

Minimal Contention-free Matrices with Application to Multicasting

Johanne Cohen, Pierre Fraigniaud, and Margarida Mitjana

ABSTRACT. In this paper, we show that the multicast problem in trees can be expressed in term of arranging rows and columns of boolean matrices. Given a $p \times q$ matrix M with 0-1 entries, the *shadow* of M is defined as a boolean vector x of q entries such that $x_i = 0$ if and only if there is no 1-entry in the i th column of M , and $x_i = 1$ otherwise. (The shadow x can also be seen as the binary expression of the integer $x = \sum_{i=1}^q x_i 2^{q-i}$. Similarly, every row of M can be seen as the binary expression of an integer.) According to this formalism, the key for solving a multicast problem in trees is shown to be the following. Given a $p \times q$ matrix M with 0-1 entries, finding a matrix M^* such that:

1. M^* has at most one 1-entry per column;
2. every row r of M^* (viewed as the binary expression of an integer) is larger than the corresponding row r of M , $1 \leq r \leq p$; and
3. the shadow of M^* (viewed as an integer) is minimum.

We show that there is an $O(q(p+q))$ algorithm that returns M^* for any $p \times q$ boolean matrix M .

The application of this result is the following: Given a *directed* tree T whose arcs are oriented from the root toward the leaves, and a subset of nodes D , there exists a polynomial-time algorithm that computes an optimal multicast protocol from the root to all nodes of D in the all-port line model.

1. Introduction

1.1. Motivations. Recent advances in telecommunication systems enhanced standard point-to-point communication protocols to multi-point protocols. These latter protocols are of particular interest for group applications. Those groups involve more than two users (some may even involve thousands of users) sharing a common application, as video-conferences, distributed data-bases, media-spaces, games, etc. Several protocols have been proposed to handle and to control a large group of users. We refer to [DDC97, MS98] for surveys on multi-point applications and protocols. Solutions differ according to the type of traffic that is induced by the shared application, and according to the quality of service required by the users. Multi-point architectures are often based on tree-networks [Win87], either a single tree connecting all the users (*e.g.*, Core-Based Tree [BFC93]), or several

The first author is supported by the AFFDU and the Editions Gauthier-Villars. Additional support by the DRET of the DGA. A preliminary version of this work was announced in [CFM99].

trees (e.g., PIM [DEF⁺94]). The traffic between the users is then routed along the edges of the tree(s).

One of the major communication problem related to multi-point applications consists to broadcast a message from one user to all the users of the application. This operation is called *broadcast* at the application level, though it is actually a *multicast* at the network level. The repetition of point-to-point connections between the source and the several destinations would significantly increase the traffic in the network, and it makes this solution not applicable in practice [DDC97]. Thence, the source must require the help of other nodes to relay messages. A broadcast message will then reach the destinations after having been relayed by several intermediate nodes (each intermediate node may possibly get one copy of the message if it belongs to the group). In order to preserve the broadcast application from transmission errors, and to bound the interval between successive receptions of consecutive packets, the number of hops between the source and each destination must be as small as possible.

The aim of this paper is to provide a polynomial algorithm which, for any tree T , and for any source $u \in V(T)$, returns a multicast protocol from u to an arbitrary subset of nodes of T that minimizes the number of hops under the all-port line model. Actually, we consider multicasting from the root to a set of destination nodes in a *directed* tree T whose arcs are oriented from the root toward the leaves. We focus our work on oriented trees because, although a bidirectional channel can be reserved between members of a group to facilitate bidirectional exchanges, it frequently happens that the bandwidth reserved in each direction differs from each other as the application is often not symmetric. For instance, consider members connected to a video server: the main point is to insure a fast broadcast of the multi-media traffic *from* the server, and thus the bandwidth of the connections from or toward the server may differ of a few order of magnitude.

1.2. Models. We will consider both 1-port and all-port models. In the 1-port model, we assume that, at any given time, each node of the tree can *call* at most one other node of the tree. In the all-port model, a node can call many other nodes simultaneously, up to one call for every of its output ports. Moreover, according to modern communication facilities (e.g., circuit-switched, wormhole, WDM, or, in some sense, ATM), long-distance calls are allowed, in the sense that the receiver of a call is not necessarily a neighboring node of the initiator of the call, and a message crossing a non-destination node can cut-through that node. This model is often called *line model* in the literature.

As a restriction though, we want the calls performed at the same time to not share any edge. This latter restriction is set to avoid contention on the links. In particular, the line model implies that, in the all-port case, a node x cannot initiate more than $\deg^+(x)$ calls, where $\deg^+(x)$ is the out-degree of node x . For instance, on Figure 1(b), the source node u cannot inform more than one other node at a time.

The set of all calls performed at the same time is called a *round*. For instance, on Figure 1(a), the first round is composed of one call, the second round is composed of two calls, and the third round is composed of four calls. We will express the cost of our broadcast protocols in terms of number of rounds. (That is we will be interested in minimizing the latency of the protocol rather than its throughput. Note that the pipeline technique may then be applied to our protocols in order to decrease the




FIGURE 1. A broadcast in the 1-port line model (a), and a multicast in the all-port line model (b). Destination nodes are colored in black.

throughput for broadcasting long messages [FL94].) The aim of this paper is to show that there exist polynomial-time algorithms that compute the multicast time of any directed tree T under the all-port line model. Comparing the two protocols on Figure 1(a) and (b) makes clear that the constraints 1-port and all-port give rise to similar types of problems. Actually, it will be shown that both problems can be solved by using a reduction to a problem on boolean matrices. However, the multicast problem can be completely solved in the all-port model using the tools introduced in this paper, whereas the 1-port version of the problem requires some more works that make its solution out of the scope of this paper.

1.3. Previous works. A huge literature has been devoted to group-communication problems under different hypotheses [DDC97, FL94, HHL86, HKMP95, MS98]. The related decision problems are often NP-complete for general networks [Mid93, SCH81], and this gave rise to several approximation algorithms [BNGNS98, KP92, Rav94] and heuristics [FV97, SW84]. Tree-networks deserved a specific interest in this context. Proskurowski [Pro81] has shown that computing the broadcast time of a tree is polynomial in the 1-port model when only neighbor-to-neighbor calls are allowed. Still in the neighbor-to-neighbor model, Slater, Cockayne and Hedetniemi [SCH81] have derived a polynomial algorithm to find the center-nodes of undirected trees, that is nodes having minimal broadcast time among all nodes of the tree. Farley and Proskurowski [FP81] have also studied the broadcast problem in undirected trees when, at the beginning of the process, more than one node know the information to broadcast. Finally, Harutyunyan and Labahn independently showed that, for any n , there exists an undirected tree-network whose broadcast time from any source is at most roughly $1.44 \lceil \log_2 n \rceil$ [Har, Lab89].

When long-distance calls are allowed, Cohen [Coh98] has shown that there exists a polynomial-time algorithm to compute an optimal *broadcast* protocol in directed trees under the all-port line model. However, although this algorithm can be extended to the *multicast* problem in which the set of destinations is a subset of the nodes of the tree, it yields an inefficient protocol. In the 1-port

line model, Farley [Far80] has shown that every *undirected* n -node network has a broadcast time of $\lceil \log_2 n \rceil$ (see also [HKUW97]). This result has been extended in [CFKR98] to the case in which the routes are chosen according to a shortest path routing function. However, the results of [CFKR98, Far80] do not hold in directed networks: take as a counter example the digraph in which a node u has a unique outgoing arc to a node v which has in turn $n - 2$ outgoing arcs to $n - 2$ vertices w_1, \dots, w_{n-2} , each connected by an outgoing arc to node u . Actually, broadcasting in a directed network gives rise to an NP-complete decision problem in the 1-port line model.

Some authors have also considered the vertex-disjoint constraint. In this context, the broadcasting problem was studied for specific architectures [HKS96, HKSH95], and approximation algorithms have been derived [KP92]. Actually, vertex-disjoint hypotheses also yield complex problems, and the broadcast problem is still open for trees (see [BET96] for a first attempt in this direction).

1.4. Our results. First, we will show that the broadcast problem in directed trees under the line model gives rise to the following matrix problem (Lemma 2.4 in Section 2). Given a $p \times q$ matrix M with p rows, q columns, and 0-1 entries, the *shadow* of M is defined as a 1-dimensional boolean vector x of q entries such that $x_i = 0$ if and only if there is no 1-entry in the i th column of M , and $x_i = 1$ otherwise. According to this formalism, the key for solving a multicast problem in directed trees is shown to be the following.

Minimal contention-free matrix problem. Given a $p \times q$ matrix M with 0-1 entries, finding a matrix M^* such that¹:

1. M^* has at most one 1-entry per column;
2. every row r of M^* is larger than the corresponding row r of M , $1 \leq r \leq q$;
and
3. the shadow of M^* is minimum.

Such matrix M^* is called a *minimal contention-free version* of M . Note that the minimal contention-free version of a matrix is not necessarily unique, even up to a permutation of the rows. On the other hand, the shadow of a minimal contention-free version of a matrix is unique.

As an example, let us consider Figure 1(a). The corresponding matrix is

$$M = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

because there are respectively 2, 4, and 1 nodes in the three branches (this correspondence will be formally established in Section 2). Since M has a single 1-entry per column, a minimal contention-free version of M is M itself, and the shadow is $7 = (111)_2$. Now, assume that the rightmost branch of the tree of Figure 1(a) contains three nodes instead of only one. Then the corresponding matrix is

$$(1.1) \quad M = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix},$$

¹Since the shadow can also be seen as the binary expression of an integer, and since, similarly, every row of M can be seen as the binary expression of an integer, the comparison of shadows and rows must be understood as comparing the corresponding integers.

and the reader can check that a minimal contention-free version of M is

$$(1.2) \quad M^* = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

M^* has a shadow equal to $14 = (1110)_2$. We will show that the matrix M^* determines a broadcast protocol from the root according to the 1-entries of the matrix. For instance, at round 1, v calls the second (middle) branch; at round 2, v calls the third (rightmost) branch; and, at round 3, v calls the first (leftmost) branch. At round 4, v is idle. We will show that there is an $O(q(p+q))$ -time algorithm that computes a minimal contention-free version of M , for any $p \times q$ boolean matrix M (Theorem 3.1 in Section 3).

Using the previous result, we will show that multicasting from the root of an arbitrary directed tree under the all-port line model can be solved in polynomial time (Corollary 4.1 in Section 4).

Let us first formalize the relationship between contention-free matrices and the broadcast problem.

2. Broadcast problems and contention-free boolean matrices

In this section, we consider the 1-port line model. Indeed, although our multicast problem is stated under the all-port line model, the 1-port model helps to understand the relationship between broadcasting on one hand, and contention-free version of matrices on the other hand. A broadcast protocol B can be described by the list of all calls performed by B . The construction of our broadcast algorithms for trees is based on the so-called *shadow* of a broadcast protocol. Let $T = (V, E)$ be any oriented tree, and let B be a broadcast protocol in T performing in r rounds.

DEFINITION 2.1. The *shadow* of B on an arc $e \in E$ is the r -dimensional vector (x_1, \dots, x_r) , $x_i \in \{0, 1\}$, such that $x_i = 1$ if and only if there is a call passing through e at round i . The *restriction* of B on a vertex $u \in V$ with d outgoing links e_1, \dots, e_d is the $d \times r$ matrix with entries in $\{0, 1\}$ such that there is a 1 at entry i, j if and only if u gives a call through link e_i at round j of B . The *shadow* of B on $u \in V$ is then the r -dimensional vector (x_1, \dots, x_r) such that $x_i = 1$ if and only if there is a 1-entry in column i of the restriction of B on u , and 0 otherwise.

The shadow of a broadcast protocol B on an arc e (resp. on a vertex u) is denoted by $\text{shad}(B, e)$ (resp. $\text{shad}(B, u)$). As shadows can be seen as binary representations of integers, we denote by $\text{bin}(B, e)$ (resp. $\text{bin}(B, u)$) the integer whose binary representation is $\text{shad}(B, e)$ (resp. $\text{shad}(B, u)$). Let B be a broadcast protocol in T performing in r rounds. For any vertex u , and for any link e , we have $\text{bin}(B, u) \leq 2^r - 1$, and $\text{bin}(B, e) \leq 2^r - 1$. The previous inequalities suggest the following definition.

DEFINITION 2.2. Let $T = (V, E)$ be any directed tree, and let B be a broadcast protocol from the root in T . Let $u \in V$, and $e \in E$. B is said *lexicographically optimal* in u (resp. in e) if $\text{bin}(B, u) \leq \text{bin}(B', u)$ (resp. $\text{bin}(B, e) \leq \text{bin}(B', e)$) for any broadcast protocol B' in T .

2.1. Broadcasting in a path. Let P_n be the path of n nodes, and let u be one extremity of the path. An optimal broadcast protocol B from u performs in $d = \lceil \log_2 n \rceil$ rounds as follows. Let us label the nodes consecutively from 0 to $n-1$,

starting at u labeled 0. If $n = 2^d$ then u calls node $n/2$ at the first round, and we are let with two simultaneous broadcasts from the extremity of a path of length 2^{d-1} . The algorithm is then defined by induction. Note that, in the case $n = 2^d$, the source u needs to call at every round so that the broadcast can complete in $\lceil \log_2 n \rceil$ rounds. In the general case, let us decompose $n - 1$ in base 2, that is $n - 1 = \sum_{i=0}^{d-1} x_i 2^i$. The $\lceil \log_2 n \rceil$ -rounds algorithm B performs as follows. Node u gives a call at round j , $j = 1, \dots, d$, if and only if $x_{d-j} = 1$. Moreover, if u does give a call at round j , then it calls node v_j labeled $n - 1 - \sum_{i=d-j}^{d-1} x_i 2^i$. Upon reception of a call from u at round j , node v_j starts a broadcast to the sub-path of P_n composed of nodes lying between node v_j and node v_k where $k = n - 2 - \sum_{i=d-j-1}^{d-1} x_i 2^i$. This sub-path is of size 2^{d-j} .

LEMMA 2.3. *The broadcast protocol B is lexicographically optimal in u .*

PROOF. When an internal node receives a call at round j , $j = 1, \dots, d$, it can inform at most $2^{d-j} - 1$ other nodes during the $d - j$ remaining rounds. Thus, any broadcast algorithm B' from u satisfies $\sum_{i=1}^d \text{shad}(B', u)_i 2^{d-i} \geq n - 1$. Since, by definition, $\sum_{i=1}^d \text{shad}(B, u)_i 2^{d-i} = n - 1$, we get $\text{bin}(B, u) \leq \text{bin}(B', u)$. \square

2.2. Broadcasting in a star. Let T be a star of p branches rooted at u , and let n_i be the number of nodes of the i th branch, $i = 1, \dots, p$. T has $n = \sum_{i=1}^p n_i + 1$ nodes in total. Assume w.l.g. that $n_1 \geq n_2 \geq \dots \geq n_p$. We denote by v_i the neighbor of u in the i th branch, and $e_i = (u, v_i)$, $i = 1, \dots, p$. Let $q = \lceil \log_2(n_1 + 1) \rceil$. A broadcast from u to T takes at least q rounds.

Let B_i be the lexicographically optimal broadcast protocol from u to the i th branch, $i = 1, \dots, p$, as defined in Section 2.1. Let M be the $p \times q$ matrix whose i th row is $\text{shad}(B_i, e_i)$. As it is defined, M is a ‘‘merging’’ of shadows, but it cannot be directly recognized as the restriction of a 1-port broadcast protocol from u to T since there might be contentions between the several shadows. For instance, if T is a star of two branches of one node each, then $\text{shad}(B_1, e_1) = \text{shad}(B_2, e_2) = [1]$, and $M = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ is not a restriction in u of a broadcast protocol since u would then have to call two nodes simultaneously, which is in contradiction with the 1-port hypothesis. However, M can be transformed in

$$M^* = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

which is the restriction of the broadcast protocol from u in T which performs as follows: at the first round u calls the node of the first branch, and, at the second round, u calls the node of the second branch. A similar example has been considered before when matrix M of Equation 1.1 was transformed into the matrix M^* of Equation 1.2.

LEMMA 2.4. *Let T be a star of p branches of length at most $2^q - 1$ nodes each, and rooted in u . Let M be the $p \times q$ matrix whose p rows are the p shadows $\text{shad}(B_i, e_i)$ of p broadcast algorithms from u to the p branches of T . Assume that all B_i 's are lexicographically optimal in u . Then any contention-free version M^* of M determines a broadcast protocol B from u , and conversely. Moreover, if M^* is minimal, then B is lexicographically optimal in u , and conversely.*

In other words, in the context of this section, there is a one-to-one correspondence between contention-free matrices and broadcast protocols, and between minimal contention-free matrices and lexicographically optimal broadcast protocols.

PROOF. Let M^* be a contention-free version of M . To show that M^* is the restriction of a broadcast protocol B from u , we give a broadcast protocol from u as a function of the structure of M^* . For every r , $1 \leq r \leq p$, the r th row of M^* is larger than the corresponding row in M . Therefore, consider a particular row L^* of M^* , and let L be the corresponding row in M . Assume both rows correspond to the r th branch of the star. If $L^* = L$ then L^* is indeed the shadow of a broadcast protocol in the r th branch. Thus assume that $L \neq L^*$, and let i be the leftmost bit position for which L and L^* differ. Note that, in this case, $L_i^* = 1$ and $L_i = 0$ because $L_i^* \geq L_i$. L^* defines a broadcast protocol in the r th branch of the star as follows. From round 1 to round $i - 1$, do as in the original broadcast protocol L . At round i , u calls its neighbor v_r in the r th branch. During the remaining rounds, u does not call the r th branch anymore. However, v_r simulates the calls of u according to L . That is, if u calls w at round $j > i$ in L , then v_r calls node w at round j . Therefore, L^* is the shadow of a broadcast protocol in the r th branch of the star. M^* has at most one 1-entry per column, thus B satisfies the 1-port model.

Conversely, given a broadcast protocol B from u in T , its restriction M^* in u satisfies that there is at most one 1-entry per column (this is because of the 1-port model). Moreover, every row of M^* is larger than the corresponding row in M because all the B_i 's are lexicographically optimal (Lemma 2.3). Therefore, M^* is a contention-free version of M .

With the same notations as before, M^* is minimal if and only if B is lexicographically optimal because $\text{shad}(M^*) = \text{shad}(B, u)$. \square

According to the previous lemma, the key to find an optimal broadcast protocol in a star is to solve the minimal contention-free matrix problem as stated in Section 1.4. Actually, we will see in Section 4 that solving the minimal contention-free matrix problem is also the key to solve the broadcast and multicast problems in any arbitrary directed tree. Therefore, the next section is entirely devoted to solving the minimal contention-free matrix problem.

3. A polynomial algorithm for the minimal contention-free boolean matrix problem

Let M be a $p \times q$ boolean matrix. Our algorithm will transform M in a $p \times q^*$ minimal contention-free version of M denoted by M^* . The total number of columns of any minimal contention-free version of M is denoted by $q^*(M)$. $q^*(M)$ and M^* will be computed by a sequence of elementary operations of two types: *insertion* of a zero-column at position 0, and *shifting* of an existing zero-column from position $t - 1$ to position t (columns are labeled from left to right). The shift operation has an important consequence on the 1-entries of the matrix. When a zero-column is shifted one position to the right, from position $t - 1$ to position t , that is when the two columns $t - 1$ and t are exchanged, the entries of the matrix are modified according to the following rule:

Rule 1.: for every i , $1 \leq i \leq p$, if there is a 1-entry originally at position t of row i , then, after the exchange, all 1-entries of row i at position $> t$ are switched to 0.

This rule comes from the simple fact that, for any k , $2^{k+1} > \sum_{i=0}^k a_i 2^i$ for any $a_i \in \{0, 1\}$, $i = 0, \dots, k$. Therefore, any row modified according to rule 1 is larger than the original row, whatever are the entries of the row left to position t .

Our algorithm is formally described in Algorithm 1, in the Appendix. An example is provided on Figure 2. Informally, Algorithm 1 performs as follows. The q columns of M are considered from left to right. Problems occur when there are two or more 1-entries in the current column (Instruction 6). On Figure 2(a), this occurs at column 4 since there is a single 1-entry in each of the three leftmost columns of M . Algorithm 1 then tries to increase the number of zero-columns by shifting existing zero-columns from their current position to the left of the current column, and applying rule 1 (Instruction 13). Possibly, one zero-column is inserted at position 0 (Instruction 18). The goal is to obtain enough zero-columns on the left of the current column to spread out the contending 1's over these zero-columns.

On Figure 2(a), there is no zero-column at the current phase of the algorithm, and thus a zero-column is inserted at position 0, as shown on Figure 2(b). Then the two first columns are exchanged. This exchange has a major consequence: according to rule 1, all 1-entries, but the leading 1, of the first row are switched to 0. This creates a new zero-column, and one of the two contending 1's of column 4 vanishes (see Figure 2(c)).

The algorithm then considers position 5 (now the 6th column from the left). Four 1-entries are contending at position 5 of the matrix. The rightmost zero-column is then shifted to the right. It is worth to notice that it is always the rightmost zero-column not next to the current column that is considered. Choosing this column instead of any zero-column has a tremendous effect on the shadow of the resulting matrix. The effect of this shift in the example is to delete one contending 1-entry (see Figure 2(d)). The zero-column is then shifted once more to the right. Again, it deletes one contending 1-entry (see Figure 2(e)). Once there are enough zero-columns to solve all conflicts between 1-entries in the current column, the contending 1's are spread out over these columns. Note that if after all possible shifting, there is still not enough zero-columns to absorb the contending 1's, then some zero-columns are inserted again (Instruction 23). In our example, there are one zero-column and two contending 1's, so there is no need to insert new zero-column (see Figure 2(e)). Now, the choice of the unique 1-entry of column 5 which is *not* moved to a zero-column matters. Algorithm 1 keeps in place the 1-entry which corresponds to the row with the minimum lexicographic order, starting from the current column (Instruction 25). In our example, it means that the 1-entry of row 5 will be let in place, while the 1-entry of row 4 will be moved to the zero-column. Indeed, from the current position, row 4 is 110 whereas row 5 is 100.

After that, we are left with the matrix on Figure 2(f) in which the last 1-entry of row 4 has been switched to 0. The effect of the choice of the smallest row is to postpone other conflicts with this row as far as possible. In the example, it transforms the penultimate column into a zero-column. Therefore, the conflict appearing at position 7 can be easily solved.

We will prove that the resulting matrix is a minimal contention-free version of the original matrix. Its shadow is $(10111111)_2$.

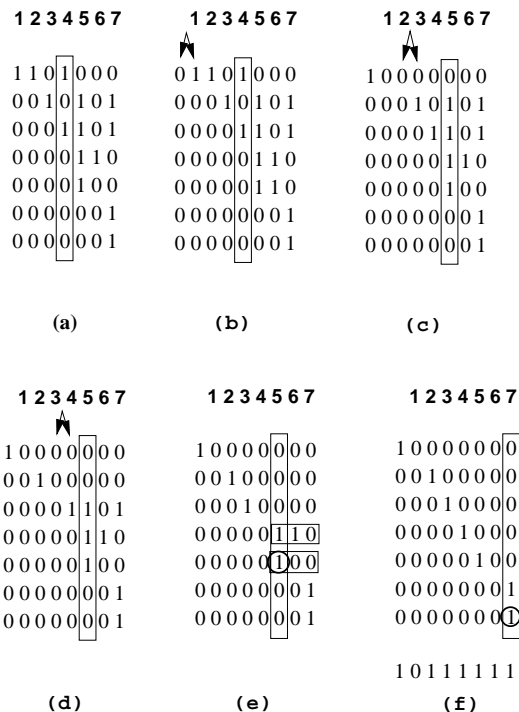


FIGURE 2. An example of the execution of Algorithm 1.

Remark. Note that it is not difficult to approximate $q^*(M)$ up to an additive factor of 1. Indeed, let M_0 be the matrix obtained from M by switching all 1-entries, but the leading 1-entry of each row, to zero. The reader can check that computing $q^*(M_0)$ and M_0^* is easy. For instance, on the example of Figure 2, $q^*(M_0) = 6$, and M_0^* is the identity matrix. Moreover, we have $q^*(M_0) \leq q^*(M) \leq q^*(M_0) + 1$. Indeed, eventually, we have to solve all contentions induced by leading 1's, that is $q^*(M) \geq q^*(M_0)$. Now, let M_1 be the $p \times (q + 1)$ matrix obtained from M_0 by adding one zero-column at position $q + 1$. All rows of M_1 are larger than the corresponding rows of M , therefore a minimal contention-free version of M_1 will give a contention-free version of M . Therefore, $q^*(M) \leq q^*(M_1) = q^*(M_0) + 1$. Unfortunately, approximating $q^*(M)$ up to an additive factor of 1 is not enough to provide a good approximation algorithm for the broadcast time of a tree. Indeed, we will see in Section 4 that one often need to solve the minimal contention-free matrix problem at all levels of the tree, and thus one would cumulate the error at each level.

THEOREM 3.1. *Algorithm 1 is an $O(q(p + q))$ -time algorithm that computes a minimal contention-free version of any $p \times q$ boolean matrix.*

First, let us show that Algorithm 1 performs in $O(q(q + p))$ steps.

LEMMA 3.2. *Algorithm 1 is an $O(q(p + q))$ -time algorithm.*

PROOF. The *for*-loop is executed q times, but the part “else” (Instruction 5) is not performed more than p times because there are p rows, and solving a contention

between 1-entries creates at least one row whose all entries are 0 after the current position. Let i be an index of the *for*-loop for which there is a contention. From what was said before, there are at most p such indices. Let k_i be the number of contending 1-entries: $\sum_i k_i \leq 2(p-1)$. All instructions before the *while*-loop do not require more than $O(p+q)$ time units. The *while*-loop is executed at most $q k_i$ times because each execution of the loop corresponds to a right-shift of a zero-column, and one cannot move a zero-column to the right more than q times, this for each of the k_i 1-entries. Actually, one can slightly modify the algorithm so that there are no more than q right-shifts in total, for all conflicts. Indeed, when shifting the zero-columns to the right, one can jump columns that were already exchanged with a zero-column since rule 1 was already applied. Altogether, rule 1 cannot be applied more than q times. Application of rule 1 has a cost of $O(q)$ since at most one row is updated after a right-shift. All other instructions inside the *while*-loop have a cost of $O(p+q)$. Instruction 25 has a cost of $O(q k_i)$, same as Instruction 27. Therefore, in total, the complexity is $O(q(q+p) + \sum_i q k_i)$ that is $O(q(q+p))$. \square

The fact that Algorithm 1 computes a minimal contention-free version of any $p \times q$ boolean matrix M is based on the following lemmas.

LEMMA 3.3. *If every rows A_i and B_i of two matrices A and B satisfy $A_i \leq B_i$, then $\text{shad}(A^*) \leq \text{shad}(B^*)$.*

PROOF. Every row of B^* is larger than the corresponding row of B , and so it is for A . Hence B^* is a contention-free version of A . Thus $\text{shad}(A^*) \leq \text{shad}(B^*)$. \square

Notation. Given two matrices A and B of the same number of rows p , and of respectively q and q' columns, AB denotes the $p \times (q+q')$ matrix obtained by putting A and B next to each other.

LEMMA 3.4. $\text{shad}((AB)^*) \leq \text{shad}\left(\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}^*\right)$.

PROOF. Let

$$\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}^* = \begin{pmatrix} X & 0 \\ Y & Z \end{pmatrix}.$$

For any row i , we have $X_i \geq A_i$. Also, for any row j , we have $Y_j Z_j \geq B_j$. $X' = X + Y$ has at most one 1-entry per column, that is $X'Z$ has at most one 1-entry per column. Moreover, $X'Z$ satisfies that, for any row i , $(X'Z)_i \geq (AB)_i$. Since $\text{shad}(X'Z) = \text{shad}\left(\begin{pmatrix} X & 0 \\ Y & Z \end{pmatrix}\right)$, the lemma holds. \square

Note that the inequality in Lemma 3.4 can be strict. For instance

$$\text{shad}\left(\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}^*\right) = 101$$

whereas

$$\text{shad}\left(\begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 1 \end{pmatrix}^*\right) = 111.$$

LEMMA 3.5. *If $(AB)^* = AB'$ where B' has the same number of columns as B , then $\text{shad}((AB)^*) = \text{shad}\left(\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}^*\right)$.*

PROOF. By lemma 3.4, we just have to show that $\text{shad}((AB)^*) \geq \text{shad}\left(\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}^*\right)$.

Let

$$C = \begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix} \text{ and } C' = \begin{pmatrix} A & 0 \\ 0 & B' \end{pmatrix}.$$

We have, for any row i , $C'_i \geq C_i$, and there is at most one 1-entry per column of C' . Since $\text{shad}(C') = \text{shad}((AB)^*)$, the lemma holds. \square

LEMMA 3.6. *Let X, X', Y, Y' be 1-dimensional vectors, and let A and A' be 2-dimensional matrices. Let*

$$M = \begin{pmatrix} X' & 0 & 1 & X \\ Y' & 0 & 1 & Y \\ A' & 0 & 0 & A \end{pmatrix}, \quad M_X = \begin{pmatrix} X' & 0 & 1 & X \\ Y' & 1 & 0 & 0 \\ A' & 0 & 0 & A \end{pmatrix}$$

and

$$M_Y = \begin{pmatrix} X' & 1 & 0 & 0 \\ Y' & 0 & 1 & Y \\ A' & 0 & 0 & A \end{pmatrix}$$

where there is at most one 1-entry per column in $\begin{pmatrix} X' \\ Y' \\ A' \end{pmatrix}$. Then

$$\text{shad}(M^*) = \min\{\text{shad}(M_X^*), \text{shad}(M_Y^*)\}.$$

PROOF. From Lemma 3.3, $\text{shad}(M^*) \leq \min\{\text{shad}(M_X^*), \text{shad}(M_Y^*)\}$. The equality holds because at least one 1-entry in the block $\begin{pmatrix} X' & 0 & 1 \\ Y' & 0 & 1 \end{pmatrix}$ must be moved to the left. \square

LEMMA 3.7. *With the same notations as in lemma 3.6, if $Y \leq X$ then $\text{shad}(M_Y^*) \leq \text{shad}(M_X^*)$.*

PROOF. Assume for the purpose of contradiction that $\text{shad}(M_Y^*) > \text{shad}(M_X^*) = \text{shad}(M^*)$. We get

$$M_X^* = \begin{pmatrix} X' & 0 & 1 & X'' \\ Z & z & 0 & 0 \\ B' & b & 0 & B \end{pmatrix}$$

where $Zz \geq Y'1$, $(B'b0B)_i \geq (A'00A)_i$ for every row i of these two matrices, and $X'' \geq X$. The first row of M_X^* is necessarily of the form $X'01X''$ because if this row would be larger or equal to $X'100$, then $\text{shad}(M_Y^*) \leq \text{shad}(M_X^*)$.

By the same arguments as for proving Lemma 3.5, we get

$$(3.1) \quad \text{shad}(M_X^*) = \text{shad}\left(\begin{pmatrix} X' & 0 & 0 & 0 \\ 0 & 0 & 1 & X' \\ Y' & 1 & 0 & 0 \\ A' & 0 & 0 & A \end{pmatrix}^*\right)$$

One can apply the same row-separation arguments on M_Y according to Lemma 3.4. Thus, $\text{shad}(M_Y^*) > \text{shad}(M_X^*)$ implies that

$$\text{shad}(M_X^*) < \text{shad}\left(\begin{pmatrix} X' & 1 & 0 & 0 \\ Y' & 0 & 0 & 0 \\ 0 & 0 & 1 & Y \\ A' & 0 & 0 & A \end{pmatrix}^*\right)$$

That is, by Lemma 3.3,

$$(3.2) \quad \text{shad}(M_X^*) < \text{shad}\left(\begin{pmatrix} X' & 1 & 0 & 0 \\ Y' & 0 & 0 & 0 \\ 0 & 0 & 1 & X \\ A' & 0 & 0 & A \end{pmatrix}^*\right) = \text{shad}\left(\begin{pmatrix} X' & 1 & 0 & 0 \\ 0 & 0 & 1 & X \\ Y' & 0 & 0 & 0 \\ A' & 0 & 0 & A \end{pmatrix}^*\right).$$

Equations 3.1 and 3.2 give

$$(3.3) \quad \text{shad}\left(\begin{pmatrix} X' & 0 & 0 & 0 \\ 0 & 0 & 1 & X \\ Y' & 1 & 0 & 0 \\ A' & 0 & 0 & A \end{pmatrix}^*\right) = \text{shad}\left(\begin{pmatrix} X' & 0 & 0 & 0 \\ 0 & 0 & 1 & X'' \\ Z & z & 0 & 0 \\ B' & b & 0 & B \end{pmatrix}\right)$$

$$(3.4) \quad < \text{shad}\left(\begin{pmatrix} X' & 1 & 0 & 0 \\ 0 & 0 & 1 & X \\ Y' & 0 & 0 & 0 \\ A' & 0 & 0 & A \end{pmatrix}^*\right)$$

Assume $b = 0$ and $z = 0$. Since $Zz \geq Y'1$, we have $Z > Y'$. Also, $B'_i > A'_i$ for at least one row i of A' and B' because otherwise there would have been no reason to replace $Y'1$ by $Z0$. Let us consider the row i such that $B'_i > A'_i$ and such that the rightmost bit position for which there is a 1-entry in B'_i is minimum. Let k and k' be the rightmost bit positions for which there is a 1-entry in Z and B'_i respectively. If $k < k'$ then replacing the row $(B'000)_i$ by $(A'100)_i$ from position k' would strictly decrease the shadow. If $k > k'$ then replacing $Z000$ by $Y'100$ from position k would also strictly decrease the shadow. We get a contradiction because the shadow of the matrix is supposed to be minimum.

Thus assume $b = 1$ or $z = 1$, for instance $z = 1$. This implies that the matrix

$$\begin{pmatrix} X' & 1 & 0 & 0 \\ 0 & 0 & 1 & X'' \\ Z & 0 & 0 & 0 \\ B' & 0 & 0 & B \end{pmatrix}$$

is a contention-free version of

$$\begin{pmatrix} X' & 1 & 0 & 0 \\ 0 & 0 & 1 & X \\ Y' & 0 & 0 & 0 \\ A' & 0 & 0 & A \end{pmatrix}$$

because $Z1 \geq Y'1 \Rightarrow Z \geq Y'$. This yields a contradiction with the strict inequality 3.4.

Thus $b = 1$. However then, the same argument as for $z = 1$ yields another contradiction with the strict inequality 3.4.

Therefore $\text{shad}(M_Y^*) > \text{shad}(M_X^*)$ cannot be satisfied. \square

LEMMA 3.8. *Let $M = AxByC$ where A is a matrix with at most one 1-entry per column, x is a zero-column, B is a matrix with exactly one 1-entry per column, y is a column with two contending 1-entries, and C is an unspecified boolean matrix. Let M' be the matrix resulting from M after an exchange between x and the leftmost column of B . We have $\text{shad}(M^*) = \text{shad}(M'^*)$.*

PROOF. We know from Lemma 3.3 that $\text{shad}(M^*) \leq \text{shad}(M'^*)$. The proof of the other inequality is by induction on the number of columns q of B . Assume $q = 1$, that is

$$M = \begin{pmatrix} A_1 & 0 & 1 & 0 & C_1 \\ A_2 & 0 & 0 & 1 & C_2 \\ A_3 & 0 & 0 & 1 & C_3 \\ A_4 & 0 & 0 & 0 & C_4 \end{pmatrix}.$$

Let

$$M^* = \begin{pmatrix} A'_1 & a_1 & a_2 & a_3 & C'_1 \\ A'_2 & b_1 & b_2 & b_3 & C'_2 \\ A'_3 & c_1 & c_2 & c_3 & C'_3 \\ A'_4 & d_1 & d_2 & d_3 & C'_4 \end{pmatrix}$$

with $A'_1 a_1 a_2 a_3 \geq A_1 0 1 0$, $A'_2 b_1 b_2 b_3 \geq A_2 0 0 1$, $A'_3 c_1 c_2 c_3 \geq A_3 0 0 1$, and $A'_4 d_1 d_2 d_3 C'_4 \geq A_4 0 0 0 C_4$. If $A'_1 a_1 a_2 a_3 \geq A_1 1 0 0$, then, by Lemma 3.3, $\text{shad}(M^*) = \text{shad}(M'^*)$. If $A_1 0 1 0 \leq A'_1 a_1 a_2 a_3 < A_1 1 0 0$, then we can assume w.l.g. that $A'_2 b_1 b_2 b_3 \geq A_2 1 0 0$. Actually, we can assume that

$$M^* = \begin{pmatrix} A'_1 & 0 & 1 & 0 & C'_1 \\ A'_2 & 1 & 0 & 0 & 0 \\ A'_3 & 0 & 0 & 1 & C'_3 \\ A'_4 & 0 & 0 & 0 & C'_4 \end{pmatrix}.$$

Let

$$M'' = \begin{pmatrix} A'_1 & 1 & 0 & 0 & 0 \\ A'_2 & 0 & 1 & 0 & 0 \\ A'_3 & 0 & 0 & 1 & C'_3 \\ A'_4 & 0 & 0 & 0 & C'_4 \end{pmatrix}.$$

We have $\text{shad}(M''^*) \leq \text{shad}(M^*)$. Now,

$$M' = \begin{pmatrix} A_1 & 1 & 0 & 0 & 0 \\ A_2 & 0 & 0 & 1 & C_2 \\ A_3 & 0 & 0 & 1 & C_3 \\ A_4 & 0 & 0 & 0 & C_4 \end{pmatrix},$$

and we get $\text{shad}(M'^*) \leq \text{shad}(M''^*) \leq \text{shad}(M^*)$, that is the lemma holds for $q = 1$.

Assume the lemma holds for every q , $1 \leq q < q_0$, and let us show that it holds for q_0 . A 1-entry in AxB must be moved to the left. For any move of a 1-entry in A , one can find a move of a 1-entry in B that preserves the shadow. Therefore, one can assume that it is a 1-entry in B that is moved to the left. Moreover, we can assume that this 1-entry, denoted by $\mathbf{1}$, is moved in xB .

- If $\mathbf{1}$ is moved in B at least one column to the right of the leftmost column of B , then one can apply the induction hypothesis, that is exchanging the first column of B with x , and then putting back $\mathbf{1}$ to its original position, without changing the shadow.
- If $\mathbf{1}$ is moved to the leftmost column of B , then we first apply Lemma 3.6, and then put back $\mathbf{1}$ to its original position. The result of these operations is just as exchanging x with the leftmost column of B . The shadow is preserved.

- If $\mathbf{1}$ is moved in x , then we can first exchange $\mathbf{1}$ with the 1-entry on the leftmost column of B , and then put back $\mathbf{1}$ to its original position, without changing the shadow.

Thus the result hold for q_o too. \square

We have now enough material to prove Theorem 3.1.

Proof of Theorem 3.1. Algorithm 1 constructs a finite sequence of matrices $M_0 = M, M_1, \dots, M_k$, such that M_i is obtained from M_{i-1} either by shifting a zero-column to the right, or by distributing 1-entries over zero-columns. Lemmas 3.7 and 3.8 insure that $\text{shad}(M_i^*) = \text{shad}(M_{i-1}^*)$, that is $\text{shad}(M^*) = \text{shad}(M_k^*)$. Since M_k is a contention-free version of M , we get that $\text{shad}(M^*) = \text{shad}(M_k)$. \square

4. Application to the multicast problem in tree-networks

As an example of application of Theorem 3.1 to the multicast problem in trees, let us consider the following problem. We are given a directed tree whose arcs are oriented from the root u toward the leaves, and a set D of nodes of the tree. We want to compute the minimum number of rounds that are required to multicast an information from u to all nodes in D . We are considering the all-port line communication model. In this context, Cohen [Coh98] has shown that there exists a polynomial-time algorithm that computes an optimal *broadcast* protocol from u to all nodes of T . To directly extend this algorithm to the multicast problem, we would make use of intermediate nodes that are not destination nodes, and this is not desirable in general. Combining Theorem 3.1 and the protocol in [Coh98] allows to overcome that problem.

COROLLARY 4.1. *There exists a polynomial-time algorithm that computes an optimal multicast protocol from any source to any destination set in any directed tree under the all-port line model, and such that only the source and the destination nodes participate to the protocol.*

PROOF. Let u be the source, and D be the destination set. The algorithm in [Coh98] proceeds bottom-up from the leaves to the source. Each node x has a list of calls stating when and to whom x gives a call in its subtree, and when and by who x is informed. This list is constructed from the lists of all the children of x in the tree. When the multicast problem is considered, the algorithm in [Coh98] fails in the following case: assume a node $x \in D$ has one of its children y not in D , and that y has k children z_1, \dots, z_k in D . The algorithm in [Coh98] requires the help of $y \notin D$. If we do not want y to be involved in the protocol, then x can be required to successively call z_1, z_2 up to z_k . More importantly, x cannot give a call simultaneously in the subtrees of the z_i 's, whereas y is able to do so in the all-port model. Therefore, giving the set of calls of y , one must schedule these calls so that x can simulate the behavior of y . One can represent the set of calls from y to the subtrees of the z_i 's by a matrix M such that $M_{i,j} = 1$ if and only if y gives a call to the subtree of z_i at round j . Theorem 3.1 gives a polynomial-time algorithm to schedule optimally these calls. Note that since this procedure must be applied at all the levels of the tree, one does not only need to compute a contention-free version of M with the minimum number of columns (i.e., number of rounds), but one also need to minimize the shadow. \square

5. Further research

We are currently working on an extension of Theorem 3.1 to make use of this result in the 1-port model. Again, the idea is to construct the protocol bottom-up from the leaves to the root. To make clear why Theorem 3.1 needs to be slightly adapted, let us consider the simple case of a *fork*, that is a particular type of directed tree in which the root u has a single child v which is the root of a star of p branches. Let X_i be the shadow on v of an optimal broadcasting algorithm applied to the i th branch, $i = 1, \dots, p$, and let M be the $p \times q$ array whose i th row is X_i .

A non necessarily optimal broadcast protocol in the 1-port line model consists in two phases: first u informs v , then v informs the p branches according to a minimal contention-free version of M . This protocol may be suboptimal because it can be more efficient to have both u and v informing the p branches (in the 1-port line model, u and v can call two distinct branches simultaneously). So the question is when to inform v ? Before v is informed, u only can inform the branches, and there is a contention in M when there is more than a single 1-entry on a column. After v has been informed, there is a contention in M when there is more than *two* 1-entries in a column. For instance, consider the following fork: u is connected to v , and v has two branches, composed of two nodes w_1, w_2 , and four nodes w'_1, w'_2, w'_3, w'_4 , respectively. One can broadcast from u in three rounds in this fork under the 1-port line model: (1) u calls w'_1 , (2) u calls w_1 , and w'_1 calls w'_3 , and (3) u calls v , w'_1 calls w'_2 , w'_3 calls w'_4 , and w_1 calls w_2 . If u calls v before the third round, then one more round is required. We are currently working on an extension of Algorithm 1 to solve that problem.

References

- [BET96] B.D. Birchler, A. Esfahanian, and E. Torng, *Information dissemination in restricted routing networks*, proceedings of the International Symposium on Combinatorics and Applications, 1996, pp. 33–44.
- [BFC93] T. Ballardie, P. Francis, and J. Crowcroft, *Core based tree (CBT): an architecture for scalable inter-domain multicast routing*, proceedings of SIGCOMM, ACM press, 1993, pp. 85–95.
- [BNGNS98] A. Bar-Noy, S. Guha, J. Naor, and B. Schieber, *Multicasting in heterogeneous networks*, ACM Symposium on Theory of Computing (STOC), ACM press, 1998.
- [CFKR98] J. Cohen, P. Fraigniaud, J-C. König, and A. Raspaud, *Optimized broadcasting and multicasting protocols in cut-through routed networks*, IEEE Transaction on Parallel and Distributed Systems **9** (1998), no. 8, 788–802.
- [CFM99] J. Cohen, P. Fraigniaud, and M. Mitjana, *Scheduling calls for multicasting in tree-networks*, 10th ACM-SIAM Symp. on Discrete Algorithms (SODA '99), 1999, pp. 881–882.
- [Coh98] J. Cohen, *Broadcasting, multicasting and gossiping in trees under the all-port line model*, 10th ACM Symposium on Parallel Algorithms and Architectures (SPAA), ACM press, 1998, pp. 164–171.
- [DDC97] C. Diot, W. Dabbous, and J. Crowcroft, *Multipoint communication: a survey of protocols, functions, and mechanisms*, IEEE journal on selected areas in communications **15** (1997), no. 3, 277–290.
- [DEF⁺94] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C-G. Liu, and L. Wei, *An architecture for wide-area multicast routing*, SIGCOMM, ACM press, 1994, pp. 126–135.
- [Far80] A. Farley, *Minimum-time line broadcast networks*, Networks **10** (1980), 59–70.
- [FL94] P. Fraigniaud and E. Lazard, *Methods and Problems of Communication in Usual Networks*, Discrete Applied Mathematics **53** (1994), 79–133.
- [FP81] A. Farley and A. Proskurowski, *Broadcasting in trees with multiple originators*, SIAM Journal Alg. Disc. Meth. **2** (1981), no. 4, 381–386.

- [FV97] P. Fraigniaud and S. Vial, *Approximation algorithms for broadcasting and gossiping*, Journal of Parallel and Distributed Computing **43** (1997), 47–55.
- [Har] H. Harutyunyan, Doctoral Thesis (in Russian), Armenia.
- [HHL86] S.M. Hedetniemi, S.T. Hedetniemi, and A. Liestman, *A survey of gossiping and broadcasting in communication networks*, Networks **18** (1986), 319–349.
- [HKMP95] J. Hromkovič, R. Klasing, B. Monien, and R. Peine, *Dissemination of information in interconnection networks (broadcasting and gossiping)*, Combinatorial Network Theory (Ding-Zhu Du and D. Frank Hsu, eds.), Kluwer Academic, 1995, pp. 125–212.
- [HKS96] J. Hromkovič, R. Klasing, and E. Stohr, *Dissemination of information in vertex-disjoint paths mode*, Computer and Artificial Intelligence **15** (1996), no. 4, 295–318.
- [HKS95] J. Hromkovič, R. Klasing, E. Stohr, and Wagener H., *Gossiping in vertex-disjoint paths mode in d-dimensional grids and planar graphs*, Information and Computation **123** (1995), no. 1, 17–28.
- [HKUW97] J. Hromkovič, R. Klasing, W. Unger, and H. Wagener, *Optimal algorithms for broadcast and gossip in the edge-disjoint path modes*, Information and Computation **133** (1997), no. 1, 1–33.
- [KP92] G. Kortsarz and D. Peleg, *Approximation algorithms for minimum time broadcast*, Proceedings of the Israel Symposium on Theory of Computer Science, Lecture Notes in Computer Science, no. 601, Springer-Verlag, 1992.
- [Lab89] R. Labahn, *Extremal broadcasting problems*, Discrete Applied Mathematics **23** (1989), 139–155.
- [Mid93] M. Middendorf, *Minimum broadcast time is NP-complete for 3-regular planar graphs and deadline 2*, Information Processing Letters **46** (1993), 281–287.
- [MS98] W. Mostafa and M. Singhal, *A taxonomy of multicast protocols for internet applications*, Computer communications **20** (1998), 1448–1457.
- [Pro81] A. Proskurowski, *Minimum broadcast trees*, IEEE Transactions on Computers **c-30** (1981), 363–366.
- [Rav94] R. Ravi, *Rapid rumor ramification: approximating the minimum broadcast time*, 35-th Symposium on Foundations of Computer Science, IEEE Computer Society Press, 1994, pp. 202–213.
- [SCH81] P. Slater, E Cockayne, and S. Hedetniemi, *Information dissemination in trees*, SIAM Journal on Computing **10** (1981), no. 4, 692–701.
- [SW84] P. Scheuermann and G. Wu, *Heuristic algorithms for broadcasting in point-to-point computer networks*, IEEE Transactions on Computers **C-33** (1984), no. 9, 804–811.
- [Win87] P. Winter, *Steiner problem in networks: a survey*, IEEE Networks **17** (1987), no. 2, 129–167.

Appendix A. Formal description of the algorithm

Algorithm 1

```

1  For  $i:=1$  to  $q$  do
   /* We sparse the columns from column 1 to column  $q$  */
2    $C_i :=$  current column;
3   if  $C_i$  is a zero-column then
4      $Z := Z \cup \{C_i\}$ ;
   /*  $Z$  currently denotes the set of zero-columns left to the current column */
5   else
6     if there is more than a single 1-entry in  $C_i$  then
7        $nb_1 :=$  # 1's in  $C_i$ ;
8        $W :=$  set of consecutive zero-columns to the left of  $C_i$ ;
9       not_yet_inserted := true;
10      While  $nb_1 - 1 > |W|$  and ( $Z \neq W$  or not_yet_inserted) do
   /* while there is still not enough zero-column immediately to the left of  $C_i$ , but */
   /* still some zero-columns that can be put immediately to the left of  $C_i$  */
11       $Z' := Z \setminus W$ ;
12       $c :=$  rightmost zero-column in  $Z'$ ;
13      Shift  $c$  one column to the right, and apply rule 1;
14       $Z :=$  set of zero-columns left to  $C_i$ ;
15       $W :=$  set of consecutive zero-columns immediately to the left of  $C_i$ ;
16       $nb_1 :=$  # 1's in  $C_i$ ;
17      if  $nb_1 - 1 > |W|$  and  $W = Z$  and not_yet_inserted then
   /* One needs to insert a zero-column */
18        Insert a zero-column at position 0;
19        not_yet_inserted := false;
20         $Z :=$  set of zero-columns left to  $C_i$ ;
21      EndIf
22    EndWhile
   /* Now, either there is enough zero-columns to solve all contentions, */
   /* or all zero-columns are immediately to the left of  $C_i$  */
23    if  $nb_1 - 1 > |W|$  then insert  $nb_1 - |W| - 1$  zero-columns left to  $C_i$ ;
   /* The  $nb_1$  1's will be spread out over the zero-columns of  $W$  */
24    Truncate each row with a 1 in  $C_i$  to keep only entries to the right of  $C_i$ ;
25     $\ell :=$  index of the row of min. lex. order among the truncated rows;
26     $W' :=$   $nb_1 - 1$  rightmost columns of  $W$ 
27    Spread out the  $nb_1$  1's of  $C_i$  over  $W'$ ; the 1-entry of row  $\ell$  stays in  $C_i$ ;
28     $Z :=$  set of zero-columns left to the current column;
29  EndIf
30 EndIf
31 EndFor

```

LABORATOIRE DE RECHERCHE EN INFORMATIQUE, BÂT. 490, UNIV. PARIS-SUD, 91405 ORSAY
CEDEX, FRANCE.

E-mail address: Johanne.Cohen@lri.fr

LABORATOIRE DE RECHERCHE EN INFORMATIQUE, BÂT. 490, UNIV. PARIS-SUD, 91405 ORSAY
CEDEX, FRANCE.

E-mail address: Pierre.Fraignaud@lri.fr

DEPT. MATEMATICA I TELEMATICA, UNIVERSITAT POLITECNICA DE CATALUNYA, 08034 BARCELONA,
SPAIN.

E-mail address: margarida@ma1.upc.es