Minimum-Time Broadcast under Edge-Disjoint Paths Modes*

PIERRE FRAIGNIAUD
CNRS-LRI, Université Paris-Sud, pierre@lri.fr

Abstract

This paper aims to study broadcast and multicast communication problems in networks under several variants of the edge-disjoint path mode. We derive upper and lower bounds for the approximability ratios of these problems (i.e., the worst-case ratios of the time required for a protocol computed in polynomial time to complete, over the time of an optimal protocol). These bounds show that slight modifications of the model (graphs vs. digraphs, all-port vs. single port, standard regimen vs. restricted regimen, etc.) have a tremendous impact on the difficulty, and of course on the complexity, of the problems.

1 Introduction

Given a node s of a network, and a set of nodes D, *multicasting* from s to D consists to transmit a piece of information from s to all nodes in D using the communication facilities of the network. *Broadcasting* is the specific case of multicasting in which the destination set consists of all nodes of the network. Multicasting and broadcasting are two of the basic operations upon which network applications are frequently based nowadays, and it hence gave rise to a vast literature, covering both applied and fundamental aspects of the problem (cf. [11, 21] and [18], respectively).

As far as fundamental aspects are concerned, the multicast problem can be expressed as follows. We are given a graph $G = (V,E)$ (whose edges model the communication links, and vertices model the network nodes), a source $s \in V$, and a destination set $D \subseteq V$. We are looking for the most efficient multicast protocol from s to D in G. The nature of the "protocol", and the measure of its "efficiency" depend on the communication model.

The *local model* assumes that transmissions proceed by synchronous *calls* between the nodes of the network. It is assumed moreover that (1) a call involves

*This work was partially done while the author was visiting Carleton University at Ottawa, and Concordia University at Montréal. Additional supports from Carleton, Concordia, NATO, and the RNRT project ROM.

exactly two neighboring nodes (locality constraint), (2) a node can participate to at most one call at a time (single-port constraint), and (3) the duration of a call is 1 (atomic constraint). A multicast protocol is then described by the list of calls placed between the nodes of the networks. The efficiency of the protocol is measured in term of number of *rounds*, where round t is defined as the set of all calls performed between time $t-1$ and time t, $t = 1,2,\ldots$ This model has been intensively investigated, for both specific and arbitrary topologies (cf. [15, 19] and [1, 2, 20, 23, 25], respectively).

The *line model* relaxes the locality constraint, and allows calls to be placed between non-neighboring nodes. A call is then a path in the network, whose two extremities are the "caller" and the "callee". (Note that, as opposed to the models in [5, 6], the intermediate nodes between the caller and the callee do not receive the information which just "cuts through" the routers attached to them.) Several variants of the line model have then been investigated. The *vertex-disjoint paths mode* specifies that, at any given round, the paths joining the two participants of every call must be pairwise vertex-disjoint (cf. [14, 20] and the references therein). This paper is concerned with another variant of the line-model, called *edge-disjoint paths mode*, i.e., at any given round, the paths joining the two participants of every call must be pairwise edge-disjoint.

Under this latter mode, it is known [9, 12] that, if the graph is undirected, then an optimal $\lceil \log_2 |D| \rceil$-round multicast protocol can be computed in polynomial time for any instance (G,s,D). However, very little is known in the case of directed graphs. It is only known that the corresponding decision problem becomes NP-complete (see [8]), and that there exists a polynomial 2-approximation algorithm for directed trees [4] (see also [10]) where a directed tree is rooted at the source, and its edges are directed from the root toward the leaves.

One can go a bit further in relaxing the constraints of the local model. For instance several contributions on multicasting under the edge-disjoint paths mode allow "all-port" communications (i.e., a node can participate to as many calls as its degree). It yields protocols that are faster than those derived for the single-port constraint. This has been verified for specific topologies [3, 22] as well as for trees [7]. However, very little is known in the case of arbitrary graphs, whatever they are directed or not. The only result the author is aware of is that the corresponding decision problems are NP-complete (see [8]).

This paper aims to investigate the multicast and broadcast problems in arbitrary graphs and digraphs under the single-port and all-port variants of the edge-disjoint paths mode. Actually, it is also worth to study a more sophisticated version of the problem, in which nodes of a given set F, specified *a priori*, cannot be involved to relay messages. Indeed, the need for efficient, secure, and fault-tolerant protocols requires that some nodes should not be involved, as they could be faulty, malicious, selfish, or just too slow. The setting of a model in which nodes in a given set F are forbidden to relay messages is called the *restricted regimen* of that model.

	All-port		Single-port	
	Bounds	References	Bounds	References
Graph	$\rho \leq O(\log\Delta)$	[Co. 1]	$\rho = 1$	[9, 12]
Digraph	$\rho \leq O(\log\Delta)$	[Co. 2]	$\rho \leq O(\frac{\Delta_{\min}+\log n}{\log(\Delta_{\min}+\log n)})$	[Th. 2]
			$\rho \geq \Omega(\log n)$	[Th. 3]

Table 1: Approximation ratios for the multicast problem.

	All-port		Single-port	
	Bounds	References	Bounds	References
Graph	$\rho \leq O(\log\Delta)$	[Co. 3]	$\rho = 1$	[9]
Digraph	$\rho \geq \Omega(\log\Delta)$	[Th. 1]	$\rho \leq \rho_{\text{VD}}$	[Th. 4]

Table 2: Restricted regimen. ρ_{VD} is the approximation ratio for the single-port multicast problem under the vertex-disjoint paths mode.

Notation. We denote by $b_s(G,D)$ the minimum number of rounds required to perform multicast from s to $D \subseteq V$ in $G = (V,E)$. For a shake of simplicity, we assume, w.l.o.g., that $s \in D$. If $D = V$, i.e., as far as the broadcast problem is concerned, $b_s(G,V)$ is simplified in $b_s(G)$. Δ is the maximum degree (resp., out-degree) of the considered graph (resp., digraph). $\Delta_{\min}$ is the minimum k such that there is a spanning tree of G with maximum (out-)degree k. If G is a digraph, then the arcs of the tree must be directed away from the root. We denote by n the number of nodes. Unless otherwise specified, all log's are in base 2. All ln's in base e.

Definition 1 An algorithm for the multicast problem is a ρ-approximation algorithm if, for any instance (G,s,D) of the problem, it returns a multicast protocol from s to D in G which completes in at most $\rho \cdot b_s(G,D)$ rounds.

We aim to derive approximation algorithms for the multicast and broadcast problems under the several variants of the edge-disjoint paths mode.

Our contributions. Our main contributions are summarized in Tables 1 and 2. That is, we give some upper and lower bounds on the approximation ratio ρ for the multicast problem in the considered variants of the edge-disjoint paths mode.

Structure of the paper. The remaining of this paper mostly consists of two parts, one dedicated to the all-port variant of the edge-disjoint paths mode (Section 2), and the other dedicated to the single-port variant (Section 3). Each of these sections is split in two subsections, one for the standard regimen, and the other for the restricted regimen. Finally, Section 4 contains some concluding remarks.

2 All-Port Multicast Problem

In order to make clear the all-port edge-disjoint paths mode, let us recall the model: communications proceed by a sequence of synchronous calls. During each round, nodes aware of the information can call other nodes in the graph, and, for each call, the callee can be at distance greater than one from the caller. A call is therefore a path in the graph linking the caller to the callee. The model specifies that two calls (i.e., two paths) performed during the same round must be edge-disjoint (resp., arc-disjoint if the network is modeled by a digraph). In particular, it implies that a node u aware of the information cannot call more that $d(u)$ nodes at each round, where $d(u)$ is the degree (resp., out-degree) of u.

The following is folklore:

Lemma 1 $b_s(G,D) \geq \lceil \log_{\Delta+1} |D| \rceil$ *where* Δ *is the maximum degree (resp. out-degree) of the graph (resp., digraph)* G*.*

Proof. Since the maximum (out-)degree of the (di)graph is Δ, the number of informed nodes can be multiplied by at most $\Delta + 1$ at each round. □

We consider first the case in which every node is allowed to receive and place calls. The restricted regimen is treated in Section 2.2.

2.1 Standard regimen

As said in the introduction, we have:

Lemma 2 (Cohen et al. [9], Farley [12].) *There is a polynomial-time algorithm which returns a* $\lceil \log_2 |D| \rceil$*-rounds* single-port *multicast protocol from* s *to* D *in the* graph G *for any instance* (G,s,D)*. Moreover, this protocol involves only nodes in* D *as callers and callees.*

As a trivial consequence of Lemmas 1 and 2, we get:

Corollary 1 *There exists a polynomial-time* $O(\log \Delta)$*-approximation algorithm for the all-port multicast problem in graphs.*

We show below that roughly the same result holds for directed graphs.

Lemma 3 *Let* $G = (V,E)$ *be an* n*-node digraph,* $D \subseteq V$*, and* $s \in D$*. There exists an all-port multicast protocol from* s *to* D *performing in at most* $2\lceil \log_2 |D| \rceil$ *rounds. Moreover, this protocol can be constructed in* $O(n \log n)$ *time.*

Proof. The protocol is constructed inductively. Let T be any directed tree rooted at s and spanning D. (Every edge of T is directed from the root s toward the leaves.) We denote by T_x the subtree of T rooted at x. Let x be a node of T such that T_x contains at least $|D|/2$ destination nodes, and every subtree T_y rooted at any child y of x contains less than $|D|/2$ destination nodes. Node x is called centroid. Let us consider the following multicast protocol from s. At the first round, s calls x. At the second round, s is idle, and x calls simultaneously all its children. Then we are let with subtrees containing no more that $|D|/2$ destination nodes each (i.e., $T \setminus T_x$ and the T_y's). There is at most $\lceil \log_2 |D| \rceil$ induction phases, and therefore the whole protocol completes in $2\lceil \log_2 |D| \rceil$ rounds.

Computing the centroid of a tree of k nodes takes $O(k)$ times by DFS. Thus the centroids of Phase i can be computed in $O(n)$ times. The result follows since there are $O(\log n)$ phases. □

Remark. One can decrease the constant factor in front of $\lceil \log_2 |D| \rceil$, for instance by choosing x such that $|T_x| > |D|/3$ and $|T_y| \leq |D|/3$ for every child y of x, and by letting s starts multicast in $T \setminus T_x$ as soon as it has called x.

From Lemmas 1 and 3, we immediately get:

Corollary 2 *There exists an $O(n \log n)$-time $O(\log \Delta)$-approximation algorithm for the all-port multicast problem in digraphs.*

2.2 Restricted regimen

The restricted regimen specifies that there is a set F of "forbidden" nodes. These nodes can be used to transmit calls (i.e., they can be crossed by paths), but not to transmit messages (i.e., they cannot be extremity of a path, neither as a caller nor as a callee). Of course, $F \cap D = \emptyset$.

The next result shows that $O(\log \Delta)$ is the best approximation ratio one can hope for the multicast problem if $F \neq \emptyset$. For that purpose, we need to recall the definition of the Minimum Set Cover problem (MSC for short).

Definition 2 (Minimum Set Cover problem.) We are given a collection $\mathbf{C}$ of subsets of a set S, and we want to find a smallest set $\mathbf{C}_{\min} \subseteq \mathbf{C}$ such that every element of S belongs to at least one set of $\mathbf{C}_{\min}$.

The following result is due to Feige.

Lemma 4 (Feige [13]) *Unless* $\mathrm{NP} \subset \mathrm{DTIME}(n^{\log \log n})$, *the optimal solution of the MSC problem is not approximable in polynomial time within* $(1-\varepsilon)\ln|S|$ *for any* $\varepsilon > 0$.

Observation 1 *Feige's result holds even if one restricts the MSC problem to instances $(S, \mathbf{C})$ satisfying $|\mathbf{C}| \leq |S|$.*

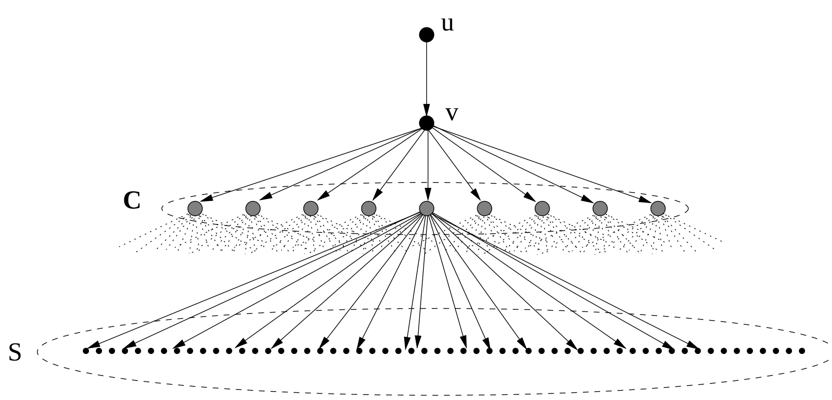

Figure 1: Gadget in the proof of Theorem 1.

Theorem 1 *Unless* NP $\subset$ DTIME$(n^{\log\log n})$*, the optimal solution of the all-port multicast problem in digraphs under the restricted regimen cannot be approximated in polynomial time within less than* $\Omega(\log\Delta)$*. Moreover, this is true even if the set of forbidden nodes contains a single node.*

Proof. Let $(S,\mathbf{C})$ be an instance of the Minimum Set Cover problem. Let us construct an instance of the multicast problem obtained from $(S,\mathbf{C})$ as follows (see Figure 1). The digraph $G=(V,E)$ has $|S|+|\mathbf{C}|+2$ vertices:

$$V = S \cup \mathbf{C} \cup \{u,v\},$$

where u and v are two nodes, and

$$E = \{(u,v)\} \cup \{(v,C), C\in\mathbf{C}\} \cup \{(C,s), C\in\mathbf{C}, s\in C\}.$$

We complete the instance of the multicast problem by setting the destination-set to S, the source-node to u, and the forbidden set to $\{v\}$.

We claim that the optimal solution $\mathbf{C}_{\min}$ of $(S,\mathbf{C})$ satisfies $|\mathbf{C}_{\min}| \leq b_u(G,S) \leq |\mathbf{C}_{\min}|+1$. Indeed, let $\mathbf{C}_{\min} \subseteq \mathbf{C}$ be an optimal solution of Minimum Set Cover, $\mathbf{C}_{\min} = \{C_1,\ldots,C_k\}$. Then u performs multicast to S as follows: for $i \in \{1,\ldots,k\}$, u calls C_i at round i. Once informed, C_i calls simultaneously all the $s \in C_i$. This multicast protocol completes in $k+1$ rounds. On the other hand, during an optimal multicast protocol from u to S, at most $b_u(G,S)$ sets in $\mathbf{C}$ are either traversed by a call from u, or destination of a call from u. Since every $s\in S$ is aware of the message after $b_u(G,S)$ rounds, these $b_u(G,S)$ sets cover S, and therefore $b_u(G,S) \geq k$.

Therefore, $|\mathbf{C}_{\min}| \leq b_u(G,S) \leq 2|\mathbf{C}_{\min}|$. Since we know from Feige's lemma that the optimal solution of the MSC problem is not approximable within $(1-$

$\varepsilon)\ln|S|$, we get that $b_u(G,S)$ cannot be approximated within a ratio $(\frac{1}{2}-\varepsilon)\ln|S|$ for any $\varepsilon > 0$. We conclude the proof by noticing that $\Delta \leq |S|$ because the instance can be chosen so that $|\mathbf{C}| \leq S$ (cf. Observation 1). □

As far as upper bounds are concerned, the following result is a trivial consequence of Lemmas 1 and 2:

Corollary 3 *There exists a polynomial-time $O(\log\Delta)$-approximation algorithm for the all-port multicast problem in graphs under the restricted regimen.*

For digraphs, we just mention a much weaker result. Let us introduce some notation. Let $G=(V,E)$ be a strongly connected digraph, let $s \in V$, $D \subset V$, and $F \subset V$, with $F \cap D = \emptyset$. For every $v \in V$, we define $\Gamma^+(v)$ as the set of nodes reachable from v in $G \setminus F$. Let $N^+(F)$ be the set of the out-neighbors of the nodes in F, and let $\{u_1,\ldots,u_k\} = \{s\} \cup (N^+(F)\setminus F)$. For $i=1,\ldots,k$, let $C_i = \Gamma^+(u_i) \cap D$. Let $\mathbf{C} = \{C_1,\ldots,C_k\}$. Let $\mathbf{C}^*$ be an approximated solution of the MSC problem with instance $(D,\mathbf{C})$. Let $\mathbf{C}^* = \{C_i, i \in I\}$. For every $i \in I$, let $\widehat{C_i} = C_i \setminus \cup_{j\in I, j<i} C_j$ so that for every $i,j \in I$, $i \neq j$, $\widehat{C_i} \cap \widehat{C_j} = \emptyset$. Given that, one can multicast from s to D in G avoiding F in two phases. During Phase 1, the source s calls u_i for every $i \in I$. During Phase 2, every u_i, $i \in I$, multicasts to $\widehat{C_i}$. Phase 1 takes at most $|\mathbf{C}^*|$ rounds. From Lemma 3, Phase 2 takes at most $2\lceil\log_2|D|\rceil$ rounds. By application of Lemma 1, we get that there exists a polynomial-time $O((1+\frac{|C^*|}{\log|D|})\log\Delta)$-approximation algorithm for the all-port multicast problem in digraphs under the restricted regimen. This strategy is not expected to perform efficiently in general, but note however, that, since $|C^*| \leq \mathbf{C} \leq \Delta|F|$, it yields an $O(\log\Delta)$-approximation algorithm for instances of the multicast problem in which $\Delta|F| = O(\log|D|)$. More interestingly, if $|F|=1$, then multicasting takes at least $\Omega(|C_{\min}| + \log_\Delta|D|)$ where $\mathbf{C}_{\min}$ is an optimal solution of the MSC problem $(D,\mathbf{C})$. Therefore, since $|\mathbf{C}_{\min}|$ can be easily approximated up to a multiplicative factor of $O(\log|D|)$, we get an $O(\log|D| + \log\Delta)$-approximation algorithm, which is optimal for a single forbidden nodes (cf. the proof of Theorem 1).

3 Single-Port Multicast Problem

Let us now consider the single-port model. In this setting, communications still proceed by a sequence of synchronous calls, but, as opposed to the all-port model, nodes aware of the information can call at most one other node during each round. Again, for each call, the callee can be at distance greater than one from the caller, and a call is hence again a path in the graph linking a caller and a callee. The model specifies that two calls performed during the same round must be edge-disjoint (resp., arc-disjoint in digraphs). For instance, let $\vec{F}_n$, $n \geq 3$, be the "Fork" digraph of n nodes: $V(\vec{F}_n) = \{u,v,w_1,\ldots,w_{n-2}\}$ and $E(\vec{F}_n) = \{(u,v),(v,w_i),(w_i,u), i = 1,\ldots,n-2\}$. The undirected Fork graph F_n is obtained from $\vec{F}_n$ by removing the orientation of the arcs. We have $b_u(F_n) = \lceil\log_2 n\rceil$, but $b_u(\vec{F}_n) = \lceil n/2\rceil$.

The following is folklore:

Lemma 5 $b_s(G,D) \geq \lceil \log_2 |D| \rceil$ *for any graph or digraph* $G = (V,E)$*, any* $D \subseteq V$*, and any* $s \in D$*.*

Proof. The single-port constraint implies that the number of informed nodes can at most double at each round. □

As in the previous section, we consider first the case in which every node is allowed to receive and place calls. The restricted regimen is treated in Section 3.2.

3.1 Standard regimen

If G is undirected, then Lemma 2 applies, i.e., $b_s(G,D) = \lceil \log_2 |D| \rceil$ for every instance (G,s,D), and the corresponding multicast protocol can be computed in polynomial time.

The example of the Fork digraph shows however that there are instances of the multicast problem in digraphs for which $b_s(G,D) = \Omega(n)$. On the other hand, it is trivial observation that multicasting in Hamiltonian digraphs can be done in $\lceil \log_2 |D| \rceil$ rounds by just using an Hamiltonian cycle. This observation applies to Eulerian digraphs as well. Actually, multicasting can also be achieved in $\lceil \log_2 |D| \rceil$ rounds if there is an Eulerian or Hamiltonian partial subgraph of G spanning D.

If fact, one can easily derive an approximation algorithm for the broadcast problem, based on approximated solutions of the Minimum Degree Spanning Tree problem (MDST for short). The MDST problem is defined as follows. Given any digraph $G = (V,E)$, and any set $D \subseteq V$, let $\Delta_{\min}(G,D)$ be the smallest integer k such that there exists a rooted directed tree of maximum out-degree k, spanning D in G. (The arcs of the tree are directed away from the root.) $\Delta_{\min}(G,D)$ is abbreviated in $\Delta_{\min}(G)$ if $D = V$.

Definition 3 (Minimum Degree Spanning Tree problem). We are given a (di)graph G, and we want to compute a spanning tree of maximum (out-)degree $\Delta_{\min}(G)$. In the Steiner version of the MDST problem, we are given G and a set of nodes D, and want to compute a tree spanning D and of maximum (out-)degree $\Delta_{\min}(G,D)$.

The following result is due to Fürer and Raghavachari.

Lemma 6 (Fürer and Raghavachari [16].)
There is a polynomial-time algorithm which, given any digraph $G = (V,E)$ *and node* $r \in V$*, returns a directed spanning tree of* G *rooted at* r *and of maximum out-degree* $O(\Delta_{\min}(G) + \log n)$*.*

The previous lemma will be used in combination with the next one:

Lemma 7 *There is a polynomial-time algorithm which, given any directed tree $T=(V,E)$ rooted at r, and of maximum out-degree Δ, and given any $D\subseteq V$, computes a multicast protocol from r to D performing in $O(\frac{\Delta}{\log\Delta}\cdot\log|D|)$ rounds.*

Proof. A multicast protocol M from r to D in T is constructed inductively, and is similar to the protocol described in the proof of Lemma 3. Again, we denote by T_x the subtree of T rooted at x. Let x be a node such that T_x contains at least $|D|/2$ destination nodes, and every subtree T_y rooted at any child y of x contains less than $|D|/2$ destination nodes. The first round is a call from r to x. Once r has performed this call, it starts multicasting in $T\setminus T_x$. Once x is informed, it spends at most Δ rounds to inform its children as follows. Let $y_1,\dots,y_q$ be the $q\leq\Delta$ children of x whose subtrees contain at least one destination node, and let d_i be the number of destination nodes in T_{y_i}. Assume, w.l.o.g., that $d_1\geq d_2\geq\cdots\geq d_q\geq 1$. Then x call y_1 first, then y_2, and so on until y_k. As soon as y_i is informed, it starts multicasting in T_{y_i}.

At this point, we are let with subtrees containing at most $|D|/2$ destination nodes each, and we can apply the same strategy inductively in each subtree. Since there are $O(\log|D|)$ phases of induction, and since each phase requires at most $\Delta+1$ rounds, M takes at most $O(\Delta\log|D|)$ rounds. Let us show that the order in which x calls its children allows to save a logarithmic factor in this analysis.

Let us show by induction on $\lceil\log_2|D|\rceil$ that M completes in at most $(\frac{\Delta}{\log\Delta}+O(1))\log|D|$ rounds. Let $b(k)$ be the maximum of the completion time of M for destination sets D of cardinality at most k. We have

$$b(2k)\leq 1+\max_{1\leq i\leq\Delta}(i+b(2k/i))$$

since the ith subtree T_{y_i} cannot contain more than $2k/i$ destination nodes. Actually, since T_{y_1} cannot contain more than k destination nodes, we have:

$$b(2k)\leq 1+\max_{2\leq i\leq\Delta}(i+b(2k/i)).$$

Therefore, assuming $b(k)\leq\frac{\Delta}{\log\Delta}\log k+\alpha(k)$ yields

$$b(2k)\leq 1+\frac{\Delta}{\log\Delta}\log 2k+\max_{2\leq i\leq\Delta}(i-\frac{\Delta}{\log\Delta}\log i+\alpha(2k/i))$$

and hence

$$b(2k)\leq 1+\alpha(k)+\frac{\Delta}{\log\Delta}\log 2k+\max_{2\leq i\leq\Delta}(i-\frac{\Delta}{\log\Delta}\log i).$$

Since $\max_{2\leq i\leq\Delta}(i-\frac{\Delta}{\log\Delta}\log i)\leq 0$, we get

$$b(2k)\leq 1+\alpha(k)+\frac{\Delta}{\log\Delta}\log 2k$$

that is $b(2k) \leq \frac{\Delta}{\log\Delta}\log 2k + \alpha(2k)$ where $\alpha(2k) = \alpha(k) + 1 = O(\log k)$. □

We get:

Theorem 2 *Let $t(n)$ be the time complexity of an algorithm which, given any n-node digraph $G = (V,E)$ and any node $r \in V$, returns a directed spanning tree of G rooted at r and of maximum out-degree $O(\Delta_{\min}(G) + \log n)$. There exists a $t(n)$-time $O(\frac{\Delta_{\min}(G)+\log n}{\log(\Delta_{\min}(G)+\log n)})$-approximation algorithm for the single-port multicast problem in digraphs.*

Proof. Given $G = (V,E)$, $D \subseteq V$ and $s \in V$, we compute in $t(n)$-time a tree rooted at s, spanning D and of maximum degree $\Delta \leq \Delta_{\min}(G) + \log n$. Applying Lemma 7, we get a multicast protocol performing in at most $O(\frac{\Delta}{\log\Delta} \cdot \log|D|)$ rounds, that is at most $O\left(\frac{\Delta_{\min}(G)+\log n}{\log(\Delta_{\min}(G)+\log n)} \cdot \log|D|\right)$ rounds since $x/\log x$ is non decreasing for x large enough. The result then follows from Lemma 5. □

From Lemma 6, $t(n)$ can be chosen to be polynomial.

Remarks. Theorem 2 could be improved by using a spanning tree of maximum out-degree $\Delta_{\min}(G,D)$ instead of $\Delta_{\min}(G)$. This would require to extend Lemma 6 to the Steiner version of the MDST problem in directed networks. Note however that it is known [14] that, unless NP $\subset$ DTIME$(n^{\log\log n})$, any polynomial-time approximation algorithm for the Steiner version of the MDST problem in digraphs, returning a tree of maximum out-degree at most $\rho \cdot \Delta_{\min}(G,D) + r$, satisfies $\rho + r = \Omega(\log n)$. (This impossibility result holds for directed graphs only, and it is shown in [17] that the Steiner version of the MDST problem in graphs can be approximated up to an additive factor of 1.)

One can also give a lower bound on ρ.

Theorem 3 *Unless* NP $\subset$ DTIME$(n^{\log\log n})$, *the optimal solution of the single-port multicast problem in digraphs cannot be approximated in polynomial time within less than* $\Omega(\log n)$.

Proof. From an instance $(S,\mathbf{C})$ of the MSC problem we construct an instance of the multicast problem as follows (see Figure 2). We construct a digraph $G = (V,E)$ by taking k copies of the digraph $G' = (V',E')$ where $V' = \{v\} \cup \mathbf{C} \cup S$ and $E' = \{(v,C), C \in \mathbf{C}\} \cup \{(C,s), C \in \mathbf{C}, s \in C\}$, merged at v, and by replacing every star $\{(C,s), s \in C\}$ by a (possibly pruned) binomial tree [24] of at most $2|C|$ nodes rooted at C. We set D to be the k copies of S, and $s = v$. Let $\mathbf{C}_{\min} = \{C_1,\ldots,C_\ell\}$ be an optimal solution of MSC for the instance $(S,\mathbf{C})$. We have $b_v(G,D) \leq k|\mathbf{C}_{\min}| + \lceil\log_2|S|\rceil$ as follows: for $i = 1$ to ℓ, and for $j = 1$ to k, v calls the jth copy of C_i at round $i + \ell(j-1)$. Once informed, a node C_i multicasts to the node-set $C_i \subset S$ in at most $\lceil\log_2|C_i|\rceil$ rounds using the binomial tree rooted at C_i.

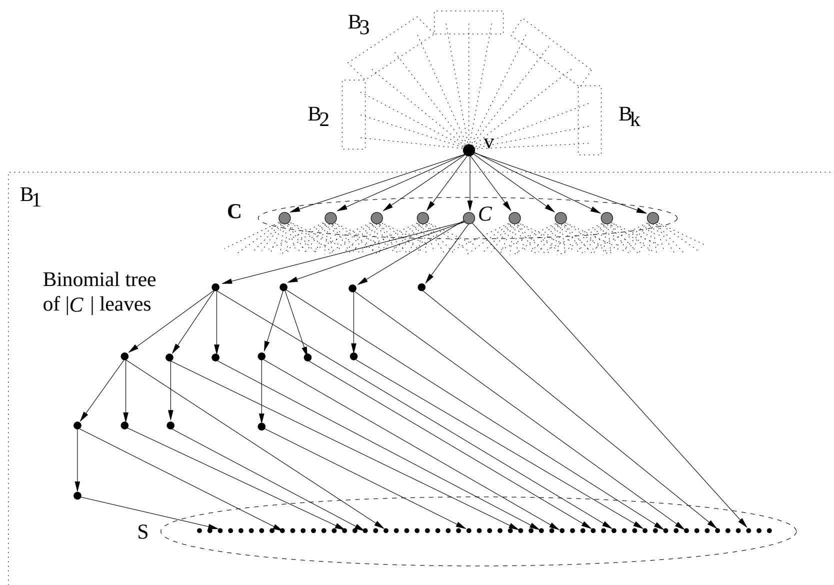

Figure 2: Gadget in the proof of Theorem 3.

Let A be a ρ-approximation algorithm for the multicast problem. Given (G,v,D), A returns a multicast protocol M from v to D in G performing in at most $\rho \cdot b_v(G,D)$ rounds, i.e., at most $\rho \cdot (k|\mathbf{C}_{\min}| + \lceil \log_2 |S| \rceil)$ rounds. M determines a collections $\mathbf{C}_1^*, \ldots, \mathbf{C}_k^*$ of subsets of $\mathbf{C}$ where $\mathbf{C}_i^*$ is the set of nodes of the ith copy of $\mathbf{C}$ that are either traversed by a call from v, or destination of a call from v. Each $\mathbf{C}_i^*$ covers S. Let $\mathbf{C}^* \in \{\mathbf{C}_1^*, \ldots, \mathbf{C}_k^*\}$ such that $|\mathbf{C}^*| = \min\{|\mathbf{C}_i^*|, i = 1, \ldots, k\}$. M completes in at least $k|\mathbf{C}^*|$ rounds. Therefore, we get: $k|\mathbf{C}^*| \leq \rho(k|\mathbf{C}_{\min}| + \lceil \log_2 |S| \rceil)$. Hence $|\mathbf{C}^*|/|\mathbf{C}_{\min}| \leq \rho(1 + \lceil \log_2 |S| \rceil / (k|C_{\min}|)) \leq \rho(1 + \lceil \log_2 |S| \rceil / k)$. By choosing $k = \lceil \log_2 |S| \rceil$, and by applying Feige's lemma, we get $\rho \geq (1-\varepsilon) \ln |S| / 2$. Since $n \leq 1 + k|\mathbf{C}| + k|\mathbf{C}||S| + k|S|$, and since $|\mathbf{C}| \leq |S|$ from Observation 1, we get $n \leq |S|^3$ and hence $\rho \geq (\frac{1}{6} - \varepsilon) \ln n$. □

3.2 Restricted regimen

In this section, we denote by F the set of "forbidden" nodes. From Lemma 2, we still have $b_s(G,D) = \lceil \log_2 |D| \rceil$ for every instance (G,s,D,F) of the single-port multicast problem in graphs under the restricted regimen. Actually, the result holds even is $F = V \setminus D$.

On the other hand, we are not aware of any result about the single-port multicast problem in digraphs under the restricted regimen. One can easily show though that there is a polynomial reduction from this problem to the same problem in the vertex-disjoint paths mode. This shows the strong interplay between the two variants of the line model.

Theorem 4 *The single-port multicast problem in digraphs under the restricted regimen of the edge-disjoint paths mode is polynomially reducible to the single-port multicast problem in digraphs under the restricted regimen of the vertex-disjoint paths mode.*

Proof. Let $G=(V,E)$ be any digraph, and let H be the digraph obtained from G as follows. Recall that the *line digraph* of G is the digraph $L(G)$ whose set of nodes is the set of arcs of G, and there is an arc from a node e of $L(G)$ to a node e' of $L(G)$ if and only if e' is incident to e in G, that is $e=(x,y)$ and $e'=(y,z)$. Every node v of G yields a complete bipartite digraph $K_{\deg^-(v),\deg^+(v)}$ in $L(G)$. H is the digraph obtained from $L(G)$ by adding one vertex to every such complete bipartite digraph. More precisely, assume that $L(G)=(E,X_E)$. Then $H=(V\cup E,X_V\cup X_E)$ where, for every $v\in V$, there is an arc in X_V from every $e\in\Gamma^-(v)$ to v, and from v to every $e\in\Gamma^+(v)$. It is easy to check that, under the single-port edge-disjoint paths mode, for every $D\subseteq V$, every $s\in V$, and every $\mathcal{F}\subseteq V\setminus D$, multicasting from s to D in G avoiding nodes in $\mathcal{F}$ takes the same number of rounds as multicasting from s to D in H avoiding nodes in $E\cup\mathcal{F}$, under the vertex-disjoint paths mode. □

Note that the proof above works for $\mathcal{F}=\emptyset$ as well, that is under the standard single-port model. Unfortunately, very little is known about multicasting under the restricted regimen of the single-port vertex-disjoint paths mode (see [14] for the most recent advances). Theorem 4 shows that progresses in this domain would help for the design of solutions for the edge-disjoint paths mode as well.

Remark. Observe that, as a consequence of Theorem 2, we trivially get that there exists a polynomial-time $O(|\mathcal{F}|\frac{\Delta_{\min}(G)+\log n}{\log(\Delta_{\min}(G)+\log n)})$-approximation algorithm for the single-port multicast problem in digraphs under the restricted regimen. Indeed, any protocol executed in the standard model can be modified so that nodes in $\mathcal{F}$ are not involved, as follows. At each round, every call whose callee is in $\mathcal{F}$ is not performed, and every call whose caller is in $\mathcal{F}$ is performed by the source. Every round of the original protocol can involve at most $|\mathcal{F}|$ forbidden nodes as caller, therefore the modified protocol has its number of rounds multiplied by at most $|\mathcal{F}|$.

4 Conclusion

As the reader could notice while reading this paper, many problems remain open. In almost all variants of the edge-disjoint path mode, there is plenty of room for improvements, and better approximation algorithms for the multicast problem are expected (see Tables 1 and 2).

In particular, a natural question that arises in this context is whether one can derive better approximation algorithms than the somewhat straightforward $O(\log\Delta)$-approximation algorithms for the all-port cases in both graphs and digraphs.

It would also be quite interesting to derive an $O(\log n)$-approximation algorithm for the single-port case in digraphs.

The more constrained restricted regimen seems to be quite challenging and deserves deeper investigations.

Acknowledgements. The author is thankful to Lali Barrière for her advice on preliminary versions of this paper. He is also thankful to David Barrington and Micah Adler for their helpful comments.

References

[1] A. Bar-Noy, S. Guha, J. Naor, and B. Schieber. Multicasting in heterogeneous networks. In *30th ACM Symposium on Theory of Computing (STOC '98)*, pages 448–453, 1998.

[2] D. Barth and P. Fraigniaud. Approximation algorithms for structured communication problems. In *9th ACM Symposium on Parallel Algorithms and Architectures (SPAA '97)*, pages 180–188, 1997. (Tech. Rep. LRI-1239, Univ. Paris-Sud. http://www.lri.fr/˜pierre).

[3] J.-C. Bermond, A. Bonnecaze, T. Kodate, S. Prennes, and P. Sole. Broadcasting in hypercubes under the circuit-switched model. In *IEEE Int. Parallel and Distributed Processing Symp. (IPDPS)*, 2000.

[4] B. Birchler, A.-H. Esfahanian, and E. Torng. Toward a general theory of unicast-based multicast communication. In *21st Workshop on Graph-Theoretic Concepts in Computer Science (WG '95)*, volume 1017 of *LNCS*, pages 237–251, 1995.

[5] C.-T. Chou and I. Gopal. Linear broadcast routing. *Journal of Algorithms* Vol. 10, (1989), 490–517.

[6] I. Cidon, I. Gopal and S. Kutten. New models and algorithms for future networks. *IEEE Trans. on Inf. Th.* Vol IT-41, (1995), 769–780.

[7] J. Cohen. Broadcasting, multicasting and gossiping in trees under the all-port line model. In *10th ACM Symposium on Parallel Algorithms and Architectures (SPAA '98)*, pages 164–171, 1998.

[8] J. Cohen. *Communications multipoints dans le modèle commuté*. PhD thesis, LRI, Université Paris-Sud, 91405 Orsay, France, 1999.

[9] J. Cohen, P. Fraigniaud, J.-C. Konig, and A. Raspaud. Broadcasting and multicasting in cut-through routed networks. In *11th IEEE Int. Parallel Processing Symposium (IPPS '97)*, pages 734–738, 1997.

[10] J. Cohen, P. Fraigniaud, and M. Mitjana. Scheduling calls for multicasting in tree-networks. In *10th ACM-SIAM Symp. on Discrete Algorithms (SODA '99)*, pages 881–882, 1999.

[11] C. Diot, W. Dabbous, and J. Crowcroft. Multipoint communication: a survey of protocols, functions, and mechanisms. *IEEE Journal on Selected Areas in Communications*, 15(3):277–290, 1997.

[12] A. Farley. Minimum-time line broadcast networks. *Networks*, 10:59–70, 1980.

[13] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45:634–652, 1998.

[14] P. Fraigniaud Approximation Algorithms for Collective Communications with Limited Link and Node-Contention. Tech. Rep. LRI-1264, Univ. Paris-Sud, France. (www.lri.fr/~pierre)

[15] P. Fraigniaud and E. Lazard. Methods and Problems of Communication in Usual Networks. *Discrete Applied Mathematics*, 53:79–133, 1994.

[16] M. Fürer and B. Raghavachari. Approximating the minimum degree spanning tree within one from the optimal degree. In *3rd Annual ACM-SIAM Sumposium on Discrete Algorithms (SODA '92)*, pages 317–324, 1992.

[17] M. Fürer and B. Raghavachari. Approximating the minimum-degree steiner tree to within one of optimal. *Journal of Algorithms*, 17:409–423, 1994.

[18] S.M. Hedetniemi, S.T. Hedetniemi, and A. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18:319–349, 1986.

[19] J. Hromković, R. Klasing, B. Monien, and R. Peine. Dissemination of information in interconnection networks (broadcasting and gossiping). In Ding-Zhu Du and D. Frank Hsu, editors, *Combinatorial Network Theory*, pages 125–212. Kluwer Academic, 1995.

[20] G. Kortsarz and D. Peleg. Approximation algorithms for minimum time broadcast. *SIAM Journal on Discrete Methods*, 8:401–427, 1995.

[21] W. Mostafa and M. Singhal. A taxonomy of multicast protocols for internet applications. *Computer Communications*, 20:1448–1457, 1998.

[22] J. Peters and M. Syska. Circuit-switched broadcasting in torus networks. *IEEE Transactions on Parallel and Distributed Systems*, 7(3):246–255, 1996.

[23] R. Ravi. Rapid rumor ramification: approximating the minimum broadcast time. In *35-th IEEE Symposium on Foundations of Computer Science*, pages 202–213, 1994.

[24] J. Vuillemin. A data structure for manipulating priority queues. *Comm. of the ACM*, 21:309–315, 1978.

[25] C. Schindelhauer. On the inapproximability of broadcasting time. In *3rd Int. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, LNCS. Springer-Verlag, 2000.