bold entries in the table denote optimal values. We note that the entry 111 for $n = 55$ was claimed in [25] but cannot be obtained via the doubling procedure with the data listed in that paper.

6. Comments

The most difficult case for any of the methods above, as far as generating an upper bound for $B(n)$ is concerned, is the case where $n = 2^m - 1$. Despite our refinements the results in this paper are rather far from the conjectured values of $B(n)$. It seems necessary to attack this case directly.

Table 3
Best known upper bounds on $B(n), 17 \leqslant n \leqslant 127$

n	Bound	Ref	Compound	n	Bound	Ref	Compound	n	Bound	Ref	Compound
17	**22**	[19]		54	103		(27,2)	91	179		(23,4,1)
18	**23**	[3]	(9,2)	55	111		(28,2,1)	92	180		(23,4)
19	**25**	[3]		56	111	[22]	(28,2)	93	188		(24,4,3)
20	**26**	[18]		57	123		(58,1,1)	94	188		(24,4,2)
21	**28**	[18]		58	**121**	[22]		95	188		(24,4,1)
22	**31**	[18]	(11,2)	59	**124**	[22]		96	188		(24,4)
23	34	[18]		60	**130**	[22]		97	203		(14,7,1)
24	36	[3]		61	**136**	[22]		98	203		(14,7)
25	40	[3]		62	**155**	[7,15]		99	220		(25,4,1)
26	**42**	[22]		63	**162**	[16]		100	220		(25,4)
27	**44**	[22]		64	**192**	[9]	(2,32)	101	228		(13,8,3)
28	**48**	[22]		65	101		(5,13)	102	228		(13,8,2)
29	**52**	[22]		66	105		(6,11)	103	228		(13,8,1)
30	**60**	[3]		67	107		(17,4,1)	104	228		(13,8)
31	**65**	[3]		68	108		(17,4)	105	236		(27,4,3)
32	**80**	[9]	(2,16)	69	111		(5,14,1)	106	236		(27,4,2)
33	48	[25]	(3,11)	70	112		(5,14)	107	236		(27,4,1)
34	49	[25]	(17,2)	71	115		(9,8,1)	108	236		(27,4)
35	51	[2]	(5,7)	72	116		(9,8)	109	252		(14,8,3)
36	52	[2]	(9,4)	73	121		(5,15,2)	110	252		(14,8,2)
37	56	[21]		74	122		(5,15,1)	111	252		(14,8,1)
38	57	[3]	(19,2)	75	123		(5,15)	112	252		(14,8)
39	59	[21]		76	128		(19,4)	113	281		(2,59,5)
40	60	[3]	(20,2)	77	131		(6,13,1)	114	280		(57,2)
41	65	[2]	(6,7,1)	78	132		(6,13)	115	278		(58,2,1)
42	66	[2]	(6,7)	79	135		(20,4,1)	116	276		(58,2)
43	70	[21]		80	136		(20,4)	117	283		(59,2,1)
44	72	[2]	(11,4)	81	142		(3,27)	118	281		(59,2)
45	78		(23,2,1)	82	145		(6,14,2)	119	292		(60,2,1)
46	79		(23,2)	83	146		(6,14,1)	120	290	[22]	(60,2)
47	83	[2]	(24,2,1)	84	147		(6,14)	121	317		(61,2,1)
48	83	[2]	(24,2)	85	157		(6,15,5)	122	315		(61,2)
49	94	[25]		86	158		(6,15,4)	123	346		(62,2,1)
50	95	[2]	(25,2)	87	159		(6,15,3)	124	341		(62,2)
51	99		(26,2,1)	88	160		(6,15,2)	125	379		(2,63,1)
52	99	[22]	(26,2)	89	161		(6,15,1)	126	378	[7,15]	
53	103		(27,2,1)	90	162		(6,15)	127	417		(2,64,1)

The widely believed conjecture that $B(n) \leqslant B(n+1)$ if $n \neq 2^m$ is still unproven.

As mentioned after Proposition 4.1, it is unknown whether the iterated doubling procedure can give better upper bounds on $B(n)$ when center node reduction is used.

Acknowledgements

We thank Art Liestman for providing his current table of best known values of $B(n)$, $n \leqslant 64$. The thoughtful comments of the referees helped us to improve the paper. The second author was partially supported by the National Science Foundation under Grant

DDM 90-57066. The third author was supported by a NZST Postdoctoral Fellowship. The fourth author was supported by FRST grant UOA403.

References

[1] M. Barborak, M. Malek, A. Dahbura, The consensus problem for fault-tolerant computing, ACM Comput. Surv. 25 (1993) 171–220.

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

[3] J.-C. Bermond, P. Hell, A.L. Liestman, J.G. Peters, Sparse broadcast graphs, Discrete Appl. Math. 36 (1992) 97–130.

[4] J.-C. Bermond, P. Hell, A.L. Liestman, J.G. Peters, Broadcasting in bounded degree graphs, SIAM J. Discrete Math. 5 (1992) 10–24.

[5] S.C. Chau, A.L. Liestman, Constructing minimal broadcast networks, J. Combin. Inform. System. Sci. 10 (1985) 110–122.

[6] X. Chen, An upper bound for the broadcast function $B(n)$, Chinese J. Comput. 13 (1990) 605–611.

[7] M.J. Dinneen, M.R. Fellows, V. Faber, Algebraic constructions of efficient broadcast networks, in: Applied Algebra, Algebraic Algorithms and Error-Correcting Codes 9, Lecture Notes in Computer Science, vol. 539, Springer, Berlin, 1991, pp. 152–158.

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

[9] A.M. Farley, S.T. Hedetniemi, A. Proskurowski, S. Mitchell, Minimum broadcast graphs, Discrete Math. 25 (1979) 189–193.

[10] P. Fraigniaud, E. Lazard, Methods and problems of communication in usual networks, Discrete Appl. Math. 53 (1994) 79–133.

[11] L. Gargano, U. Vaccaro, On the construction of minimal broadcast networks, Networks 19 (1989) 673–689.

[12] M. Grigni, D. Peleg, Tight bounds on minimum broadcast networks, SIAM J. Discrete Math. 4 (1991) 207–222.

[13] S.M. Hedetniemi, S.T. Hedetniemi, A.L. Liestman, A survey of gossiping and broadcasting in communication networks, Networks 18 (1988) 319–349.

[14] J. Hromkovič, R. Klasing, B. Monien, R. Peine, Dissemination of information in interconnection networks (broadcasting and gossiping), in: Combinatorial Network Theory, Applied Optimization, vol. 1, Kluwer, Dordrecht, 1996, pp. 125–212.

[15] L.H. Khachatrian, H.S. Haroutunian, Construction of new classes of minimal broadcast networks, in: Proceedings 3rd International Colloquium on Coding Theory, Armenia, 1990, pp. 69–77.

[16] R. Labahn, A minimum broadcast graph on 63 vertices, Discrete Appl. Math. 53 (1994) 247–250.

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

[18] M. Maheo, J.-F. Saclé, Some minimum broadcast graphs, Discrete Appl. Math. 53 (1994) 275–285.

[19] S. Mitchell, S. Hedetniemi, A census of minimum broadcast graphs, J. Combin. Inform. System Sci. 5 (1980) 141–151.

[20] A. Pelc, Fault-tolerant broadcasting and gossiping in communication networks, Networks 28 (1996) 143–156.

[21] Ridwan, Construction of minimal broadcast networks, Master's thesis, Pennsylvania State University, August 1995.

[22] J.-F. Saclé, Lower bounds for the size in four families of minimum broadcast graphs, Discrete Math. 150 (1996) 359–369.

[23] P.J. Slater, E. Cockayne, S.T. Hedetniemi, Information dissemination in trees, SIAM J. Comput. 10 (1981) 692–701.

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

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