

Logarithmic Inapproximability of the Radio Broadcast Problem

Michael Elkin ^{*†} Guy Kortsarz [‡]

Abstract

We show that the radio broadcast problem is $\Omega(\log n)$ -inapproximable unless $NP \subseteq BPTIME(n^{O(\log \log n)})$. This is the first result on the hardness of approximation of this problem. Our reduction is based on the reduction from the Label-Cover problem to the Set Cover problem due to Lund and Yannakakis [LY94], and uses some new ideas.

^{*}Department of Computer Science, Yale University, New Haven, CT, USA, 06520-8285, elkin@cs.yale.edu.

[†]Part of this work was done in School of Mathematics, Institute for Advanced Study, Princeton, NJ, USA, 08540.

[‡]Computer Science department, Rutgers University, Camden, NY, USA. Email: guyk@crab.rutgers.edu

1 Introduction

Consider a network of processors modeled by an undirected n -vertex graph $G = (V, E)$. Assume that the communication in the network is synchronous, i.e., occurs in discrete pulses, called *rounds*. Suppose also that there is a designated vertex $s \in V$, called *source*, that generates a message μ . On each round of the communication each processor that knows the message μ is allowed to send it to all its neighbors. Note that a processor can either send the message μ to all of them or not send μ at all, but it can transmit to no strict subset of its set of neighbors. Furthermore, only a processor that receives the message from precisely one neighbor in a certain round is considered to be informed in this round. The intuition is that processors, that receive a message from more than one neighbor in the same round, get it corrupted. The *radio broadcast* problem requires to compute a schedule with minimal number of rounds that delivers a message from the source to all the other processors in the network.

The radio broadcast is a basic primitive in distributed computing and computer communication theory, and is used as a building block for various more complicated tasks in these areas (see, e.g., [ABLP91, ABLP92, BGI91, CK85, CK87, CK2-85, CW87, GVF87, GM95, KGBK78, KP02, SC82]).

Let $rad(G, s)$ be the largest distance between the source s and any other vertex in the graph G . Bar-Yehuda et al. [BGI91] and Gaber and Mansour [GM95] have shown that for any graph G and vertex s in G , a radio broadcast that was originated from the source s can be completed in $O(\min\{rad(G, s) + \log^5 n, (rad(G, s) + \log n) \log n\})$ rounds. On the other hand, Alon et al. [ABLP91] have shown that for any $r = 2, 3, \dots$, there is an infinite family of graphs of radius r , on which the radio broadcast requires $rad(G, s) + \Omega(\log^2 n)$ rounds. While the aforementioned upper bound of [BGI91] can be interpreted as a (randomized) $O(\log^2 n)$ -approximation algorithm for the problem, to the best of our knowledge, no lower bound on its *hardness of approximation* was known prior to the current paper. (A *deterministic* $O(\log^2 n)$ -approximation algorithm for the problem was devised in [CW87]; the exact version of the problem is known to be NP-hard even when restricted to disc graphs [CK85, GPM03].) However, an evidence of its hardness was recently discovered by Bar-Yehuda [B01]. It is shown there that for any $c > 0$ the existence of a $(c \cdot \log n)$ -approximation algorithm for the radio broadcast problem implies an existence of an algorithm that colors any 3-colorable graph with $O(n^c)$ colors. The latter is considered to be hard for small c , and the best known upper bound on the required number of colors is $O(n^{3/14})$, due to Blum and Karger [BK97].

We show that the problem is $\Omega(\log n)$ -*inapproximable* unless $NP \subseteq BPTIME(n^{O(\log \log n)})$ (see Theorem 6.5), showing, in particular, that it is unlikely that the reduction of [B01] may ever be used to prove a polylogarithmic upper bound on the number of colors that are required to color any 3-colorable graph. We note that the *multiplicative* $\Omega(\log n)$ -inapproximability that we establish holds only for the general case of the radio broadcast problem, that is the case when the instances that admit a schedule with only a constant number of rounds are allowed. Indeed, an approximation algorithm with only a *constant multiplicative* ratio, but with an *additive error term* of $O(\log^5 n)$ for the radio broadcast problem was devised in [GM95]. It follows from their result that there exists an algorithm that provides a *constant* approximation guarantee for the restriction of the problem to the set of instances for which any admissible schedule requires $\omega(\log^5 n)$ rounds. On the other hand, it follows from our result that unless $NP \subseteq BPTIME(n^{O(\log \log n)})$, the *additive* $o(\log n)$ approximation for the radio broadcast problem is impossible, even when the problem is

restricted to the set of instances for which all admissible schedules require at least polynomial in n number of rounds.

2 Preliminaries

2.1 The radio broadcast problem

Let $G = (V, E)$ be an unweighted undirected graph. For a vertex $x \in V$, let $\Gamma_G(x) = \{u \in V \mid (x, u) \in E\}$ denote the set of *neighbors* of the vertex x in the graph G .

For a subset $R \subseteq V$ of vertices, the set of vertices that are *informed* by R , denoted $I(R)$, is $I(R) = \{v \mid \exists! x \in R \text{ s.t. } v \in \Gamma_G(x)\}$ (the notation $\exists! x$ stands for “there exists a unique x ”). For a singleton set $R = \{x\}$, $I(R) = I(\{x\}) = I(x) = \Gamma_G(x)$. In other words, a set R informs a vertex v if v has exactly one neighbor in R .

A sequence of vertex sets $\Pi = (R_1, R_2, \dots, R_q)$, $q = 1, 2, \dots$, is called a *radio broadcast schedule* (henceforth, *schedule*) if $R_{i+1} \subseteq \bigcup_{j=1}^i I(R_j)$ for every $i = 1, 2, \dots, q-1$. Intuitively, the vertices of the set R_i send the message on round i . The condition $R_{i+1} \subseteq \bigcup_{j=1}^i I(R_j)$ means that the vertices that send a message in a certain round were informed in one of the previous rounds.

The *vertex set* of the schedule Π , denoted $V(\Pi)$, is $V(\Pi) = \bigcup_{R \in \Pi} R$. The set of vertices *informed by a schedule* Π , denoted $I(\Pi)$, is $I(\Pi) = \bigcup_{R \in \Pi} I(R)$.

An *instance* of the *radio broadcast problem* \mathcal{G} is a pair $(G = (V, E), s)$, where G is a graph, and $s \in V$ is a vertex. Given a graph $G = (V, E)$ and a vertex $s \in V$, a schedule Π is *admissible* with respect to the instance (G, s) , if $R_1 = \{s\}$ and $V = I(\Pi)$. The *length* of the schedule $\Pi = (R_1, R_2, \dots, R_q)$ is $|\Pi| = q$. The goal of the *radio broadcast problem* is to compute an admissible schedule Π of minimal length.

For any schedule $\Pi = (R_1, R_2, \dots, R_q)$, the set R_i is called the *i th round* of Π , $i = 1, 2, \dots, q$.

2.2 The MIN-REP problem

The MIN-REP problem was introduced in [K98]. This problem is a variant of the Label-Cover problem (see, e.g., [AL96]), and it is equivalent to the *symmetric Label-Cover* problem (see, e.g., [DK99]).

Its input consists of a bipartite graph $G = (V_1, V_2, E)$, and two partitions \tilde{V}_1 and \tilde{V}_2 of V_1 and V_2 respectively, into disjoint unions of subsets, $V_j = \bigcup_{X \in \tilde{V}_j} X$, for $j = 1, 2$. Denote $n = |V_1| + |V_2|$.

An instance \mathcal{M} as above induces a bipartite supergraph $\tilde{G} = (\tilde{V}_1, \tilde{V}_2, \tilde{E})$ with $\tilde{E}(\mathcal{M}) = \tilde{E} = \{(A, B) \in \tilde{V}_1 \times \tilde{V}_2 \mid a \in A, b \in B, (a, b) \in E\}$. In other words, the vertices of the supergraph are the sets that belong to one of the partitions, and there is a superedge between a pair of sets (A, B) if and only if the graph G contains an edge between some pair of vertices (a, b) , with $a \in A$ and $b \in B$. It is convenient to visualize the supergraph with the sets $A \in \tilde{V}_1$ on the left-hand side, and the sets $B \in \tilde{V}_2$ on the right-hand side.

Let $\tilde{V} = \tilde{V}_1 \cup \tilde{V}_2$ denote the vertex set of the supergraph.

A pair of vertices $x_1, x_2 \in \tilde{V}_1 \cup \tilde{V}_2$ is called a *matching pair with respect to a superedge* $\tilde{e} = (A, B) \in \tilde{E}$ (henceforth, *\tilde{e} -m.p.*) if $(x_1, x_2) \in \tilde{E} \cap (A \times B)$, where $A \times B$ stands for the set of pairs $\{(x_1, x_2) \mid x_1 \in A, x_2 \in B\}$.

A subset $C \subseteq V_1 \cup V_2$ of vertices is said to *cover* a superedge $\tilde{e} = (A, B)$ if it contains an \tilde{e} -m.p.. A subset C that covers all the superedges is called an *admissible solution* for the instance \mathcal{M} of the MIN-REP problem. A subset $C \subseteq V_1 \cup V_2$ that satisfies $|C \cap X| = 1$ for every $X \in \tilde{V}$ is called a *MAX-cover*.

An instance \mathcal{M} of the MIN-REP problem is called a *YES-instance* if there exists a MAX-cover that covers all the superedges. Such a MAX-cover is called a *proper MAX-cover*.

An instance \mathcal{M} of the MIN-REP problem is called a *t-NO-instance*, for some number $t > 1$, if for any subset $C \subseteq V_1 \cup V_2$ such that $|C| \leq (|\tilde{V}_1| + |\tilde{V}_2|) \cdot t$, there exists a superedge $\tilde{e} \in \tilde{E}$ that is not covered by C .

Intuitively, in a YES-instance there is a way to select one vertex from each supervertex in such a way that the resulting set C of vertices will cover all the superedges. In contrast, in a *t-NO-instance* the set C has to contain at least t vertices per supervertex *on average* in order to cover all the superedges.

We next impose an additional restriction on the instances of the MIN-REP problem. Specifically, for every superedge $\tilde{e} = (A, B) \in \tilde{E}$, for every vertex $b \in B$ there exists exactly one vertex $a \in A$, denoted $\tilde{e}(b)$, such that $(a, b) \in E$. This property is henceforth called the *star property*.

The star property implies that for every superedge $\tilde{e} = (A, B) \in \tilde{E}$, the graph $E(\tilde{e})$ induced by $A \cup B$ is a collection of vertex-disjoint stars (a star is a tree of depth 1). The *head* of the star is the unique vertex that belongs to the intersection of the vertex set of the star with the set A . The other vertices are called the *leaves* of the star. Two vertices $b_1, b_2 \in B$ are said to belong to *different stars* with respect to the superedge \tilde{e} if $\tilde{e}(b_1) \neq \tilde{e}(b_2)$; otherwise, they are said to belong to the *same star* with respect to the superedge \tilde{e} . See Figure 1 for an example of an instance of the MIN-REP problem that obeys the star property. In this example the vertices a and b form a matching pair with respect to the superedge $\tilde{e} = (A_1, B_1)$, but the vertices a' and b do not form a matching pair with respect to the superedge \tilde{e} .

Theorem 2.1 [R98] *No polynomial time algorithm may distinguish between YES-instances and $\log^{10} n$ -NO-instances of the MIN-REP problem unless $NP \subseteq DPTIME(n^{O(\log \log n)})$, even when the instances of the MIN-REP problem satisfy the star property.*

The constant 10 is arbitrary, and its effect is hidden by the O-notation in the complexity-theoretic assumption.

3 The reduction

Next we describe our randomized reduction from the MIN-REP problem to the radio broadcast problem, that shows that the approximability threshold for the latter is $\Omega(\log n)$. This reduction is based on the reduction of Lund and Yannakakis [LY94] from the maximization variant of the Label-Cover problem to the Set-Cover problem (see, e.g., [AL96]).

Consider an instance $\mathcal{M} = (G, \tilde{V}_1, \tilde{V}_2)$, with $G = (V_1, V_2, E)$ of the MIN-REP problem.

The instance $\mathcal{G} = (s, \hat{G})$, $\hat{G} = (\hat{V}, \hat{E})$ of the radio broadcast problem is formed in the following way.




Figure 1: Supervertices are depicted by large circles, and vertices are depicted by small black circles. Edges are depicted by solid lines, and superedges are depicted by dashed lines. Observe that for every superedge (A_i, B_j) , the set $A_i \cup B_j$ forms a collection of vertex-disjoint stars.

The vertex set \hat{V} :

1. **The set \mathcal{V}_1 :** This set contains one *copy* \bar{x} for every vertex $x \in V_1 \cup V_2$. Let $\bar{V} = \bar{V}_1 \cup \bar{V}_2$ denote the set of these copies. Thus, each set A and B in $\tilde{V}_1 \cup \tilde{V}_2$ has a copy set \bar{A} and \bar{B} in \bar{V} . We will also use the notation $A = \text{copy}(\bar{A})$ to denote the relation between the sets A and \bar{A} .
2. **The vertex set \mathcal{V}_2 :** For every superedge $\tilde{e} = (A, B) \in \tilde{E}$, the set \mathcal{V}_2 contains $L = n^2$ *ground sets* $M_{\tilde{e}}^j$ of size $M = n$ each. Let $\hat{M}_{\tilde{e}} = \bigcup_{j=1}^L M_{\tilde{e}}^j$ denote the union of all ground sets that correspond to the same superedge \tilde{e} .
3. **The source s :** The vertex set of the graph \hat{G} contains one additional vertex, designated as a *source*, and denoted s .

Therefore, the vertex set \hat{V} is defined by $\hat{V} = \{s\} \cup \mathcal{V}_1 \cup \mathcal{V}_2$, with $\mathcal{V}_1 = \bar{V}_1 \cup \bar{V}_2$, $\mathcal{V}_2 = \bigcup_{\tilde{e} \in \tilde{E}} \hat{M}_{\tilde{e}}$, $\hat{M}_{\tilde{e}} = \bigcup_{j=1}^L M_{\tilde{e}}^j$, $|M_{\tilde{e}}^j| = n$ for every $\tilde{e} \in \tilde{E}$, $j = 1, 2, \dots, L$.

The edge set \hat{E} :

1. The source s is connected to all of the vertices of \mathcal{V}_1 .

2. For every vertex $\bar{a} \in \bar{A}$ and for every superedge $\tilde{e} \in \tilde{E}$ that is incident to the supervertex A , and for every $j = 1, 2, \dots, L$, the vertex \bar{a} is connected to a random half $H_{\tilde{e}}^j(a)$ of $M_{\tilde{e}}^j$. Specifically, the set $H_{\tilde{e}}^j(a)$ is chosen uniformly at random from the collection of all the subsets H of the set $M_{\tilde{e}}^j$ that have cardinality $|M_{\tilde{e}}^j|/2$, and for different vertices a these choices are independent.
3. All the vertices \bar{b} that satisfy $\tilde{e}(b) = a$ are connected to all the vertices of the complementary half of the set $M_{\tilde{e}}^j$, that is, the set $H_{\tilde{e}}^j(b) = M_{\tilde{e}}^j \setminus H_{\tilde{e}}^j(a)$.
4. In addition, all the vertices \bar{b} that satisfy $\tilde{e}(b) = a$ are connected independently at random to all the vertices of a random half $Q_{\tilde{e}}^j(b)$ of the set $H_{\tilde{e}}^j(a)$. Analogously to the choice of $H_{\tilde{e}}^j(a)$, the set $Q_{\tilde{e}}^j(b)$ is chosen uniformly at random from the collection of all the subsets of $H_{\tilde{e}}^j(a)$ of cardinality $|H_{\tilde{e}}^j(a)|/2$.

Intuitively, every vertex $\bar{a} \in \bar{A}$ “chooses” a random half of the vertices of the ground set $M_{\tilde{e}}^j$ and connects to them. All the leaves of the star of \bar{a} are connected to all the vertices of the complementary half of the ground set. In addition, each leaf in the star of a “chooses” at random half of the neighbors of \bar{a} . Thus, each vertex $\bar{b} \in \bar{B}$ is actually connected to a randomly chosen three quarters of the vertices of the ground-set $M_{\tilde{e}}^j$; to all the non-neighbors of the vertex \bar{a} , and to one half of its neighbors. In contrast, each $\bar{a} \in \bar{A}$ is connected to exactly half of $M_{\tilde{e}}^j$.

See Figure 2 for an illustration. In this example the vertices a, b_1, b_2 belong to the same star in G , with the vertex a serving as the head of the star, and b_1, b_2 serving as its leaves. This star is a part of the subgraph induced by the superedge \tilde{e} . The set $M_{\tilde{e}}^j$ in this example consists of eight vertices, v_1, v_2, \dots, v_8 . The vertex a is connected to all the vertices of a random half of the set $M_{\tilde{e}}^j$, and in this example this random half consists of the vertices v_1, v_2, v_3 and v_4 . Both vertices b_1 and b_2 are connected to all the vertices of the complementary half of the set $M_{\tilde{e}}^j$, that is, to the vertices v_5, v_6, v_7 and v_8 . In addition, the vertex b_1 is connected to all the vertices of a random half of the set of the neighbors of a , and in this example this half consists of the vertices v_1 and v_2 . The vertex b_2 is also connected to all the vertices of a random half of the set of neighbors of a , and in this example this half consists of the vertices v_1 and v_3 .

Let $m = O(n^2)$ denote the number of superedges in \tilde{G} . Let $\mathcal{V} = \mathcal{V}_1 \cup \mathcal{V}_2 = \hat{V} \setminus \{s\}$. The number of vertices in the graph \hat{G} is $|\hat{V}| = 1 + |\mathcal{V}| = 1 + n + m \cdot L \cdot n = O(n^5)$.

We remark that the bipartite graph $(\mathcal{V}_1, \mathcal{V}_2, \tilde{E})$ is similar to the graph that is constructed by the reduction of Lund and Yannakakis [LY94]. The latter reduction is used there to prove the logarithmic inapproximability of the Set-Cover problem. However, there are two differences. First, in [LY94] only one set $\hat{M}_{\tilde{e}}$ is formed for every superedge \tilde{e} , instead of L copies of it, as it is in our reduction. Second, and probably more important difference, is the introduction of the sets $Q_{\tilde{e}}^j(b)$. In the reduction to the Set-Cover problem for every $b \in V_2$, the vertex \bar{b} is connected in $\hat{M}_{\tilde{e}}$ only to the vertices of $H_{\tilde{e}}(b)$. In our reduction, in addition to $H_{\tilde{e}}(b)$, the vertex \bar{b} is connected to a random half $Q_{\tilde{e}}(b)$ of the set $H_{\tilde{e}}(a)$.

Remark about Notation: We use bars to denote copies of the vertices of the original graph, tildes to denote supervertices and superedges of the supergraph \tilde{G} that is a part of the instance \mathcal{M} of the MIN-REP problem, and hats for the vertex sets and edgesets of the graph \hat{G} that is returned by our reduction.




Figure 2: The dashed lines correspond to the edges of the graph G . The solid lines correspond to the edges of the graph \hat{G} . The wide solid lines indicate various partitions of the set $M_{\tilde{e}}^j = \{v_1, \dots, v_8\}$.

4 Intuition: YES-instance

To provide some intuition of the reduction, we next analyze its behavior on a YES-instance \mathcal{M} of the MIN-REP problem.

Consider a YES-instance $\mathcal{M} = (G, \tilde{V}_1, \tilde{V}_2)$, $G = (V_1, V_2, E)$ of the MIN-REP problem, and the instance $\mathcal{G} = (s, \hat{G})$ of the radio broadcast problem that is obtained out of \mathcal{M} via our reduction. We show that there is a short schedule for \mathcal{G} .

Lemma 4.1 *There exists (with probability 1) an admissible schedule Π for \mathcal{G} of length 3.*

Proof: As \mathcal{M} is a YES-instance of the MIN-REP problem, there exists a proper MAX-cover C for \mathcal{M} . Recall that such a cover contains exactly one vertex from every supervertex from $\tilde{V}_1 \cup \tilde{V}_2$.

We next show that the schedule $\Pi = (R_1 = \{s\}, R_2 = \bar{C} \cap \bar{V}_1, R_3 = \bar{C} \cap \bar{V}_2)$ satisfies the requirements of the lemma.

First, $I(\{s\}) = \Gamma_{\hat{G}}(s) = \mathcal{V}_1$. In other words, all the vertices of \mathcal{V}_1 are informed after the first round. It remains to argue that every vertex $v \in \mathcal{V}_2$ is informed either by the second or by the third rounds of the schedule Π .

As $\mathcal{V}_2 = \bigcup_{\tilde{e} \in \tilde{E}} \hat{M}_{\tilde{e}}$, it follows that there exists a superedge $\tilde{e} = (A, B) \in \tilde{E}$ such that $v \in \hat{M}_{\tilde{e}} = \bigcup_{j=1}^L M_{\tilde{e}}^j$. Hence $v \in M_{\tilde{e}}^j$ for some $j = 1, 2, \dots, L$. As C is a proper MAX-cover for \mathcal{M} , it contains some \tilde{e} -m.p. $a \in C \cap A, b \in C \cap B$ such that $(a, b) \in E$. Note that as $(a, b) \in E, a \in A, b \in B$, it follows that $a = \tilde{e}(b)$. Hence $M_{\tilde{e}}^j = H_{\tilde{e}}^j(a) \cup H_{\tilde{e}}^j(b)$.

Since $R_2 = \bar{C} \cap \bar{V}_1$, and C is a MAX-cover, the vertex \bar{a} is the only vertex of $\bar{A} \cup \bar{B}$ that transmits on round 2, and, analogously, since $R_3 = \bar{C} \cap \bar{V}_2$, the vertex \bar{b} is the only vertex of $\bar{A} \cup \bar{B}$ that transmits on round 3. Therefore, no collision occurs, and $M_{\tilde{e}}^j = H_{\tilde{e}}^j(a) \cup H_{\tilde{e}}^j(b) = I(\{\bar{a}\}) \cup I(\{\bar{b}\}) \subseteq I(R_2) \cup I(R_3) \subseteq I(\Pi)$, completing the proof. ■

Remark: Note that since the vertices of \bar{V}_1 and \bar{V}_2 share many common neighbors in \mathcal{V}_2 , it may happen that no admissible schedule with less than three rounds exists.

5 Analysis of the construction

All the probabilities in this section are with respect to the sample space that is determined by all possible random choices of the neighborhoods of the vertices of the vertex set \bar{V} of the graph \hat{G} .

Notation: In the sequel, for simplicity of notation, we say that $x \mathcal{N} v$, if the vertices x and v are not connected, that is, if $(x, v) \notin \hat{E}$. Similarly, we say that $X \mathcal{AN} v$ if no vertex in the set X is connected to v . We also say that $X \mathcal{AC} v$ if *all* the vertices of the set X are connected to the vertex v . Finally, we say that $x \mathcal{C} v$ if x and v are connected, that is, if $(x, v) \in \hat{E}$.

Additionally, for a set $B' \subseteq \mathcal{V}_1$, let $heads(B')$ be the set of all star heads of B' , namely,

$$heads(B') = \{\bar{a} \mid \bar{b} \in B', a = \tilde{e}(b)\}.$$

Remark: Throughout the sequel, we assume that n is sufficiently large whenever necessary. All the logarithms are over the base 2 unless specified otherwise.

5.1 Positive and negative witnesses of a schedule

Consider a schedule $\Pi = (\tilde{R}_1, \tilde{R}_2, \dots, \tilde{R}_q)$, and let $R_i = \tilde{R}_i \cap (\bar{A} \cup \bar{B})$ (recall that the superedge $\tilde{e} = (A, B)$ is fixed). In what follows we derive a sufficient condition that guarantees that the schedule Π *does not* inform the vertex v .

To this end we define two disjoint subsets $N_\Pi, P_\Pi \subseteq \bar{A} \cup \bar{B}$ of elements called the sets of *negative and positive witnesses* for the schedule Π , respectively. We then show that if a vertex v is connected to *no* negative witness, but *is* connected to *all* the positive witnesses, then the vertex v is necessarily not informed by the schedule Π . Furthermore, we show that for any schedule Π , there exist relatively small witness sets.

The set N_Π of negative witnesses is defined as follows. If the schedule Π contains no round R that contains only one single element (i.e., $R = \{x\}$ for some vertex x ; such a round is called a *singleton* round), then the set N_Π of negative witnesses is empty. Otherwise, the set N_Π is defined by Procedure NEG_WITNESS described below.

Procedure NEG_WITNESS

Input: schedule Π ;

1. Set $\Pi' \leftarrow \Pi$;
2. **Do**
 Choose a singleton round $R = \{x\}$ of the schedule Π' , and add the vertex x into N_Π ; Remove all the occurrences of the vertex x from the schedule Π' ;
Until Π' contains no singleton rounds.
3. **Return**(N_Π, Π');

Let Π' be the schedule that is returned by Procedure NEG_WITNESS. Observe that every round in this schedule contains at least two vertices. Choose two arbitrary vertices x_i, y_i , from each non-empty round i and add them to the set P_Π of positive witnesses.

In the next lemma we analyze the properties of the sets of positive and negative witnesses.

Lemma 5.1

1. $N_\Pi \cap P_\Pi = \emptyset$.
2. $|N_\Pi| \leq |\Pi| = q$.
3. $|P_\Pi| \leq 2 \cdot q$.
4. If $(N_\Pi \mathcal{AN} v)$ and $(P_\Pi \mathcal{AC} v)$, that is, none of the elements in N_Π is connected to v but all the elements in P_Π are connected to v , then the vertex v is not informed by Π .

Proof: Claim 1 is immediate. For Claims 2 and 3 note that each time an element x is added to the set N_Π , a singleton round $R = \{x\}$ of the schedule Π' becomes empty. Therefore, the size of N_Π is at most q . Note also that $|\Pi'| \leq |\Pi| = q$. Hence, $|P_\Pi| = 2 \cdot |\Pi'| \leq 2q$.

To prove Claim 4, let Π_0 denote the sub-schedule of the schedule Π that contains only rounds that were emptied during the execution of Procedure `NEG_WITNESS`, and, consequently, do not appear in the schedule Π' . Let $\Pi'' = \Pi \setminus \Pi_0$ denote the complementary subschedule of Π . Note also that for every round R in Π_0 , $R \subseteq N_\Pi$. Hence, the condition $(N_\Pi \mathcal{AN} v)$ implies that the subschedule Π_0 does not inform the vertex v .

In addition, recall that any round in the schedule Π' contains at least two vertices that belong to the set P_Π of positive witnesses. Assuming that all the positive witnesses are connected to v , it follows that every round of the schedule Π' contains at least two elements that are connected to v . Hence, under this assumption, the schedule Π' does not inform the vertex v , and, therefore, the subschedule Π'' informs v neither. Since neither Π_0 nor Π'' inform v , it follows that the schedule Π does not inform v , proving the claim. \blacksquare

5.2 Basic properties

In this section we state and prove several basic properties of our reduction.

Rounds $i = 2, 3, \dots$: Observe that after round 1, all the vertices of the set \mathcal{V}_1 are informed. Hence, without loss of generality, assume that $R_1 = \{s\}$, and for $i = 2, 3, \dots$, $R_i \subseteq \mathcal{V}_1$.

Dependencies in the construction: Observe that the events $\{x \mathcal{C} v\}$ and $\{y \mathcal{C} v\}$ for two vertices $x, y \in \mathcal{V}_1$ and a vertex $v \in \mathcal{V}_2$ are not necessarily independent. For example, if the vertex $x = \bar{a}$ is the head of a star, and the vertex $y = \bar{b}$ is a leaf of the same star, then if the head x is not connected to v , then the leaf y is necessarily connected to v .

On the other hand, if two vertices \bar{b}_1, \bar{b}_2 belong to different stars in G , by construction, all the events that involve \bar{b}_1 and \bar{b}_2 are independent. Events that involve \bar{b} depend only on the events that involve the head of the star of b , and the other leaves in the same star as b .

Finally, suppose that the vertices \bar{b}' and \bar{b}'' belong to the same star (i.e., $a' = \tilde{e}(b') = \tilde{e}(b'')$). If the edge (\bar{a}', v) belongs to \hat{E} , then the events $\{\bar{b}' \mathcal{C} v\}$ and $\{\bar{b}'' \mathcal{C} v\}$ are *independent*. (This follows from Step 4 of the construction of the edge set \hat{E} .) To summarize,

Corollary 5.2 Fix a superedge $\tilde{e} = (A, B)$, an index $j = 1, 2, \dots, L$, and a vertex $v \in M_{\tilde{e}}^j$.

1. Suppose that the vertices $b_1, b_2 \in B$ belong to two different stars in $E(\tilde{e})$. Then the events $\{\bar{b}_1 \mathcal{C} v\}$ and $\{\bar{b}_2 \mathcal{C} v\}$ are independent.

2. Let $a_1, a_2 \in A$. The events $\{\bar{a}_1 \mathcal{C} v\}$ and $\{\bar{a}_2 \mathcal{C} v\}$ are independent.

3. **The separation principle:** Let $B' \subseteq B$ and suppose that the event $\{\text{heads}(B') \mathcal{AN} v\}$ occurs. Then each vertex in B' is connected or not connected to v independently with probability $1/2$.

Definition 5.3 For a superedge $\tilde{e} = (A, B)$, and an index $j = 1, 2, \dots, L$, and (\tilde{e}, j) -history H is a set of events of the form $\{x \mathcal{C} v\}$ or $\{x \mathcal{N} v\}$, where $x \in \bar{A} \cup \bar{B}$, and $v \in M_{\tilde{e}}^j$. The (\tilde{e}, j) -history H is said to be independent of the subset $C \subseteq \bar{A} \cup \bar{B}$, if there exists no pair of vertices $x \in C$, $v \in M_{\tilde{e}}^j$ such that $\{x \mathcal{C} v\} \in H$ or $\{x \mathcal{N} v\} \in H$.

In the next lemma we analyze the probability that all the vertices of a given subset are connected to a fixed vertex v .

Lemma 5.4 Fix a superedge $\tilde{e} = (A, B)$, and an index $j = 1, 2, \dots, L$. Let $B' \subseteq B$ be a subset of vertices, and H be an (\tilde{e}, j) -history that is independent of the subset B' . Let v be a vertex in $M_{\tilde{e}}^j$. Then $\mathbb{P}(\bar{B}' \mathcal{AC} v \mid H) \geq \frac{1}{2^{|B'|}}$, and $\mathbb{P}(\bar{B}' \mathcal{AN} v \mid H) \leq \frac{1}{2^{|B'|}}$.

Proof: Fix some vertex $b \in B'$. Let $a = \tilde{e}(b)$. By construction, if \bar{a} is not connected to v then \bar{b} is connected to v with probability 1. Otherwise, by the separation principle, regardless of the history, the vertex \bar{b} is connected to v with probability $1/2$. Hence, $\mathbb{P}(\bar{b} \mathcal{C} v \mid H) \geq 1/2$ and $\mathbb{P}(\bar{b} \mathcal{N} v \mid H) \leq 1/2$.

Now, the lemma follows from parts 1 and 3 of Corollary 5.2. ■

We next analyze the probability that when all the vertices of a fixed set X transmit in the same round, a fixed vertex v becomes informed. The following lemma provides an upper bound on the probability of this event. This upper bound is useful mostly when the set X has large cardinality.

Lemma 5.5 For a subset $X \subseteq A \cup B$, an index $j = 1, 2, \dots, L$, a superedge $\tilde{e} = (A, B)$, and a vertex $v \in M_{\tilde{e}}^j$,

$$\mathbb{P}(v \in I(\bar{X})) \leq \frac{|X|}{2^{|X|-1}}.$$

Proof: Recall that $v \in I(\bar{X})$ if and only if there exists a unique vertex $\bar{x} \in \bar{X}$ that is connected to v . For a fixed vertex $\bar{x} \in \bar{X}$, let $X_A = (\bar{X} \setminus \{\bar{x}\}) \cap \bar{A}$, and $X_B = (\bar{X} \setminus \{\bar{x}\}) \cap \bar{B}$. Note that $X_A \cup X_B = \bar{X} \setminus \{\bar{x}\}$, and $X_A \cap X_B = \emptyset$. By Bayes formula,

$$\begin{aligned} \mathbb{P}((\bar{x} \mathcal{C} v) \text{ and } ((\bar{X} \setminus \{\bar{x}\}) \mathcal{AN} v)) &\leq \mathbb{P}((X_B \cup X_A) \mathcal{AN} v) = \\ \mathbb{P}((X_B \mathcal{AN} v) \text{ and } (X_A \mathcal{AN} v)) &= \mathbb{P}((X_B \mathcal{AN} v) \mid (X_A \mathcal{AN} v)) \cdot \mathbb{P}(X_A \mathcal{AN} v). \end{aligned}$$

Observe that $\mathbb{P}(X_A \mathcal{AN} v) = \frac{1}{2^{|X_A|}}$. Let H denote the singleton set $\{(X_A \mathcal{AN} v)\}$. Note that the (\tilde{e}, j) -history H is independent of the subset X_B . Thus, by Lemma 5.4,

$$\mathbb{P}(X_B \mathcal{AN} v \mid H) = \mathbb{P}((X_B \mathcal{AN} v) \mid (X_A \mathcal{AN} v)) \leq \frac{1}{2^{|X_B|}}.$$

Hence, altogether,

$$\mathbb{P}((x \mathcal{C} v) \text{ and } ((\bar{X} \setminus \{\bar{x}\}) \mathcal{AN} v)) \leq \frac{1}{2^{|X_A|+|X_B|}} = \frac{1}{2^{|X|-1}}.$$

The statement of the lemma now follows by union-bound. ■

A basic lemma: As was discussed above, we are going to define a set P of positive witnesses and a set N of negative witnesses, and will be interested in the probability of the event $\{(P \mathcal{A} \mathcal{C} v) \text{ and } (N \mathcal{A} \mathcal{N} v)\}$.

Lemma 5.6 *Let $N, P \subseteq \bar{A} \cup \bar{B}$, $N \cap P = \emptyset$ be two disjoint subsets of $\bar{A} \cup \bar{B}$. Suppose also that $N \cup P$ contains no \tilde{e} -m.p.. Then*

$$\mathbb{P}((N \mathcal{A} \mathcal{N} v) \text{ and } (P \mathcal{A} \mathcal{C} v)) \geq \frac{1}{4^{|N|+|P|}}.$$

Proof: Let $N_A = N \cap \bar{A}$, $N_B = N \cap \bar{B}$, $P_A = P \cap \bar{A}$ and $P_B = P \cap \bar{B}$. Let $A' = \text{heads}(N_B \cup P_B)$. Observe that as $N \cup P$ contains no \tilde{e} -m.p., it follows that $(P_A \cup N_A) \cap A' = \emptyset$. Note that by Corollary 5.2(2), $\mathbb{P}(A' \mathcal{A} \mathcal{N} v) = \frac{1}{2^{|A'|}}$. Hence,

$$\begin{aligned} \mathbb{P}((N \mathcal{A} \mathcal{N} v) \text{ and } (P \mathcal{A} \mathcal{C} v)) &\geq \mathbb{P}((N \mathcal{A} \mathcal{N} v) \text{ and } (P \mathcal{A} \mathcal{C} v) \text{ and } (A' \mathcal{A} \mathcal{N} v)) \\ &= \mathbb{P}((N \mathcal{A} \mathcal{N} v) \text{ and } (P \mathcal{A} \mathcal{C} v) \mid (A' \mathcal{A} \mathcal{N} v)) \cdot \mathbb{P}(A' \mathcal{A} \mathcal{N} v) \\ &= \frac{\mathbb{P}((N \mathcal{A} \mathcal{N} v) \text{ and } (P \mathcal{A} \mathcal{C} v) \mid (A' \mathcal{A} \mathcal{N} v))}{2^{|A'|}}. \end{aligned}$$

Note that $\text{heads}(P_B) \subseteq A'$, and so the condition $(A' \mathcal{A} \mathcal{N} v)$ implies the condition $(\text{heads}(P_B) \mathcal{A} \mathcal{N} v)$. In turn, by the construction, the latter condition implies the condition $(P_B \mathcal{A} \mathcal{C} v)$. Recall also that $P_A \cup P_B = P$. Hence $\mathbb{P}(((N \mathcal{A} \mathcal{N} v) \text{ and } (P \mathcal{A} \mathcal{C} v)) \mid (A' \mathcal{A} \mathcal{N} v)) = \mathbb{P}(((N \mathcal{A} \mathcal{N} v) \text{ and } (P_A \mathcal{A} \mathcal{C} v)) \mid (A' \mathcal{A} \mathcal{N} v))$.

Recall that $P_A \cap A' = \emptyset$, $P_A \cap N = \emptyset$, and that the set $P_A \cup N$ contains no \tilde{e} -m.p.. Hence,

$$\begin{aligned} \mathbb{P}(((N \mathcal{A} \mathcal{N} v) \text{ and } (P_A \mathcal{A} \mathcal{C} v)) \mid (A' \mathcal{A} \mathcal{N} v)) &= \mathbb{P}((N \mathcal{A} \mathcal{N} v) \mid (A' \mathcal{A} \mathcal{N} v)) \cdot \mathbb{P}(P_A \mathcal{A} \mathcal{C} v) \\ &= \frac{\mathbb{P}((N \mathcal{A} \mathcal{N} v) \mid (A' \mathcal{A} \mathcal{N} v))}{2^{|P_A|}}. \end{aligned}$$

Since $N = N_A \cup N_B$, and $N_A \cap A' = \emptyset$, and $N_A \cup A' \subseteq \bar{A}$ (and, therefore, contains no \tilde{e} -m.p.), it follows that

$$\begin{aligned} \mathbb{P}((N_A \mathcal{A} \mathcal{N} v) \text{ and } (N_B \mathcal{A} \mathcal{N} v) \mid (A' \mathcal{A} \mathcal{N} v)) &= \mathbb{P}((N_B \mathcal{A} \mathcal{N} v) \mid (A' \mathcal{A} \mathcal{N} v)) \cdot \mathbb{P}(N_A \mathcal{A} \mathcal{N} v) \\ &= \frac{\mathbb{P}((N_B \mathcal{A} \mathcal{N} v) \mid (A' \mathcal{A} \mathcal{N} v))}{2^{|N_A|}}. \end{aligned}$$

Hence

$$\frac{\mathbb{P}(((N \mathcal{A} \mathcal{N} v) \text{ and } (P_A \mathcal{A} \mathcal{C} v)) \mid (A' \mathcal{A} \mathcal{N} v))}{2^{|A'|}} = \frac{\mathbb{P}((N_B \mathcal{A} \mathcal{N} v) \mid (A' \mathcal{A} \mathcal{N} v))}{2^{|A'|+|N_A|+|P_A|}}.$$

Finally, by the separation principle (Corollary 5.2(3)), $\mathbb{P}((N_B \mathcal{A} \mathcal{N} v) \mid (A' \mathcal{A} \mathcal{N} v)) = \frac{1}{2^{|N_B|}}$. It follows that

$$\mathbb{P}((N \mathcal{A} \mathcal{N} v) \text{ and } (P \mathcal{A} \mathcal{C} v)) \geq \frac{1}{2^{|A'|+|N_A|+|P_A|+|N_B|}}.$$

Since $|A'| = |\text{heads}(P_B \cup N_B)| \leq |P_B| + |N_B|$, it follows that $|A'| + |N_A| + |P_A| + |N_B| \leq 2(|P| + |N|)$, completing the proof. \blacksquare

5.3 Main lemmas

In this section we analyze the probabilities of existence of schedules that inform all the vertices of $\hat{M}_{\tilde{e}}$, and prove several lemmas that will be most useful in the analysis of our reduction.

Lemma 5.7 *For a superedge $\tilde{e} \in \tilde{E}$, let Π be a schedule of length $q \leq \log n/30$, such that the set $V(\Pi)$ contains no \tilde{e} -m.p.. For a vertex $v \in \hat{M}_{\tilde{e}}$, the probability that the schedule Π does not inform v is at least $1/n^{1/5}$.*

Proof: Let P_Π and N_Π be the sets of positive and negative witnesses of Π , respectively. By Lemmas 5.6 and 5.1, the probability for that Π does not inform v is at least

$$\left(\frac{1}{4}\right)^{|N_\Pi|+|P_\Pi|} \geq \left(\frac{1}{4}\right)^{3q} \geq n^{1/5},$$

proving the claim. \blacksquare

Definition 5.8 *For a superedge $\tilde{e} = (A, B)$, a round R in a schedule Π is called large (resp., small) with respect to \tilde{e} , if $|R \cap (\bar{A} \cup \bar{B})| > 20 \log n$, (resp., $|R \cap (\bar{A} \cup \bar{B})| \leq 20 \log n$). The subschedule of the schedule Π that contains only large (resp., small) rounds with respect to the superedge \tilde{e} is denoted $L(\Pi, \tilde{e})$ (resp., $S(\Pi, \tilde{e})$). When the second parameter \tilde{e} can be deduced from the context, we omit it from this notation.*

(The constant 20 in the above definition is somewhat arbitrary.) The next lemma states that large rounds are unlikely to be of any help to the broadcast.

Lemma 5.9 *For a schedule Π of length $O(\log n)$, a superedge $\tilde{e} \in \tilde{E}$, and a vertex $v \in \hat{M}_{\tilde{e}}$, $\mathbb{P}(v \in I(L(\Pi, \tilde{e}))) = O(\frac{\log n}{n^{19}})$.*

Proof: Let R be a round in the schedule $L(\Pi) = L(\Pi, \tilde{e})$. By Lemma 5.5, the probability that the vertex v is informed by R is at most $\frac{|R|}{2^{|R|-1}} < \frac{n}{2^{20 \cdot \log n}} = \frac{1}{n^{19}}$. Since $|L(\Pi)| \leq |\Pi| = O(\log n)$, by union-bound, the probability that v belongs to $I(L(\Pi))$ is $O(\frac{\log n}{n^{19}})$. \blacksquare

For a superedge \tilde{e} , a schedule Π is called \tilde{e} -proper if the vertex set $V(S(\Pi, \tilde{e}))$ contains no \tilde{e} -m.p..

Lemma 5.10 *For a superedge \tilde{e} , and an \tilde{e} -proper schedule Π of length at most $\log_2 n/30$, and a vertex $v \in \hat{M}_{\tilde{e}}$, $\mathbb{P}(v \notin I(\Pi)) \geq \frac{1/2}{n^{1/5}}$.*

Proof: Note that $\mathbb{P}(v \in I(\Pi)) = \mathbb{P}(v \in I(L(\Pi))) + \mathbb{P}(v \in I(S(\Pi)))$. By Lemma 5.9, $\mathbb{P}(v \in I(L(\Pi))) = O(\frac{\log n}{n^{19}})$. Since the set $V(S(\Pi))$ contains no \tilde{e} -m.p., Lemma 5.7 is applicable. Hence $\mathbb{P}(v \in I(S(\Pi))) \leq 1 - \frac{1}{n^{1/5}}$. Hence $\mathbb{P}(v \in I(\Pi)) \leq 1 - \frac{1}{n^{1/5}} + O(\frac{\log n}{n^{19}}) \leq 1 - \frac{1/2}{n^{1/5}}$. \blacksquare

We next compute the probability that a fixed \tilde{e} -proper schedule Π informs all the vertices of the set $\hat{M}_{\tilde{e}}$.

Lemma 5.11 *Let Π be an \tilde{e} -proper schedule of length at most $\log_2 n/30$, and $j = 1, 2, \dots, L$ be an index. Then*

$$\mathbb{P}(M_{\tilde{e}}^j \subseteq I(\Pi)) \leq 1 - \frac{1/2}{n^{1/5}}.$$

Proof: For a fixed vertex $v \in M_{\tilde{e}}^j$, by Lemma 5.10, $\mathbb{P}(v \in I(\Pi)) \leq 1 - 1/(2 \cdot n^{1/5})$. Observe that the event $(M_{\tilde{e}}^j \subseteq I(\Pi))$ is contained in the event $(v \in I(\Pi))$, and, thus, the lemma follows. \blacksquare

We conclude that

Lemma 5.12 *The probability that a fixed \tilde{e} -proper Π of length at most $\log_2 n/30$ informs all the vertices of $\hat{M}_{\tilde{e}}$ is at most $\exp(-n^{9/5})$.*

Proof: Recall that $\hat{M}_{\tilde{e}} = \cup_j M_{\tilde{e}}^j$. By Lemma 5.11, and the fact that events $\{M_{\tilde{e}}^j \subseteq I(\Pi)\}$ and $\{M_{\tilde{e}}^{j'} \subseteq I(\Pi)\}$ for $j \neq j'$ are independent, it follows that

$$\mathbb{P}(M_{\tilde{e}} \subseteq I(\Pi)) \leq \left(1 - \frac{1}{n^{1/5}}\right)^{n^2} = O(\exp(-n^{9/5})). \quad \blacksquare$$

We next strengthen the previous lemma, and show that not only a fixed \tilde{e} -proper schedule is unlikely to inform all the vertices of the set $\hat{M}_{\tilde{e}}$, but that with high probability there is no \tilde{e} -proper schedule that informs all the vertices of the set $\hat{M}_{\tilde{e}}$. In other words, even if an \tilde{e} -proper schedule Π is chosen by the adversary, it is still unlikely to be able to inform all the vertices of the set $\hat{M}_{\tilde{e}}$.

Definition 5.13 *For a subset $R \subseteq \mathcal{V}_1$, and a superedge $\tilde{e} = (A, B) \in \tilde{E}$, let $R|_{\tilde{e}} = R \cap (\bar{A} \cup \bar{B})$ denote the projection of the round R on the superedge \tilde{e} . For a schedule $\Pi = (R_1, R_2, \dots, R_q)$, and a superedge \tilde{e} , let $\Pi|_{\tilde{e}} = (R_1|_{\tilde{e}}, R_2|_{\tilde{e}}, \dots, R_q|_{\tilde{e}})$ denote the projection of the schedule Π on the superedge \tilde{e} . For a collection $\hat{\Pi}$ of schedules, and a superedge \tilde{e} , let $\hat{\Pi}|_{\tilde{e}} = \{\Pi|_{\tilde{e}} \mid \Pi \in \hat{\Pi}\}$ denote the projection of the collection $\hat{\Pi}$ on the superedge \tilde{e} .*

Let $\hat{\Pi}(\tilde{e})$ be the collection of all \tilde{e} -proper schedules of length at most $O(\log n)$.

Note that for a schedule $\Pi \in \hat{\Pi}(\tilde{e})$, and a round $R \in \Pi$, the projection $R|_{\tilde{e}}$ is a subset of the set $\bar{A} \cup \bar{B}$, and $|\bar{A} \cup \bar{B}| \leq n$. Thus,

$$|\{R|_{\tilde{e}} \mid R \in \Pi, \Pi \in \hat{\Pi}(\tilde{e})\}| \leq 2^{|\bar{A} \cup \bar{B}|} \leq 2^n.$$

Since for every schedule $\Pi \in \hat{\Pi}(\tilde{e})$, $|\Pi| = O(\log n)$, it follows that

$$|(\hat{\Pi}(\tilde{e}))|_{\tilde{e}}| \leq 2^{O(n \cdot \log n)}. \quad (1)$$

Lemma 5.14 *Let $\tilde{e} \in \tilde{E}$ be a superedge, and $\hat{\Pi}(\tilde{e})$ be a collection of schedules defined above. Then*

$$\mathbb{P}(\exists \Pi \in \hat{\Pi} \text{ s.t. } \hat{M}_{\tilde{e}} \subseteq I(\Pi)) \leq 2^{-\Omega(n^{9/5})}.$$

Proof: Note that

$$\mathbb{P}(\exists \Pi \in \hat{\Pi}(\tilde{e}) \text{ s.t. } \hat{M}_{\tilde{e}} \subseteq I(\Pi)) = \mathbb{P}(\exists \Pi \in (\hat{\Pi}(\tilde{e}))|_{\tilde{e}} \text{ s.t. } \hat{M}_{\tilde{e}} \subseteq I(\Pi))$$

By union-bound,

$$\begin{aligned} \mathbb{P}(\exists \Pi \in (\hat{\Pi}(\tilde{e}))|_{\tilde{e}} \text{ s.t. } \hat{M}_{\tilde{e}} \subseteq I(\Pi)) &\leq |(\hat{\Pi}(\tilde{e}))|_{\tilde{e}}| \cdot \max_{\Pi \in (\hat{\Pi}(\tilde{e}))|_{\tilde{e}}} \left\{ \mathbb{P}(\hat{M}_{\tilde{e}} \subseteq I(\Pi)) \right\} \\ &\leq |(\hat{\Pi}(\tilde{e}))|_{\tilde{e}}| \cdot \max_{\Pi \in \hat{\Pi}(\tilde{e})} \left\{ \mathbb{P}(\hat{M}_{\tilde{e}} \subseteq I(\Pi)) \right\}. \end{aligned} \quad (2)$$

By Lemma 5.12, for every $\Pi \in \hat{\Pi}(\tilde{e})$, $\mathbb{P}(\hat{M}_{\tilde{e}} \subseteq I(\Pi)) \leq e^{-n^{9/5}}$. The statement of the lemma follows now from (1) and (2). ■

Next we argue that the probability that there exists an admissible short schedule Π and a superedge \tilde{e} such that Π is an \tilde{e} -proper schedule is very small. Formally, let $\hat{\Pi} = \bigcup_{\tilde{e} \in \tilde{E}} \hat{\Pi}(\tilde{e})$.

Corollary 5.15 *With probability at most $2^{-\Omega(n^{9/5})}$, the collection $\hat{\Pi}$ contains no admissible schedules for the instance \mathcal{G} of the radio broadcast problem. I.e., $\mathbb{P}(\exists \Pi \in \hat{\Pi} \text{ s.t. } \mathcal{V}_2 \subseteq I(\Pi)) \leq 2^{-\Omega(n^{9/5})}$.*

Proof: By union-bound, and since for every superedge \tilde{e} , $\hat{M}_{\tilde{e}} \subseteq \mathcal{V}_2$,

$$\begin{aligned} \mathbb{P}(\exists \Pi \in \hat{\Pi} \text{ s.t. } \mathcal{V}_2 \subseteq I(\Pi)) &\leq \sum_{\tilde{e} \in \tilde{E}} \mathbb{P}(\exists \Pi \in \hat{\Pi}(\tilde{e}) \text{ s.t. } \mathcal{V}_2 \subseteq I(\Pi)) \\ &\leq \sum_{\tilde{e} \in \tilde{E}} \mathbb{P}(\exists \Pi \in \hat{\Pi}(\tilde{e}) \text{ s.t. } \hat{M}_{\tilde{e}} \subseteq I(\Pi)) . \end{aligned}$$

By Lemma 5.14, each term of this sum is at most $2^{-\Omega(n^{9/5})}$. The lemma follows since there are $|\tilde{E}| = O(n^2)$ terms. ■

6 Analysis of the reduction

The case when the original instance \mathcal{M} of the MIN-REP problem is a YES-instance was analyzed in Section 4. We next analyze the case when \mathcal{M} is a NO-instance.

Consider a NO-instance \mathcal{M} of the MIN-REP problem, and the instance \mathcal{G} of the radio broadcast problem that is obtained from \mathcal{M} via our reduction. In this section we show that with high probability there is no admissible short schedule for \mathcal{G} .

Definition 6.1 *For a supervertex $X \in \tilde{V}_1 \cup \tilde{V}_2$, a round R of a schedule Π is said to be important with respect to the supervertex X , if it is small with respect to some superedge $\tilde{e} = (X, Y)$ that is adjacent to the supervertex X . Let $\text{Imp}_{\Pi}(X)$ denote the set of rounds of Π that are important with respect to the supervertex X .*

For a supervertex X and a schedule Π , the subcover $C_{X,\Pi}$ of X that is induced by Π is defined by $C_{X,\Pi} = \bigcup_{R \in \text{Imp}_{\Pi}(X)} \text{copy}(R) \cap \tilde{X}$. The cover C_{Π} of \mathcal{M} that is induced by the schedule Π is the union of the subcovers of all the supervertices, i.e., $C_{\Pi} = \bigcup_{X \in \tilde{V}_1 \cup \tilde{V}_2} C_{X,\Pi}$.

Lemma 6.2 *With probability at least $1 - 2^{-\Omega(n^{9/5})}$, for every admissible schedule Π of length at most $\frac{\log n}{30}$ for the instance \mathcal{G} of the radio broadcast problem, the cover C_{Π} induced by the schedule Π is an admissible solution for the instance \mathcal{M} of the MIN-REP problem.*

Proof: Consider a set Φ of admissible schedules for \mathcal{G} of length at most $\frac{\log n}{30}$. Note that Φ is a random variable. Consider the set Ψ of schedules Π of length at most $\frac{\log n}{30}$ that satisfy that C_{Π} is an admissible solution for \mathcal{M} . Observe that the set Ψ is fixed, and does not depend on the random coins. Our objective is to prove that $\mathbb{P}(\Phi \subseteq \Psi) \geq 1 - 2^{-\Omega(n^{9/5})}$. Let Υ denote the set of all possible schedules Π . Consider a schedule Π that belongs to $\Upsilon \setminus \Psi$. Then there exists a superedge $\tilde{e}_{\Pi} = \tilde{e} = (A, B)$ that is not covered by the induced cover C_{Π} . In other words, the set $C_{A,\Pi} \cup C_{B,\Pi}$

contains no \tilde{e} -m.p.. By definition, $C_{A,\Pi} \cup C_{B,\Pi} = (A \cup B) \cap \bigcap_{R \in (\text{Imp}_\Pi(A) \cup \text{Imp}_\Pi(B))} \text{copy}(R)$. Note that the set $(\text{Imp}_\Pi(A) \cup \text{Imp}_\Pi(B))$ contains all the rounds of Π that are small with respect to the superedge $\tilde{e} = (A, B)$, i.e., $S(\Pi, \tilde{e}) \subseteq (\text{Imp}_\Pi(A) \cup \text{Imp}_\Pi(B))$. Since the set $C_{A,\Pi} \cup C_{B,\Pi}$ contains no \tilde{e} -m.p., neither does the subschedule $S(\Pi, \tilde{e})$, and so the schedule Π is weakly \tilde{e} -proper. Hence, by definition of the collection $\hat{\Pi}(\tilde{e})$ of schedules (see Def. 5.13), if $\Pi \in (\Upsilon \setminus \Psi) \cap \Phi$ then necessarily $\Pi \in \hat{\Pi}(\tilde{e}) = \hat{\Pi}(\tilde{e}_\Pi)$. In other words, $(\Upsilon \setminus \Psi) \cap \Phi \subseteq \bigcup_{\tilde{e} \in \tilde{E}} \hat{\Pi}(\tilde{e}) = \hat{\Pi}$. By Corollary 5.15,

$$\mathbb{P}(\exists \Pi \in (\Upsilon \setminus \Psi) \cap \Phi \text{ s.t. } \mathcal{V}_2 \subseteq I(\Pi)) \leq \mathbb{P}(\exists \Pi \in \hat{\Pi} \text{ s.t. } \mathcal{V}_2 \subseteq I(\Pi)) \leq 2^{-\Omega(n^{9/5})}.$$

By definition of the set Φ , all the schedules in Φ are admissible. Hence

$$\mathbb{P}((\Upsilon \setminus \Psi) \cap \Phi \neq \emptyset) = \mathbb{P}(\exists \Pi \in (\Upsilon \setminus \Psi) \cap \Phi \text{ s.t. } \mathcal{V}_2 \subseteq I(\Pi)) \leq 2^{-\Omega(n^{9/5})},$$

and so $\mathbb{P}(\Phi \subseteq \Psi) \geq 1 - 2^{-\Omega(n^{9/5})}$, as required. \blacksquare

Lemma 6.3 *For every admissible schedule Π of length at most $\frac{\log n}{30}$ for the instance \mathcal{G} of the radio broadcast problem, the cover C_Π induced by the schedule Π has cardinality $O(\log^2 n) \cdot (|\tilde{V}_1| + |\tilde{V}_2|)$ (with probability 1).*

Proof: Consider some supervertex $X \in \tilde{V}_1 \cup \tilde{V}_2$. It is sufficient to prove that $|C_\Pi \cap X| = O(\log^2 n)$. Note that $C_{X,\Pi} = C_\Pi \cap X$. Recall that $|\Pi| = O(\log n)$. Consider some round $R_0 \in \Pi$. Note that it is sufficient to prove that $|\text{copy}(R_0) \cap C_\Pi| = O(\log n)$. If round R_0 is not important with respect to the supervertex X then $\text{copy}(R_0) \cap C_{X,\Pi} = \emptyset$. Otherwise, there exists a superedge $\tilde{e} = (X, Y) \in \tilde{E}$ that is adjacent to the supervertex X and such that the round R_0 is small with respect to \tilde{e} . Hence $|\text{copy}(R_0) \cap X| \leq |\text{copy}(R_0) \cap (X \cup Y)| = O(\log n)$, completing the proof. \blacksquare

Corollary 6.4 *Consider a NO-instance \mathcal{M} of the MIN-REP problem, and the instance \mathcal{G} of the radio broadcast problem that is obtained from \mathcal{M} via our reduction. With probability at least $1 - 2^{-\Omega(n^{9/5})}$, there exists no admissible schedule of length at most $\frac{\log n}{30}$ for \mathcal{G} .*

Proof: By Lemma 6.2, with probability at least $1 - 2^{-\Omega(n^{9/5})}$, $\Phi \subseteq \Psi$. However, by definition of Ψ and by Lemma 6.3, for any schedule $\Pi \in \Psi$, the induced cover C_Π is an admissible solution for the instance \mathcal{M} of the MIN-REP problem, and it has cardinality $O(\log^2 n) \cdot (|\tilde{V}_1| + |\tilde{V}_2|)$. Since \mathcal{M} is $\log^{10} n$ -NO-instance of the MIN-REP problem, the set Ψ is empty. Hence, with probability at least $1 - 2^{-\Omega(n^{9/5})}$, the set Φ is empty as well. \blacksquare

Theorem 6.5 *There exists a positive real universal constant $c > 0$ such that the radio broadcast problem is $c \cdot \log n$ -inapproximable unless $NP \subseteq BPTIME(n^{O(\log \log n)})$.*

Proof: Lemma 4.1 and Corollary 6.4 show that our reduction transforms YES-instances of the MIN-REP problem into instances of the radio broadcast problem that admit schedules of length 3, and $\log^{10} n$ -NO-instances of the MIN-REP problem into instances of the radio broadcast problem for which with overwhelming probability any admissible schedule is of length $\Omega(\log n)$. It is easy to see that the reduction can be implemented in (probabilistic) polynomial time, and that the size of the obtained instance of the radio broadcast problem is at most polynomial in the size of the

original instance of the MIN-REP problem. Hence, the statement of the theorem follows from Theorem 2.1. ■

Finally, note that our proof can be easily modified to show that for any positive integer $r = 2, 3, \dots$, there is a set of instances $\mathcal{G} = (s, \hat{G})$ with $rad(\hat{G}, s) = r$ that either have an admissible schedule of length $r + 1$ or satisfy that any admissible schedule for them is of length $r + \Omega(\log n)$, and such that it is hard to distinguish between instances of these two types. This generalization is achieved by adding a path of length r , $(s, s_1, s_2, \dots, s_r)$ to the construction, and connecting s_r to all the vertices of \mathcal{V}_1 (instead of connecting s to all the vertices of \mathcal{V}_1).

Acknowledgments

The authors are grateful to an anonymous referee for suggesting a simpler proof of Lemma 5.11, and providing a tighter bound on the probability in Lemma 5.14. The first-named author also wishes to thank Oded Regev for helpful discussions.

References

- [ABLP91] N. Alon, A. Bar-Noy, N. Linial, D. Peleg. A lower bound for radio broadcast. In *Journal of Computer and System Sciences* 43, pp. 290-298, 1991.
- [ABLP92] N. Alon, A. Bar-Noy, N. Linial, D. Peleg. Single round simulation on radio networks, *Journal of Algorithms* 13 (1992), pp. 188-210.
- [AL96] S. Arora and C. Lund. Hardness of Approximations. In *Approximation Algorithms for NP-hard Problems*, Dorit Hochbaum, Ed. PWS Publishing, 1996.
- [B01] R. Bar-Yehuda. Private communication.
- [BK97] A. Blum and D. Karger. An $O(n^{3/14})$ -Coloring Algorithm for 3-Colorable Graphs. *Information Processing Letters*, 61(1):49–53, 1997.
- [BGI91] R. Bar-Yehuda, R., O. Goldreich, and A. Itai. Efficient emulation of single-hop radio network with collision detection on multi-hop radio network with no collision detection. In *Distributed Computing*, 5(2):67–72, 1991.
- [CK85] I. Chlamtac and S. Kutten. On broadcasting in radio networks - Problem analysis and protocol design, *IEEE Trans. Commun.* **33** (1985), pp. 1240-1246.
- [CK87] I. Chlamtac and S. Kutten. Tree based broadcasting in multi hop radio networks, *IEEE Trans. Commun.* **36** (1987), pp. 1209-1223.
- [CK2-85] I. Chlamtac and S. Kutten. A spatial reuse TDMA/FDMA for mobile multi-hop radio networks, in *Proc. IEEE INFOCOM*, 1985, Washington, DC, pp. 389-394.
- [CW87] I. Chlamtac and O. Weinstein. The wave expansion approach to broadcasting in multi hop radio networks, in *Proc. INFOCOM* 1987, pp. 874-881.

- [DK99] Y. Dodis and S. Khanna, Designing Networks with Bounded Pairwise Distance, *Proc. 30th ACM Ann. Symp. of Theory of Computing*, 1999.
- [GM95] I. Gaber and Y. Mansour. Broadcast in Radio Networks. In *Proc. 6th ACM-SIAM on Discrete Algorithms*, 577-585, 1995.
- [GPM03] R. Gandhi, S. Parthasarathy and A. Mishra. Minimizing Broadcast Latency and Redundancy in Ad Hoc Networks. To appear in the Proc. of the Fourth ACM International *Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC'03)*, Jun. 2003.
- [GVF87] I. Gitman, R. M. Van Slyke, and H. Frank. Routing in packet-switching broadcast radio networks, *IEEE Trans. Commun.* (1976), pp. 926-930.
- [KGBK78] R. E. Kahn, S. A. Gronemeyer, N. Madras, and P. March. Advances in the radio technology, *Proc. IEEE* **66** (1978).
- [K98] G. Kortsarz. On the hardness of approximating spanners. *The first International Workshop APPROX-98*, pages 135–146, 1998.
- [KP02] D. Kowalski and A. Pelc. Deterministic broadcasting time in radio networks of unknown topology, *Proc. 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2002)*, November 2002, Vancouver, Canada, 63-72.
- [LY94] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *J. Assoc. Comput. Mach.*, 41(5):960–981, 1994.
- [R98] R. Raz. A Parallel Repetition Theorem, *SIAM Journal of Computing* 27(3) (1998) pp. 763-803.
- [SC82] N. Shacham and E. J. Craighill, Dynamic routing for real-time transport in packet radio networks, in *Proc. INFOCOM*, 1982.