Polylogarithmic Additive Inapproximability of the Radio Broadcast Problem

Michael Elkin [*†‡] Guy Kortsarz[§]

April 27, 2005

Abstract

The input for the radio broadcast problem is an undirected n-vertex graph G and a source node s. The goal is to send a message from s to the rest of the vertices in minimum number of rounds. In a round, a vertex receives the message only if exactly one of its neighbors transmits. The radio broadcast problem admits an $O(\log^2 n)$ approximation [10, 22].

In this paper we consider the additive approximation ratio of the problem. We prove that there exists a constant c so that the problem can not be approximated within an additive term of $c \log^2 n$, unless $NP \subseteq BTIME(n^{O(\log \log n)})$.

[*]Department of Computer Science, Ben-Gurion University, Beer-Sheva, Israel. elkinm@cs.bgu.ac.il

[†]This work was done in Department of Computer Science, Yale University, New Haven, CT, USA.

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

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

1 Introduction

1.1 The Radio Broadcast Problem

1.1.1 Definition and Motivation

Consider a synchronous network of processors that communicate by transmitting messages to their neighbors, where a processor receives a message in a given step if and only if *precisely one* of its neighbors transmit. The instance of the Radio Broadcast problem, called *radio network*, is a pair $(G = (V, E), s)$, $s \in V$, where G is an unweighted undirected n-vertex graph, and s is a vertex, called *source*. The objective is to deliver one single message that the source s generates to all the vertices of the graph G using the smallest possible number of communication rounds. The prescription that tells each vertex when it should broadcast is called *schedule*; the *length* of the schedule is the number of rounds it uses, and it is called *feasible* if it informs all the vertices of the graph. From practical perspective, the interest in radio networks is usually motivated by their military significance, as well as by the growing importance of cellular and wireless communication (see, e.g., [19, 15, 4]). The Radio broadcast is perhaps the most important communication primitive in radio networks, and it is intensively studied starting from mid-eighties [9, 20, 21, 6, 5, 8, 18, 15, 19, 1, 4, 10, 7].

From theoretical perspective, the study of the Radio Broadcast problem provided researchers with a particularly convenient playground for the study of such broad and fundamental complexity-theoretic issues as the power and limitations of randomization, and of different models of distributed computation [4, 19, 21]. In this paper we study the approximation threshold of the *Radio Broadcast* problem. We believe that our results show that this problem is of a particular interest from the stand-point of the theory of Hardness of Approximation as well.

1.1.2 Previous Results

Upper bounds: The first algorithm for the Radio Broadcast problem was devised by Chlamtac and Weinstein in 1987 [10]. That algorithm, given an instance (G, s) of the problem, constructs a feasible broadcast schedule of length $O(Rad(G, s) \cdot \log^2 n)$ where $Rad(G, s)$ stands for the *radius* of the instance (G, s), that is, the maximum distance $d_G(s, v)$ in the graph G between the source s and some vertex $v \in V$. Their algorithm is *centralized*, i.e., it accepts the entire graph as input.

Soon afterwards Bar-Yehuda et al. [4] devised a distributed randomized algorithm that

provides feasible schedules of length $O(Rad(G,s)\cdot\log n+\log^2 n)$. Recently [22] a *deterministic* (albeit, centralized) algorithm with the same performance was given by Kowalski and Pelc. Alon et al. [1] have shown that the additive term of $\log^2 n$ in the result of [4, 22] is inevitable, and devised a construction of infinitely many instances (G,s) of constant radius that satisfy that any broadcast schedule for them requires $\Omega(\log^2 n)$ rounds. Kushilevitz and Mansour [19] have shown that for *distributed* algorithms, the multiplicative logarithmic term in the result of [4] is inevitable as well, and proved that for *any distributed algorithm* for the Radio Broadcast problem there exist (infinitely many) instances (G,s) on which the algorithm constructs a schedule of length $\Omega(Rad(G,s)\cdot\log(n/Rad(G,s)))$. Finally, the gap between the $\log n$ and $\log(n/Rad(G,s))$ was recently closed by Kowalski and Pelc [21], and Czumaj and Rytter [9].

Gaber and Mansour [15] devised a centralized algorithm that constructs feasible schedules of length $O(Rad(G,s)+\log^5 n)$. In [12] we improved this result providing a schedule of length $Rad(G,s)+O(\sqrt{Rad(G,s)}\cdot\log^2 n)=O(Rad(G,s)+\log^4 n)$.

Since, obviously, any schedule for an instance (G,s) requires at least $Rad(G,s)$ rounds, the algorithms for the Radio Broadcast problem [10, 4, 15, 21, 9, 22] can be interpreted as *approximation algorithms* for the problem. In particular, [10, 22] is a deterministic $O(\log^2 n)$ approximation algorithm.

Lower bounds: The NP-hardness of the Radio Broadcast problem was shown by Chlamtac and Kutten [7] already in 1985. In [16] an NP-hardness result is established for solving the problem on unit disc graphs.

A gap reduction is a reduction that maps an arbitrary NPC problem to the problem at hand giving some gaps for the optimum values resulting from a yes and a no instance. The authors of the current paper have shown [12] a gap reduction that maps a yes instance to a radio broadcast instance that admits a 3 rounds schedule, while a no instance is mapped into an $\Omega(\log n)$ schedule. This proves that there exists a constant $c>0$ such that the Radio Broadcast problem cannot be approximated within approximation ratio of $c\log n$ unless $NP\subseteq BPTIME(n^{O(\log\log n)})$.

1.2 Our Results

Note that [12, 15] can be considered as an additive approximation algorithms for the problem. Hence, we study the additive ratio of radio broadcast. We provide a gap reduction that maps a yes instance to a radio network that admits a schedule of length $O(\log n)$, and a no instance

to a radio network for which any feasible schedule is of length $\Omega(\log^2 n)$. Thus, there exists some $c > 0$ so that the radio broadcast problem admits no polynomial additive $c \log^2 n$ ratio approximation unless $NP \subseteq BPTIME(n^{O(\log \log n)})$. This fully determines the additive approximation ratio of the problem for graphs with radius at most $\log n$ as the result of [22] implies that for graphs with radius at most $\log n$, there exists a matching additive upper bound of $O(\log^2 n)$. We are not aware of any other problem that exhibits a tight additive polylogarithmic ratio. (See [17, 14] for the only example we are aware of an almost tight polylogarithmic *multiplicative* approximation. This example is the Group Steiner problem on trees.)

Remark: A big challenge seems to be designing a gap reduction that maps a yes instance to a constant number of rounds schedule and a no instance to a schedule of length $\Omega(\log^2 n)$. We leave this question open. If such a proof is possible, then the (multiplicative) best approximation ratio for the problem is $\log^2 n$ (up to constants) much like the Group Steiner problem on trees. Alternatively, the challenge is to design an $O(\log n)$ ratio approximation for small radios graphs.

2 Preliminaries

2.1 Definitions and notation

We start with introducing some definitions and notations. In all the notations, we may eventually omit some parameters, if the meaning can be deduced from the context.

Definition 2.1 *The set of neighbors of a vertex v in an unweighted undirected graph $G(V, E)$, denoted $\Gamma_G(v)$, is the set $\{u \in V \mid (v, u) \in E\}$. For a subset $X \subseteq V$, the set of neighbors of the vertex v in the subset X, denoted $\Gamma_G(v, X)$ is the set $\{u \in X \mid (v, u) \in E\}$.*

Notation 2.2 *For a positive integer number n, let $[n]$ denote the set $\{1, 2, \ldots, n\}$.*

Definition 2.3 *Let $G = (V, E)$ be an unweighted undirected graph, and $R \subseteq V$ be a subset of vertices. The set of vertices* informed *by R, denoted $I(R)$, is $I(R) = \{v \mid \exists! x \in R \ \ s.t. \ \ v \in \Gamma_G(x)\}$ (the notation $\exists! x$ stands for "there exists a unique x"). For a singleton set $R = \{x\}$, $I(R) = I(\{x\}) = I(x) = \Gamma_G(x)$.*

A sequence of vertex sets $\Pi = (R_1, R_2, \ldots, R_q)$, $q = 1, 2, \ldots$, is called a *radio broadcast schedule* (henceforth referred as a *schedule*) if $R_{i+1} \subseteq \bigcup_{j=1}^{i} I(R_j)$ for every $i = 1, 2, \ldots, q-1$.

Intuitively, this condition means that the vertices that send a message in certain round have to be informed in one of the previous rounds.

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

Given a graph $G = (V, E)$ and a vertex $s \in V$, a schedule Π is *feasible* with respect to (G, s) if $R_1 = \{s\}$ and $V \subseteq I(\Pi)$.

The *length* of the schedule $\Pi = (R_1, R_2, \ldots, R_q)$ is $|\Pi| = q$.

An instance of the *radio broadcast problem* $\mathcal{G}$ is a pair $(\bar{G} = (\bar{V}, \bar{E}), s)$, where $\bar{G}$ is a graph, and $s \in \bar{V}$ is a vertex. The goal is to compute a feasible schedule Π of minimal length. The *value* of an instance $\mathcal{G}$ of the radio broadcast problem is the length of the shortest feasible schedule Π for this instance.

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

2.2 The MIN-REP problem

Definition 2.4 *The MIN-REP problem is defined as follows. The input consists of a bipartite graph $G = (V_1, V_2, E)$. In addition, for $j = 1, 2$, the input contains a partition $\tilde{V}_j$ of V_j into a disjoint union of subsets, $V_1 = \bigcup_{A\in\tilde{V}_1} A$, $V_2 = \bigcup_{B\in\tilde{V}_2} B$. The triple $\mathcal{M} = (G, \tilde{V}_1, \tilde{V}_2)$ is an* instance *of the MIN-REP problem. The size of the instance is $n = |V_1| + |V_2|$. An instance G as above induces a bipartite super-graph $\tilde{G} = (\tilde{V}_1, \tilde{V}_2, \tilde{E})$ in which the sets A and B of the partition serve as the vertices of the super-graph. The edges of the super-graph are $\tilde{E}(\mathcal{M}) = \tilde{E} = \{(A, B) \in \tilde{V}_1 \times \tilde{V}_2 \mid a \in A, b \in B, (a, b) \in E\}$. In other words, there is a (super-)edge between a pair of sets $A \in \tilde{V}_1$, $B \in \tilde{V}_2$ if and only if the graph G contains an edge between a pair of vertices $a \in A$, $b \in B$.*

Denote $\tilde{V} = \tilde{V}_1 \cup \tilde{V}_2$. A pair of vertices $x_1, x_2 \in V_1 \cup V_2$ is called a matched pair with respect to a super-edge $\tilde{e} = (A, B) \in \tilde{E}$ *(henceforth, $\tilde{e}$-*m.p.*) if $(x_1, x_2) \in E$ and either $x_1 \in A$ and $x_2 \in B$ or vice versa.*

A subset $C \subseteq V_1 \cup V_2$ of vertices is said to cover *a super-edge $\tilde{e} = (A, B)$ if it contains an $\tilde{e}$-m.p. 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. *In other words, a MAX-cover C contains exactly one vertex from each super-vertex.*

An instance $\mathcal{M}$ of the MIN-REP problem is called a yes instance *if there exists a MAX-cover that covers all the super-edges. Such a MAX-cover is called a* perfect MAX-cover.

For a positive real number $t > 1$, an instance $\mathcal{M}$ of the MIN-REP problem is called a

t-no instance *if any C that covers at least half of the superedges must pick on average at least t elements of every A and every B (every C that covers at least half of the superedges has size at least t times the number of A, B sets).*

The maximization version of MIN-REP problem is equivalent to the maximization variant of the Label-Cover problem (see, e.g., [3]).

The parameters of the MIN-REP instance: We impose several additional (somewhat less standard) restrictions on the set of instances of the MIN-REP problem. For the rest of the paper, let n denote the number of vertices in the MIN-REP instance.

1. All the super-vertices $X \in \tilde{V}$ are of size polylogarithmic in n (namely, size at most $(\log n)^d$ for some constant d).
2. The number of super-edges is $O(n \cdot \ polylog(n))$
3. **The Star property:** For every super-edge $\tilde{e} = (A, B) \in \tilde{E}$, $A \subseteq V_1$ and $B \subseteq V_2$ and every vertex $b \in B$ there exists exactly one vertex $a \in A$, denoted $\tilde{e}(b)$, such that $(a, b) \in E$. The set of all vertices b such that $\tilde{e}(b) = a$ for the same vertex a, along with the vertex a, is called an *$\tilde{e}$-star*.

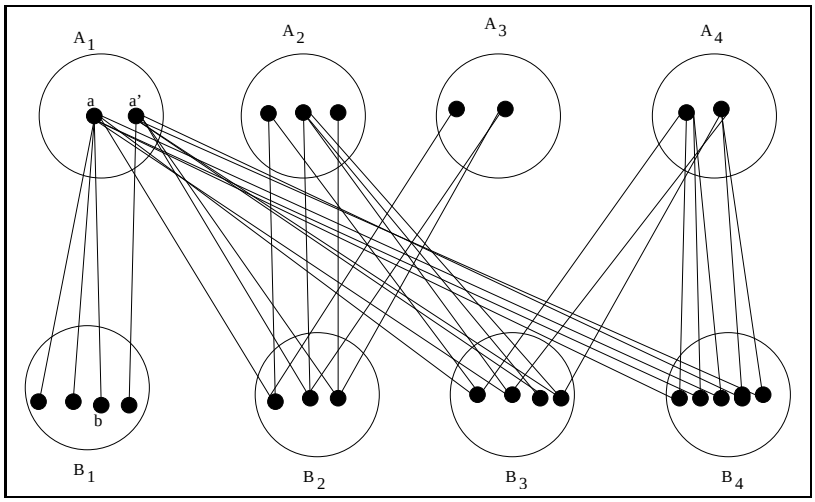

Figure 1: An example of a MIN-REP instance that satisfies the star property. Every pair of sets (A_i, B_j) in the partition, induces a collection of disjoint stars with heads in A. The vertices (a, b) form a matching pair that covers the super-edge (A_1, B_1). The vertices a', b do not form a matching pair, and do not cover the super-edge (A_1, B_1).

Essentially, the star property means that for every super-edge $\tilde{e} = (A, B) \in \tilde{E}$, the graph induced by the subset $A \cup B$ decomposes into a collection of vertex-disjoint stars (a *star* is a graph with all vertices but (maybe) one having degree 1). The vertex with degree larger than 1 is called the *head* of the star; if there are only two vertices in the star, then the head is the vertex that belongs to the super-vertex A. The other vertices of the star are called the *leaves* of the star. See Figure 1 for an example of a MIN-REP instance that obeys the star property.

Theorem 2.5 *[2, 24] No deterministic polynomial time algorithm can distinguish between the YES-instances and the* $\log^{10} n$*-no instances of the MIN-REP problem, unless* $NP \subseteq DTIME(n^{O(\log\log n)})$*, even when the instances of the MIN-REP problem satisfy the conditions (1)-(3).*

3 The construction

3.1 The high-level idea

In [23] a reduction from an arbitrary NPC language to the set-cover problem is given. The elements of the set-cover instance are grouped into a union of *ground sets*. In a yes instance, every ground set M can be covered by two "complementary" sets each covering a disjoint half of M. In a no instance, every set in the set-cover instance that contains elements of M (essentially) contains a random half of M and so $\Omega(\log |M|)$ sets are required to cover the entire set M. This is used in [11] to design a radio network that admits no schedule of length $o(\log n)$ for a no instance, and a schedule with only a constant number of rounds for a yes instance.

In [1], another special kind of set-cover is designed. The elements are partitioned to $\Theta(\log n)$ ground sets M_j. In this instance, covering *uniquely* many elements in $\bigcup M_j$ is not possible (an element is uniquely covered by a collection of sets if it belongs to exactly one set in the collection). Specifically, if a collection of sets uniquely covers "many" M_j elements, then it does not uniquely cover many elements of M_q, for any $q \neq j$. This construction is used in [1] to design a radio network that, essentially, has to inform the sets M_j "one by one" while the construction for every M_j is similar to the one from in [23, 11], namely, informing M_j by itself requires $\Omega(\log n)$ rounds. Since the number of sets M_j is $\Theta(\log n)$, a lower bound of $\Omega(\log^2 n)$ for the length of a feasible schedule follows.

We modify the construction of [1], and add a "trap door" to their construction, using ideas

of [23]. This trapdoor makes it possible to inform every M_j in 2 rounds. This guarantees that for a yes instance, a feasible schedule of logarithmic length exists: simply inform M_j one j after the other. On the other hand, the modification maintains the lower bound of $\Omega(\log^2 n)$ for a no instance. Hence, we obtain a gap between $O(\log n)$ and $\Omega(\log^2 n)$, that is, an additive gap of $\Omega(\log^2 n)$.

3.2 The construction of [1]

Since our construction relies on that of [1] we briefly sketch their construction.

Definition 3.1 *A schedule of at most* $\log^2 n/100$ *rounds is called a short schedule.*

Let $(\mathcal{X}, \mathcal{Y})$ be two sets of vertices. Let $|\mathcal{X}| = n$ and $\mathcal{Y}$ be a disjoint union $\mathcal{Y} = \bigcup \mathcal{Y}_j$ of sets $\mathcal{Y}_j$ each containing n^7 vertices for $0.4 \cdot \log_2 n \leq j \leq 0.6 \cdot \log_2 n$. Thus, $|\mathcal{Y}| = \Theta(n^7 \log n)$.

A vertex $x \in \mathcal{X}$ and a vertex $y \in \mathcal{Y}_j$ are connected with an edge with probability $1/2^j$ independently of other edges. In addition, add a source s and connect s to all the vertices of $\mathcal{X}$. Observe that, by definition, after the first round the set of informed vertices is exactly $\{s\} \cup \mathcal{X}$.

Intuitively, in the above construction any transmitting subset of $S \subseteq \mathcal{X}$ helps to inform only part of the sets $\mathcal{V}_j$. It is not possible to choose a size for S so that all $\mathcal{Y}_j$ will contain many vertices informed by S. For a given j, for any set S of size larger than 2^j, there may exist many vertices of $\mathcal{Y}_j$ having at least two neighbors in S. But if S is much smaller than 2^j, then many of the vertices of $\mathcal{Y}_j$ will not have even a single neighbor in S.

The following elegant lemma formalizes this intuition.

Lemma 3.2 [1] *Let* $\Pi = (R_1, R_2, \ldots, R_t)$ *be a short (namely,* $t \leq \log^2 n/100$*) collection of subsets of* $\mathcal{X}$*. Then there exists a subset* $S \subseteq \mathcal{X}$ *and an index* j*,* $0.4 \cdot \log_2 n \leq j \leq 0.6 \cdot \log_2 n$ *so that:*

1. $|S| \leq 2^j \cdot \log_2 n$.
2. *Let* $\Phi' = (R'_1, R'_2, \ldots) = \Phi \setminus S$*. Then for each round* R'_q *in the schedule,* $|R'_q| \geq 2^j$.
3. *Let* f_k *be the number of sets in the schedule with cardinality* $2^{j+k} \leq |R'_j| \leq 2^{j+k+1}$*. Then,*

$$\sum_{k \geq 0} \frac{f_k}{2^k} \leq \log n.$$

Indeed, how could a "short" schedule Φ cover $\mathcal{Y}_j$? The set S has size 2^j, so there is a non-negligible probability that no vertex in S is connected to a vertex in $\mathcal{Y}_j$. (We use the term *non-negligible* for a the probability which is at least $\frac{1}{poly(n)}$, where $poly(n)$ is some polynomial in n.) Thus, the task of covering $\mathcal{Y}_j$ may be left to $\Phi' = (R'_1, R'_2, \ldots) = \Phi \setminus S$. Consider some vertex $y_j \in \mathcal{Y}_j$. Observe that by Item 2 above, each R'_q is of large enough size to make the probability of R'_q not informing y_j non-negligible. Indeed, it is reasonable to expect that at least two vertices of $\mathcal{R}'_q$ will be connected to y_j in which case y_j does not get the message in round q. Now, since $\mathcal{Y}_j$ is "large" (has size n^7), with high probability there will be a vertex that is not going to be informed at any round. The paper of Alon et al. [1] formally proves this claim along these lines.

3.3 Intuition behind the random permutation step

One of the difficulties in imitating the construction of [23], and combining it with the construction of [1] is as follows. The construction of [1] requires that vertices are connected to $\mathcal{Y}_j$ with probability $1/2^j$. On the other hand, in the construction of [23] some vertices are connected to one half of the elements in every ground set M_j (see [23] for more details). Thus, the probability that a and $v \in M_j$ are connected is $1/2$.

The way we overcome this difficulty is by forming *many* copies of every vertex $x \in A \cup B$.

Consider some super-edge $\tilde{e}$ and the ground sets that correspond to $\tilde{e}$. Suppose that $M = M_{\tilde{e}}(j)$ is some ground set that corresponds to $\tilde{e}, j$ (similar to $\mathcal{Y}_j$ but dedicated to $\tilde{e}$). Every copy of a is connected in M to some random subset of size $|M|/2^{j+1}$. We ensure that neighbors in M of different $a-$copies are disjoint, and that 2^j copies of a take part in this process. This implies that altogether the copies of $a \in A$ are indeed connected to a half of M. Let M_a denote this half.

In addition, let (a, b) be an $\tilde{e}-$matching pair. Then the copies of b are similarly connected but to $M_j \setminus M_a$. Thus, the copies of b cover the complementary half of M.

This way we are able on the one hand to control the degrees of copies of a and b (that is, to make it roughly $|M|/2^j$, as required in [1]), but on the other hand to guarantee that the copies of a and b cover together disjoint halves of M (as required in [23]).

For the claims in [1] to work we need the neighbors of (copies of) a and b in M to be random. We use copies of a and b for covering random elements of M as follows. We first choose a random half of M. Then this random half is arbitrarily split into 2^j equal parts. Then match the 2^j copies of a and the 2^j parts by a random permutation. The copy of a is

connected to all vertices in its matching part. Similar construction is applied for copies of b on the complementary half.

3.4 The random permutation step: formal definition

For the rest of the paper, let n denote the number of vertices of the MIN-REP graph. Consider an instance $\mathcal{M} = (G, \tilde{V}_1, \tilde{V}_2)$, $G = (V_1, V_2, E)$ of the MIN-REP problem with $V_1 = \bigcup_{A \in \tilde{V}_1} A$, $V_2 = \bigcup_{B \in \tilde{V}_2} B$. The reduction constructs an instance $\mathcal{G} = \mathcal{G}(\mathcal{M}) = (\bar{G}, s)$, $\bar{G} = (\bar{V}, \bar{E})$, $s \in \bar{V}$, of the radio broadcast problem in the following way.

Let $N = n^{0.6}$. The vertex set $\bar{V}$ of the graph consists of the source s, and the disjoint vertex sets $\mathcal{V}_1$ and $\mathcal{V}_2$.

The vertex set $\mathcal{V}_1$ contains N copies of every vertex a or b in $V = V_1 \cup V_2$; the set of all copies of a vertex x ($x = a$ or $x = b$) is denoted by $cp(x)$, and $cp(x, j)$ is the subset that contains the first 2^j copies of x.

For a subset $X \subseteq V$, let $cp(X)$ denote $cp(X) = \bigcup_{x \in X} cp(x)$. Let $\hat{J}$ denote the set of indices $\{0.4 \log n, 0.4 \log n + 1, \ldots, 0.6 \log n\}$.

The vertex set $\mathcal{V}_2$ is of the form $\mathcal{V}_2 = \bigcup_{\tilde{e} \in \tilde{E}} M_{\tilde{e}}$, where the *ground sets* $M_{\tilde{e}}$ are disjoint, and all have equal size. Each ground set $M_{\tilde{e}}$ is a disjoint union of the sets $M_{\tilde{e}}(j, q)$, $j \in \hat{J}$, $q \in [L]$, with $L = n^{c_0+4}$, and c_0 is an integer positive universal constant that will be determined later. The sets $M_{\tilde{e}}(j, q)$ are all of equal size $M = n^{c_0}$, for the same constant c_0.

The edge set $\bar{E}$ of the graph $\bar{G}$ contains edges that connect the source s to the vertices of the set $\mathcal{V}_1$. We next construct the edge set between the vertices of $\mathcal{V}_1$ and $\mathcal{V}_2$. Fix $\tilde{e} = (A, B) \in \tilde{E}$, and the indices indices $j \in \hat{J}$, $q \in [L]$.

The random permutation step:

1. For every star head $a \in A$ let $M_a = M_{\tilde{e},a}(j, q)$ be an *exact random half* of the set $M = M_{\tilde{e}}(j, q)$.

2. For every vertex b in the star of a set $M_b = M_{\tilde{e},b}(j, q) = M \setminus M_a$.

   Remark: Steps 1-2 will be referred as the *exact partition step*.

3. For every vertex $a \in A$, partition the set $M_a = M_{\tilde{e},a}(j, q)$ arbitrarily into 2^j disjoint subsets of equal size. Randomly permute $cp(a, j)$ (the first 2^j copies of a) and connect the ith copy (in the order determined by the random permutation) of a to the ith part of M_a.

4. Similarly, for every vertex b that belongs to the star of a, cover M_b by a random permutation. The random permutations of a and of leaves in the star of a are independent.

See Figure 2 for an illustration of the random permutation step.

Remarks:

1. The parameter q in $M = M_{\tilde{e}}(j,q)$ does not affect the probability of $\mathcal{V}_1$ vertices to be adjacent to $M_{\tilde{e}}(j,q)$. This probability is $1/2^j$. Unlike [1], many "$\mathcal{Y}_j$ type" sets are required. It is important though that if $q \neq q'$ then different events for $M_{\tilde{e}}(j,q)$ and $M_{\tilde{e}}(j,q')$ are independent.

2. If b and b' belong to the star of a in $\tilde{e}$, then $M_b = M_{b'} = M \setminus M_a$. However, the random permutations of b and b' are independent, and are likely to be different.

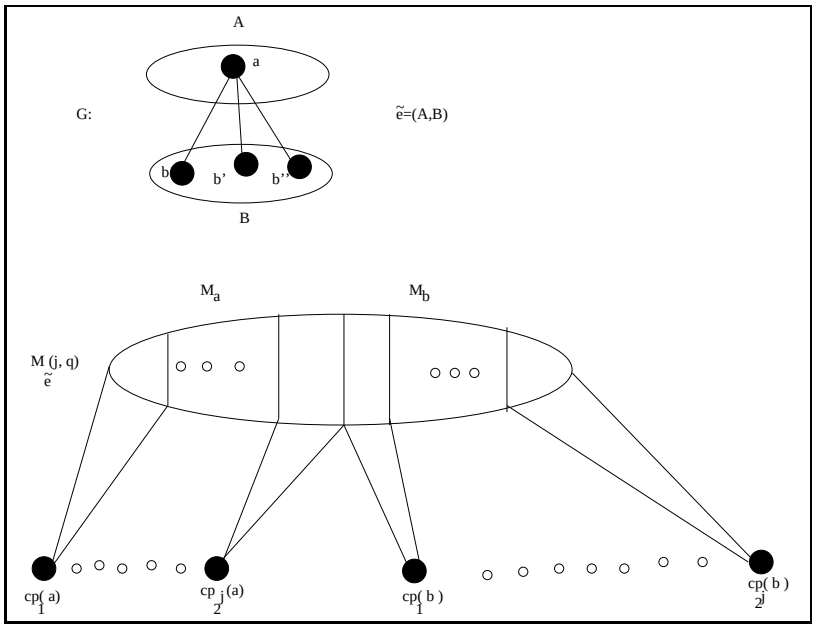

Figure 2: The figure illustrates the random permutation step. The copies of the vertex a cover an exact half of the vertices of the set $M_{\tilde{e}}$. The copies of the vertex b cover the complementary half of the vertices.

Adding dummy vertices: Recall that $cp(X)$ is the set of all copies of X vertices. Currently, $|cp(A)|, |cp(B)| = \tilde{O}(n^{0.6})$. We need to later use Lemma 3.2 with $cp(A) \cup cp(B)$ playing the role of $\mathcal{X}$. For that, we need that $|cp(A) \cup cp(B)| = n$ (otherwise, it is required to use the lemma with $\tilde{\Theta}(n^{0.6})$ playing the role of n which may be confusing). Add dummy

vertices to every $cp(A)$ and $cp(B)$ to complete its size to $n/2$ each. The dummy vertices have no connection to $\mathcal{V}_2$, but are joined to s. Thus, dummy vertices never transmit (do not belong to any round). This change only affects the constants in the ratio. Thus, we throughout assume that $|cp(A)| = |cp(B)| = n/2$ for every A, B.

3.5 The mixing step: intuition and formal definition

Intuitively, we need to build a reduction in which a short schedule for the resulting radio network necessarily reveals a good solution for the original instance of the MIN-REP problem. Namely, a round is forced to use copies of many matched pairs, otherwise, the connections are random in a way similar to [1].

By the construction so far, this goal is not yet achieved because vertices can "coordinate efforts" even if they do not belong to a matching pair. For example, observe that if b, b' both belong to the star of a then the copies of b and b' are connected in $M_{\tilde{e}}(j,q)$ to the same half (see Figure 2). This half is $M_b = M_{b'} = M_{\tilde{e}}(j,q) \setminus M_{\tilde{e},a}(j,q)$. Even though the random permutations of b and b' are independent, inserting both copies of b and b' into R increases the probability that no vertex in $M_b = M_{b'}$ remains uncovered. Therefore, so far we have not prevented b and b' from coordinating efforts.

Thus, we should modify the construction so that copies of b and copies of b' "hurt each other", and consequently, they cannot be used in the same round together to cover many vertices. This is done by adding some random edges. Copies of a are randomly connected to the other half of the vertices, namely to the vertices of the set $M_b = M_{b'} = M_{\tilde{e}}(j,q) \setminus M_{\tilde{e},a}(j,q)$. Copies of b are randomly connected to $M_a = M_{\tilde{e},a}(j,q)$. See Figure 3.

These additional edges prevent a schedule from forming very large rounds. Because if a round is very large, many elements are covered two times or more. For example, inserting many copies of b and also many copies of b' into a round leads to "over-covering" vertices.

Further, the copies of $cp(a) \setminus cp(a,j)$ pose a problem. So far, they have no edges to $M_{\tilde{e}}(j,q)$. This would imply that we may add vertices from $cp(a) \setminus cp(a,j)$ without affecting $M_{\tilde{e}}(j,q)$, and consequently, it leaves a possibility of forming large big rounds with only a small number of vertices that are connected to $M_{\tilde{e}}(j,q)$.

Hence, we need to connect every $cp(a) \setminus cp(a,j)$ vertex to every $M_{\tilde{e}}(j,q)$ vertex with probability $1/2^j$.

3.6 The mixing step: formal definition

Let $j, q, \tilde{e}$, $\tilde{e} = (A, B)$ be fixed. Fix some $\tilde{e}$-star with head a. Let $M = M_{\tilde{e}}(j, q)$. Let M_a be the set of neighbors of (the copies of) a in M and $M_b = M \setminus M_a$.

1. For every copy of a in $cp(a, j)$ and every vertex $v \in M_{\tilde{e},b}(j, q)$, add an edge between those two vertices with probability $1/2^j$.
2. Similarly, for every copy of b in $cp(b, j)$, and and every vertex $v \in M_a = M_{\tilde{e},a}(j, q)$, add an edge with probability $1/2^j$.
3. For every j and every $x \in A \cup B$ connect every vertex of $cp(x) \setminus cp(x, j)$ to every vertex of $M_{\tilde{e},a}(j, q)$ independently, with probability $1/2^j$.

Steps 1, 2, and 3 will be referred to as the *mixing step*. See Figure 3.

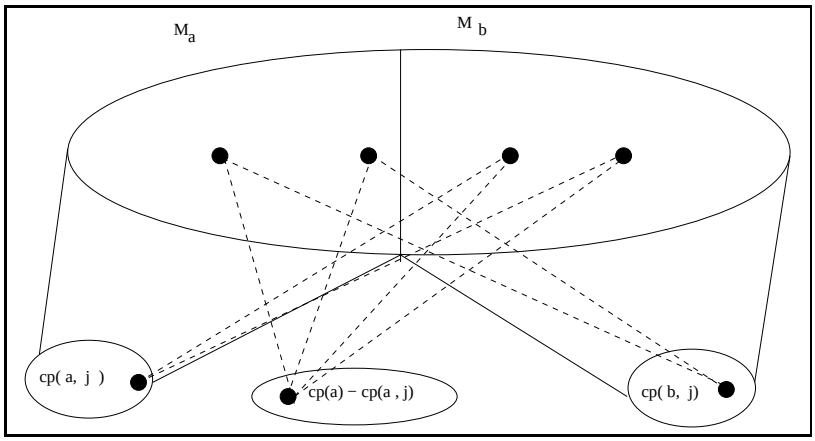

Figure 3: The figure illustrates the mixing step for some fixed q, j. The dotted edges represent random events that may result in edges. The probability for such an edge to be present is $1/2^j$. The figure also indicates that, the vertices of $cp(a, j)$ form an exact cover of $M_a = M_{a,\tilde{e}}(j, q)$, by the random permutation step. On the other hand, pairs of vertices from $cp(a) \setminus cp(a, j)$ and M are connected with probability $1/2^j$ for every pair.

3.7 Trapdoor: a schedule of logarithmic length for a yes instance

One way to explain some of the ideas behind the construction is by showing that the radio network resulting from a yes instance of the MIN-REP problem admits a schedule of length $O(\log n)$.

Let $\mathcal{M} = (G = (V, E), \tilde{G})$, $|V| = n$, be a yes instance of the MIN-REP problem, and let $(\mathcal{G}, s)$ be the instance of of the radio broadcast problem that is obtained via our reduction.

Let C be a perfect MAX-cover, that is, a subset of the set $V_1 \cup V_2$ that covers all the super-edges and contains exactly one vertex from each super-vertex. (Recall that by definition of the yes instance of the MIN-REP problem, there exists a perfect MAX-cover C for such an instance.)

Lemma 3.3 *There is a schedule of length $O(\log n)$ for the radio network $(\mathcal{G}, s)$.*

Proof: On the first round s transmits, and all the vertices of $\mathcal{V}_1$ are informed. Then, for each index j, $0.4 \log n \leq j \leq 0.6 \log n$, build two rounds. On the first one all the vertices of $\bigcup_{a \in C} cp(a, j)$ transmit in parallel, and on the second one all the vertices of $\bigcup_{b \in C} cp(b, j)$ transmit in parallel. Altogether, we obtain a schedule with $2 \cdot (02. \log n + 1) = O(\log n)$ rounds.

Claim 3.4 *The schedule informs $\mathcal{V}_2$.*

Proof: Since the set C is a perfect MAX-cover, for every super-edge $\tilde{e} = (A, B) \in \tilde{E}$, there exists some $\tilde{e}$-m.p. $a, b \in C$. Thus, when $cp(a, j)$ broadcasts, all the sets $M_{a,\tilde{e}}(j, q)$ sets are informed (observe that the mixing step does not insert edges between $cp(a, j)$ and $M_{\tilde{e},a}(j, q)$). Also, when $cp(b, j)$ transmits, all of the sets $M_{b,\tilde{e}}(j, q)$ are informed. ∎

∎

4 Analysis, part I: comparison to Lemma 3.2

For the rest of the paper, let $a, b, b', \tilde{e}, j, q, M_{\tilde{e}}(j, q)$ and $v \in M_{\tilde{e}}(j, q)$ be vertices, indices and sets that satisfy that (a, b) and (a, b') are $\tilde{e}$-matching pairs.

Since $j, q, \tilde{e}$ are fixed, for the simplicity of the notation we use M for $M_{\tilde{e}}(j, q)$ and M_a for $M_{\tilde{e},a}(j, q)$, etc. See also Figure 2.

In the next subsection we discuss a set $\mathcal{S}$ of size at most $2^j \cdot \ln n$. This set is analogous to the set S from Lemma 3.2. We need to estimate the probability that no element of $\mathcal{S}$ covers v (which we call the *probability of silence*). Later, we consider a subset $\mathcal{R}$ of size at least 2^j and discuss the probability that v has at least two neighbors in $\mathcal{R}$.

4.1 Probability of silence

Lemma 3.2 shows that there exists a relatively small set S with useful properties. The probability that no element of S is connected to v is at least $1/n^{1+o(1)}$. In the construction of Alon et al. [1] computing this probability of silence is not difficult, because each vertex of $\mathcal{Y}_j$ (for the same index j) is connected to v with probability $\frac{1}{2^j}$ *independently* of other vertices.

In our reduction it is *not necessarily true* that every small enough set $\mathcal{S}$ does not cover v with a non-zero probability. For example, suppose that $v \in M_a$, and $\mathcal{S}$ contains all the copies of a. Then, by definition of the random permutation step, the copies of a cover v *with probability 1.* Further, if $M_b = M \setminus \mathcal{S}$, and $\mathcal{S}$ contains all the copies of b, then $\mathcal{S}$ covers the entire set M with probability 1.

On the other hand, the cover of M that we have just described uses a matching pair (a, b). Thus, intuitively, our goal is to prove that any set $\mathcal{S}$ that does not use matching pairs does not cover v with a non-negligible probability.

For the probability of the event: "$\mathcal{S}$ does not cover v in the random permutation step" to be greater than zero, we need v to satisfy the following property:

The safety property: *If $v \in M_x$ then $\mathcal{S}$ contains "only a fraction of" the copies of x. Alternatively, if all the copies of some x belong to $\mathcal{S}$ then $v \notin M_x$ must hold.*

This is further formalized in the following definition:

Definition 4.1 *The partitions defined by the exact partition steps (namely, in Steps 1 and 2 of the random permutation step) are* safe *for $\mathcal{S}$ and v if for every x so that $v \in M_x$,*

$$|\mathcal{S} \cap cp(x, j)| \leq 2^j/8.$$

We shall see that if the partition is safe for $\mathcal{S}$ and v, then with a non-negligible probability all the various random permutation steps do not cover v. The following lemma formalizes this claim.

We use the notation $\mathcal{S}\ \mathcal{AN}\ v$ to denote the event "no vertex of $\mathcal{S}$ is connected to v".

Lemma 4.2 *Suppose that the partitions formed by the exact partition steps are safe for $\mathcal{S}$ and v. Suppose also $|\mathcal{S}| \leq 2^j \cdot \ln n$. Then:*

$$\mathbb{P}\left((\mathcal{S}\ \mathcal{AN}\ v)\right) \geq \frac{1}{n^4}.$$

Proof: We use the following notation in the proof: S_N is the set of copies $x_q \in \mathcal{S}$ of some vertex x that *cannot* be connected to v in the random permutation step. Namely, either

$v \in M \setminus M_x$ or $q > 2^j$ (x_q is not one of the 2^j first copies of x). The complement set is denoted by $S_Y = \mathcal{S} \setminus S_N$.

For every vertex $x \in \mathcal{S}$, let $S(x)$, be defined by $S(x) = cp(x,j) \cap S_Y$ (these are copies that can be connected to v by the random permutation step). Let $s(x) = |S(x)|$. Clearly, $s(x) \leq 2^j/8$. See an illustration for this notation in Figure 4.

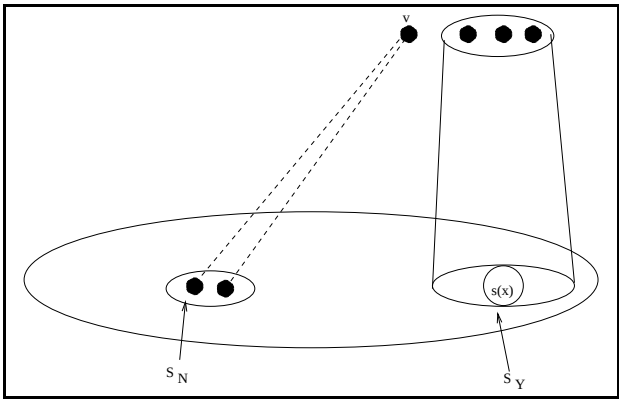

Figure 4: The effect of partition on M_x and $M \setminus M_x$, in addition to the random permutation step and mixing step on $\mathcal{S}, v$. The set $\mathcal{S}$ is partitioned to S_N and S_Y. The vertices of S_N cannot be connected to v by the random permutation step, because v does not belong to "their half" of M. Only vertices that belong to the set S_Y may be connected to the vertex v by the permutation step. For x, vertices of $S(x)$ can be connected to v by the random permutation step. The figure illustrates the "silent scenario", hence the vertices of $S(x)$ are not connected to the vertex v by the random permutation step and all the dotted lines represent non-edges.

The probability that no copy of x *that belongs to* $S(x)$ is connected to v in the random permutation step is:

$$\frac{2^j - s(x)}{2^j}.$$

This is because 2^j copies of x participate in the random permutation step and only $s(x)$ of them belong to S_Y.

For two different vertices $x, y \in S_Y$, note that the choices of the random permutations of x and y are independent (this is true even if they are leaves in the same star). Let $\mathcal{A}$ be the event that S_Y does not cover v in the random permutation step. Hence:

$$\mathbb{P}(\mathcal{A}) \;\geq\; \Pi_{x \in \mathcal{S}} \frac{2^j - s(x)}{2^j} \geq \Pi_{x \in \mathcal{S}} \left(\left(1 - \frac{s(x)}{2^j} \right) \cdot e^{-1} \right)^{s(x)/2^j}$$

The last inequality is because for any positive real $u > 0$, $(1 - 1/u)^{u-1} \geq 1/e$, and so $1 - 1/u \geq ((1 - 1/u) \cdot e^{-1})^{1/u}$.

As $s(x) \leq 2^{j-3}$,

$$\left(1 - \frac{s(x)}{2^j}\right) \cdot e^{-1} \geq \frac{7}{8 \cdot e}.$$

Hence,

$$Pr(\mathcal{A}) \geq \left(\frac{7}{8 \cdot e}\right)^{\sum_x s(x)/2^j} \geq \left(\frac{7}{8 \cdot e}\right)^{\ln n} \geq 1/n^2.$$

(The second inequality follows as $\sum_{x \in \mathcal{S}} s(x) \leq 2^j \cdot \log n$.)

The vertices of S_N are independently connected to v with probability $1/2^j$. (This follows from the mixing step of the reduction). Let $\mathcal{B}$ be the event that the mixing step does not form an edge between s and v. Hence, $\mathbb{P}(\mathcal{B}) \geq (1 - 1/2^j)^{|\mathcal{S}|} \geq (1 - 1/2^j)^{2^j \log n} \geq 1/n^2$. Finally, the event "no vertex of the set $\mathcal{S}$ is connected to the vertex v by the mixing step" is independent of the event "no vertex of the set $\mathcal{S}$ is connected to the vertex v by the permutation step". Hence,

$$\mathbb{P}(\mathcal{S} \; \mathcal{AN} \; v) \geq \frac{1}{n^2} \cdot \frac{1}{n^2} = \frac{1}{n^4} \; . \qquad \blacksquare$$

The probability for a safe partition: The partitions of M_x can be not safe with probability 1 for some "problematic" sets $\mathcal{S}$. In fact, one can easily guarantee that v *is covered* by the random permutation step. This can be achieved by taking into $\mathcal{S}$ all copies of a and all copies of b (recall that (a, b) is a $\tilde{e}$ matching pair).

The following definition utilizes this idea.

Definition 4.3 *A set $\mathcal{S}$ is* $(\tilde{e}, j)$-partial *if for every $\tilde{e}$-m.p. (x, y), the set $\mathcal{S}$ contains at most 2^{j-3} vertices of $cp(x, j)$* or *it contains at most 2^{j-3} copies of $cp(y, j)$.*

If $\mathcal{S}$ is $(\tilde{e}, j)$-partial, it is still possible that after the coins are tossed in the exact partition step, v will not be covered by $\mathcal{S}$ in the random permutation step. Namely, we shall see that if $\mathcal{S}$ is $(\tilde{e}, j)$-partial, the partition is safe with a non-negligible probability.

Definition 4.4 *Let $\tilde{e} = (A, B)$ be a superedge. Let $a \in A$. The set $star(a, \tilde{e})$ is the set of all copies of a and all copies of leaves b in an $\tilde{e}-$star of A.*

Lemma 4.5 *Let $\mathcal{S}$ be an $(\tilde{e}, j)$-partial of size $|\mathcal{S}| \leq 2^j \ln n$. The probability that the partition is safe for $\mathcal{S}, v$ is at least $1/n^8$.*

Proof: Since $|\mathcal{S}| \leq 2^j \cdot \ln n$, and the stars $star(a')$ with different vertices $a' \in A$ are all disjoint, the number of such stars that satisfy $|star(a) \cap \mathcal{S}| \geq 2^{j-3}$ is at most $8 \cdot \ln n$. Throughout the proof of this lemma, we will call such stars *dangerous*. Note that stars that are not dangerous can not make the partition unsafe.

Consider a dangerous star $star(a')$. Since the set $\mathcal{S}$ is $(\tilde{e}, j)$-partial, *either* it contains at most 2^{j-3} j-relevant copies of the vertex a', *or for all* vertices c in the star of a', the set $\mathcal{S}$ contains at most 2^{j-3} copies of the vertex c. In any case, with probability $1/2$, v belongs to the "right half" of M. (For example, if $\mathcal{S}$ contains at most $2^j/8$ copies of a' then $v \in M_{a'}$). Since the number of dangerous stars is at most $8 \cdot \ln n$, it follows that

$$Prob\,(\text{ safe partition }) \geq \left(\frac{1}{2}\right)^{8 \ln n} > \frac{1}{n^8},$$

proving the claim. ∎

The following corollary is immediate

Corollary 4.6 *Let $\mathcal{S}$ be an $(\tilde{e}, j)$-partial set of size $|\mathcal{S}| \leq 2^j \ln n$. Then $\mathbb{P}(\mathcal{S}\ \mathcal{AN}\ v) \geq 1/n^{12}$.*

Proof: By Lemma 4.5, with probability $1/n^8$, the partitions of M_x are safe with respect to the set $\mathcal{S}$, i.e., for every vertex $x \in S_Y$, $s(x) \leq 2^{j-3}$. By Lemma 4.2,

$$\mathbb{P}(\mathcal{S}\ \mathcal{AN}\ v \mid \mathcal{S} \text{ safe partition}) \ \geq\ 1/n^4\ .$$

Hence, with probability at least $1/n^{12}$, the set $\mathcal{S}$ induces a safe partition, and the event $(\mathcal{S}\ \mathcal{AN}\ v)$ holds. ∎

4.2 Probability of a collision

In this section we consider sets $\mathcal{R} \subseteq \mathcal{V}_1$ of size at least 2^j that are analogous to R'_i in Lemma 3.2. In all the following sets, we are interested in the event that $\mathcal{R}$ does not inform v because it covers v at least twice (namely, v has at least two neighbors in $\mathcal{R}$). Let $(\mathcal{R}\ 2\mathcal{C}\ v)$ denote this event.

Clearly, if all the relevant events were independent, namely, if every vertex of $\mathcal{R}$ was connected to v with probability $1/2^j$ independently of other vertices then

$$\mathbb{P}(\mathcal{R}\ 2\mathcal{C}\ v)\ =\ 1 - \left(1 - 1/2^j\right)^{|\mathcal{R}|} - \left(1 - 1/2^j\right)^{|\mathcal{R}|-1} \cdot |\mathcal{R}| \qquad (1)$$

This is in fact the case in the [1] construction.

We shall now see that in our construction, despite its dependencies, a similar inequality can be proven.

Consider a vertex x that contributes copies to $\mathcal{R}$, and suppose first that $v \notin M_x$. For such x, the edges between its copies and the vertex v are determined by the mixing step, and they behave exactly as in inequality (1). Similarly, if x_q is a copy of x, and $x_q \in cp(x) \setminus cp(x, j)$, and the probability that the edge (x, v) is in the graph is $1/2^j$ (see the mixing step).

However, the more delicate case is when $v \in M_x$. In this case the edges between the copies of x and the vertex v are determined by the random permutation step.

Let $X = \{x_1, \ldots, x_p\}$ be the set of copies of x in $\mathcal{R}$. First, we compute an upper bound on the probability of the event $(X \ \mathcal{AN} \ v)$, namely, that no vertex in X is connected to v by the random permutation step. Let $p' = \lceil p/2 \rceil$ and $X' = \{x_1, \ldots, x_{p'}\}$. Then

$$\begin{aligned} \mathbb{P}(X \ \mathcal{AN} \ v) \leq \mathbb{P}(X' \ \mathcal{AN} \ v) &= \mathbb{P}(x_1 \ \mathcal{AN} \ v) \cdot \mathbb{P}(x_2 \ \mathcal{AN} \ v \mid x_1 \ \mathcal{AN} \ v) \\ &\ldots \ \mathbb{P}(x_{p'} \ \mathcal{AN} \ v \mid \{x_1, \ldots, x_{p'-1}\} \ \mathcal{AN} \ v) \ . \end{aligned}$$

If it is known that i of the copies of x are not connected to v by the random permutation step, then all the *rest* of the $2^j - i$ copies have equal probability of covering v. Thus, the probability that the next copy x_{i+1} is connected to v is:

$$\frac{1}{2^j - i} \leq \frac{1}{2^{j-1}}.$$

This is because $i \leq p' - 1 \leq p/2 \leq 2^j/2$. Thus, we get that:

$$\mathbb{P}(X \ \mathcal{AN} \ v) \leq \left(\frac{1}{2^{j-1}}\right)^{p'}$$

In summary, the contribution of X to the probability is very similar to its contribution in inequality (1). The differences are: $1/2^{j-1}$ instead of $1/2^j$, and p' (recall that $p' \geq p/2$) instead of p. We derive an upper bound on the probability that v is informed by $\mathcal{R}$ in a similar way. Let $\rho = |\mathcal{R}|/2$. We have proved:

Lemma 4.7

$$\mathbb{P}(\mathcal{R} \ 2\mathcal{C} \ v) \geq 1 - \left(1 - 1/2^{j-1}\right)^{\rho} + \rho \cdot \left(1 - 1/2^{j-1}\right)^{\rho-1} \qquad \blacksquare$$

4.3 Deriving a lemma similar to Lemma 3.2

The pivot and the most significant index for $\tilde{e}$: For the rest of the section, consider a fixed short schedule $\Pi = (T_1, T_2, \ldots)$. We first define how to find the most "important" index j for $\tilde{e}$. Recall that $cp(A)$ (respectively, $cp(B)$) is the set of all copies of vertices of A (respectively, of vertices of B). Let $\Pi(\tilde{e})$ be the schedule $(R_1, R_2, \ldots)$ with $R_i = T_i \cap (cp(A) \cup cp(B))$. For the rest of the subsection, we use symbols R_i and $\mathcal{R}$ to denote rounds that are subsets of $cp(A) \cup cp(B)$. Recall that $|cp(A) \cup cp(B)| = n$ (because of the dummy vertices). Hence $cp(A) \cup cp(B)$ can play the role of the set $\mathcal{X}$ in Lemma 3.2.

Definition 4.8 *The index j whose existence is guaranteed by Lemma 3.2 with respect to $\mathcal{X} = cp(A) \cup cp(B)$ and the rounds $\Pi(\tilde{e})$ is called the* pivot *of $\tilde{e}$.*

For the rest of the section we adopt a notation from the paper of Alon et al. [1]; let $\mathcal{S}$ be the set whose existence is guaranteed by Lemma 3.2 (i.e., it plays the role of S from Lemma 3.2), and $R_i' = R_i \setminus \mathcal{S}$.

The probability of 2-covering: The proof of the following lemma is very similar to the proof of Lemma 3.4 from [1]. The only difference between the two proofs is that in Lemma 3.2, the following inequality holds with respect to the schedule Π

$$\mathbb{P}(\mathcal{R} \ 2\mathcal{C} \ v) \ \geq \ 1 - \left(1 - 1/2^j\right)^{|\mathcal{R}|} - \left(1 - 1/2^j\right)^{|\mathcal{R}|-1} \cdot |\mathcal{R}|.$$

This is because all the relevant events are independent. In our case, we use Lemma 4.7 that shows that though the events are not independent, a similar inequality holds. To summarize,

Lemma 4.9 *There is some universal constant c_1 so that*

$$\mathbb{P}(\ \textit{For all } i, R_i' \ 2\mathcal{C} \ v) \geq \frac{1}{n^{c_1}}. \qquad \blacksquare$$

How can we force $\mathcal{S}$ to be partial? In order to apply Corollary 4.6 to bound the probability that v is informed we need the subset $\mathcal{S}$ (from Lemma 3.2) to be $\tilde{e}-$partial. How can we guarantee that? One way of ensuring this is by requiring that $\bigcup R_i$ is $\tilde{e}-$partial.

To understand our approach, assume for the moment that indeed $\bigcup R_i$ is $\tilde{e}-$partial. Then, we can derive a lemma similar to Lemma 3.2, that is, show that an $\tilde{e}-$partial schedule cannot cover all the vertices of $M_{\tilde{e}}$. We want to use this claim to get a good MIN-REP solution along the following lines:

1. With high probability $\bigcup R_i$ cannot be $\tilde{e}-$*partial*, because of a lemma similar to Lemma 3.2.
2. This will hold in a similar way to many other super-edges.
3. As $\bigcup R_i$ is not $\tilde{e}-$partial, there should exist an $\tilde{e}-$matching pair (x, y) so that $\bigcup R_i$ contains at least $2^j/8$ copies of x and at least $2^j/8$ copies of y. If this is the case, we say that Π *chose* x, y.
4. If $\mu = |\bigcup R_i|$ is "small" then $\mu/2^j$ is "small" as well. Hence Π can choose only a few matching pairs from $A \cup B$.
5. Hence, a small subset of $A \cup B$ can be used to cover $\tilde{e}$.

6. Since this applies to "many" super-edges, we obtain a small solution for the original instance of the MIN-REP problem.

The problem with this scenario is in the case that $\mu = |\bigcup R_i|$ is "too large". In this case $\mu/2^j$ can be very large by itself, and so the the MIN-REP solution that will be derived may be large.

Indeed, it turns out that expecting that $\bigcup R_i$ is $\tilde{e}-$partial is too harsh a requirement, at least as far as very large rounds are present. The good news are, however, that large rounds have little effect because of Lemma 4.7. We next formalize this intuition. Again, we restrict our attention to $\Pi(\tilde{e})$, and consider rounds R that are subsets of $cp(A) \cup cp(B)$.

Definition 4.10 *We say that a round R is $(\tilde{e}, j)-$small if*

$$|R| \leq c \cdot 2^j \ln n \ ,$$

where c is some constant to be determined later. Let $Small(\Pi, j, \tilde{e})$ be the collection of small rounds of $\Pi(\tilde{e})$, and $Large(\Pi, j, \tilde{e})$ be the set of all other rounds.

Definition 4.11 *Let*

$$\mathcal{W}(\Pi, j, \tilde{e}) = \bigcup_{R_i \in Small(\Pi, j, \tilde{e})} R_i$$

Definition 4.12 *A schedule Π is $(\tilde{e}, j)-$partial if $\mathcal{W}(\Pi, j, \tilde{e})$ is $\tilde{e}-$partial.*

Note that even if Π is $\tilde{e}-$partial, still $\bigcup R_i$ may contain all the copies of *both* a and b, for an $\tilde{e}$-matching pair (a, b). This is because some rounds R_i may be large.

Assume $c_0 \geq c_1 + 12$, where c_1 is the constant from Lemma 4.9. The following corollary (it is analogous to Lemma 3.2) is derived from Lemma 4.9 and Corollary 4.6. However, it applies only to $Small(\Pi, j, \tilde{e})$.

Corollary 4.13 *Let Π be $\tilde{e}-$partial. With probability at least $1/n^{c_0}$, $Small(\Pi, j, \tilde{e})$ does not inform v.*

Proof: We mimic the proof of Lemma 3.2. Namely, we compute the probability for the event $\mathcal{N}$ = "no vertex of $\mathcal{S}$ is connected to v" and the event $2\mathcal{C}$ = "all small rounds $R_i' = R_i \setminus \mathcal{S}$ cover v at least twice". If both $\mathcal{N}$ and $2\mathcal{C}$ occur, then v is not informed by $Small(\Pi, j, \tilde{e})$.

Proving that the event "for every i, R_i' $2-$covers v" occurs with probability at least $1/n^{c_1}$ is done exactly as in Lemma 4.9, and in [1].

Now we deal with $\mathcal{S}'$. As Π is $\tilde{e}$-partial, by definition $\mathcal{W}$ is $\tilde{e}$-partial, and thus $\mathcal{S}'$ is

$\tilde{e}$-partial. Thus, from Corollary 4.6,

$$\mathbb{P}(\mathcal{S}' \ \ \mathcal{AN} \ \ v) \geq \frac{1}{n^{12}} \ .$$

Note that R'_i and $\mathcal{S}'$ are disjoint. But the events $\mathcal{N}$ and $2\mathcal{C}$ are *not* independent, as the two sets may contain copies of the same vertex. We need to study the correlation between $\mathcal{N}$ and $2\mathcal{C}$. We shall now see that the events are positively correlated.

Let R_x (resp., S_x) be the subset of copies of x that belong to R'_i (resp., $\mathcal{S}$). If $v \notin M_x$ then the edges between the vertices of R_x and S_x on the one hand and the vertex v on the other are defined by the mixing step and are completely independent. If $v \in M_x$, then the connection of $R_x \cup S_x$ and v is determined by the random permutation step. Now, if S_x is not connected to v this only *increases* the probability that R_x is connected to v. Hence, the correlation between these probabilities is positive.

Hence, with probability at least $\frac{1}{n^{c_1+12}} \geq \frac{1}{n^{c_0}}$ v is not informed by $Small(\Pi, j, \tilde{e})$. ∎

5 Analysis part II: deriving the result

Let Π be an $\tilde{e}-$*partial* short schedule.

5.1 How many vertices can $Small(\Pi, j, \tilde{e})$ inform?

We consider the number of vertices informed by $Small(\Pi, j, \tilde{e})$. (At this point we ignore the contribution of $Large(\Pi)$. We will deal with it later).

Let q' be some index. We say that a set $M_{\tilde{e}}(j, q')$ is *fully informed* by $Small(\Pi, j, \tilde{e})$ if all the elements of $M_{\tilde{e}}(j, q')$ are informed. Let $NF = NF(\tilde{e}, j, Small(\Pi, j, \tilde{e}))$ be the number of indices q' for which $M_{\tilde{e}}(j, q')$ is *not* fully informed by $Small(\Pi, j, \tilde{e})$.

Lemma 5.1

$$\mathbb{P}(NF < n^2) \leq exp(-\Omega(n^3)).$$

Proof: Consider a fixed index q'. By Corollary 4.13 and the Markov inequality,

$$\mathbb{P}\left(M_{\tilde{e}}(j, q') \text{ is not fully informed }\right) \geq \frac{1}{n^{c_0-1}} \ .$$

Hence, the number of not fully informed sets $M_{\tilde{e}}(j, q')$ is a Binomial variable with success probability greater or equal to $1/n^{c_0-1}$. The number of different indices q' is n^{c_0+4}. Thus,

the expected number of not fully informed $M_{\tilde{e}}(j,q')$ is at least n^3. Hence, the claim follows from the Chernoff bound. ∎

The lemma shows that $Small(\Pi, j, \tilde{e})$ not only does not inform all the vertices, but also leaves *many* indices q' for which $M_{\tilde{e}}(j,q')$ has at least one non-informed element. In fact, a simple counting argument and the union-bound imply the following corollary.

Corollary 5.2 *With probability* $1 - exp(-\Omega(n^3))$, *for* any $\tilde{e}-$*partial short schedule* Π, $NF(\tilde{e}, j, Small(\Pi, j, \tilde{e})) \geq n^2$.

Proof: The set of relevant vertices on a fixed round of Π is a subset of $cp(A) \cup cp(B)$. The size of $cp(A) \cup cp(B)$ is n. Hence, the number of subsets of $A \cup B$ is 2^n. Thus the number of short schedules is at most $2^{n \log^2 n}$. Since $2^{n \log^2 n} \ll exp(n^3)$, the claim follows from the union-bound. ∎

5.2 Large rounds are not able to inform many vertices

By Corollary 5.2, with probability $1-exp(-\Omega(n^3))$, no short partial schedule satisfies $NF \leq n^2$. We next show that for any choice of Π, with high probability $Large(\Pi)$ can not "complete the task" and leaves some vertices uninformed.

Lemma 5.3 *If* Π *is* $\tilde{e}-$*partial short schedule then with probability at least* $1 - 2^{-2n^2}$, Π *does not inform all the vertices of* $M_{\tilde{e}}(j)$.

Proof: We know that (with high probability) there is a set $\mathcal{U}$ of n^2 elements, $\mathcal{U} = \{u_1, \ldots, u_{n^2}\}$, that are not informed by $Small(\Pi, j, \tilde{e})$. The crucial property of $\mathcal{U}$ is that different vertices u_i belong to different sets $M_{\tilde{e}}(j,q')$. Hence the random events that we consider are independent. We fix the sets $Large(\Pi)$ and $\mathcal{U}$, and estimate the probability that $Large(\Pi)$ covers $M_{\tilde{e}}$.

Let $\mathcal{R} \subset A \cup B$ be a large round in $\Pi(\tilde{e})$. Let $r = \log_2 |\mathcal{R}|$. By definition, $\mu = |\mathcal{R}| \geq c \cdot 2^{j+1} \ln n$. By Lemma 4.7, setting

$$\rho = \frac{|\mathcal{R}|}{2},$$

we get:

$$\mathbb{P}(\mathcal{R}\ 2\mathcal{C}\ v) \geq 1 - \left(1 - \frac{1}{2^{j-1}}\right)^{\rho} + \rho \cdot \left(1 - \frac{1}{2^{j-1}}\right)^{\rho-1} .$$

Thus, it follows that the probability that some u_i is informed by $\mathcal{R}$ is at most $1/n^{c'}$ with c' being some universal constant (that depends on c from Definition 4.10). Since the edges between the vertices of $\mathcal{U}$ and v are independent (different u_i belong to different sets $M_{\tilde{e}}(j,q')$), the probability that the entire set $\mathcal{U}$ is informed is at most $1/n^{c'n^2}$.

Now, we count the number of possibilities to choose the sets $Large(\Pi)$ and $\mathcal{U}$. Note that $\mathcal{U}$ is a subset of size n^2 chosen out of a set of n^{c_0+4} vertices. The number of ways to do so is at most $n^{(c_0+4)n^2}$. The number of possible $Large(\Pi)$ schedules (restricted to subsets of $cp(A) \cup cp(B)$) is $O(2^{n \cdot \log^2 n})$. Thus the number of choices of $\mathcal{U}$ and $Large(\Pi)$ is at most $n^{(c_0+5)n^2}$. We set c so that $c' \geq c_0 + 7$. Now the claim follows from the union-bound. ∎

5.3 Short proper schedules can not be feasible

We need the following definition. Intuitively, it defines short schedules that are partial for *many* super-edges.

Definition 5.4 *A short schedule* Π *is called* proper *if there exists a subset* $E' \subseteq \tilde{E}$ *that contains at least one half of all the super-edges, and such that the* $Small(\Pi, j, \tilde{e})$ *is* $\tilde{e}'-$*partial for every* $\tilde{e}' \in E'$*. Otherwise, the schedule* Π *is called* non-proper.

The following lemma holds both for yes and no instances of the MIN-REP problem.

Lemma 5.5 *With probability* $1 - exp(\Omega(n^2))$*, no proper* Π *informs all the vertices of* $\mathcal{V}_2$.

Proof: First, fix E'. For a single super-edge $\tilde{e}' \in E'$, the probability that $M_{\tilde{e}}$ is fully informed is at most 2^{-2n^2} (Lemma 5.3). Naturally, this also implies an upper bound on the probability that all the vertices of $\mathcal{V}_2$ are informed.

By Restriction 2 in the definition of the MIN-REP problem, the number of super-edges is bounded by $n \cdot polylog(n)$. Thus, the number of subsets of the superedges is at most $2^{o(n^2)}$. By the union-bound the probability that *there exists* a subset E' so that for every $\tilde{e}' \in E'$, $M_{\tilde{e}'}$ is fully informed is at most:

$$2^{o(n^2)} \cdot 2^{-2n^2} = exp(-\Omega(n^2)) \ . \qquad \blacksquare$$

Remark: Note that a schedule of logarithmic length for a yes instance that was described in Section 3.7 is *not proper*.

5.4 For no instances short non-proper schedules are not feasible

Lemma 5.6 *With probability 1 there is no non-proper feasible short schedule* Π *for the instance derived out of a no instance of the MIN-REP problem.*

Proof: Suppose for contradiction that there exists a schedule Π as above. Assume, without loss of generality, that the first round of the schedule Π is the set $\{s\}$, and that all the other

rounds $R \in \Pi$ are subsets of the set $\mathcal{V}_1$ (with no dummy vertices). Let $E_N \subseteq \tilde{E}$ be a subset of super-edges, such that for every super-edge $\tilde{e} \in E_N$, the schedule Π is not $\tilde{e}$-partial. By definition, the set E_N contains at least half of the super-edges.

For a super-edge $\tilde{e} \in E_N$, let j be its pivot. Recall that $\mathcal{W}$ is the union of all small rounds. By definition, $\mathcal{W}$ contains at least $2^j/8$ copies of both x and y for some $\tilde{e}-$matching pair (x, y). We call this pair "the important pair for $\tilde{e}$".

We now define a MIN-REP solution C that is both "of small size" and covers all the super-edges of E_N. C is defined by the following procedure.

1. Go over all the super-edges in E_N in an arbitrary order.
2. For a super-edge $\tilde{e}$, let (x, y) be the important pair for $\tilde{e}$.
3. Add x and y to C.

The following claim is immediate by definition:

Claim 5.7 *C covers all the super-edges E_N.*

We now bound $|C \cap A|$ for an arbitrary super-vertex A.

Since the super-vertex A participates in several different super-edges, and each with its own pivot, we consider every index j separately. Fix some pivot j, and bound the contribution to $C \cap A$ due to j. Recall that the subschedule $Small(\Pi, j, \tilde{e})$ contains only the rounds R that are j-small with respect to the super-edge $\tilde{e}$, i.e., the rounds $R \in \Pi(\tilde{e})$, that satisfy $|R| \leq 2^{j+1} \cdot c \cdot \ln n$. Also, the number of rounds in the schedule is $O(\log^2 n)$. Hence, $|\mathcal{W}| = O(2^j \cdot \log^3 n)$ (because $\mathcal{W}$ is the union of all small rounds.)

Every super-edge (A, B) with pivot j causes the important $\tilde{e}-$pair (a, b) to be added into C. But, by definition, $\mathcal{W}$ contains at least $2^j/8$ copies of a and $2^j/8$ copies of b. In particular, the number of vertices a that can be added to C with pivot j is at most

$$\frac{|\mathcal{W}|}{2^j/8} = O(\log^3 n).$$

This bounds the contribution of j to $|A \cap C|$.

Summing over all different indices j, the total size of $A \cap C$ is bounded by $O(\log^4 n)$. Similar bound follows for every B.

In other words, we have shown that the set C covers at least one half of all the super-edges of the instance $\mathcal{M}$ of the MIN-REP problem, and contains $O(\log^4 n)$ representative

vertices from each super-vertex. It follows that no A or B sets contributes more than $\log^{10} n$ vertices to C. This contradicts Theorem 2.5 ∎

Corollary 5.8 *With high probability, for a no instance no short schedule* Π *is feasible.*

Proof: By Lemma 5.6, with probability 1, there is no non-proper feasible short schedule for the instance $\mathcal{G}$. By Lemma 5.5, with high probability there is no proper feasible short schedule for the instance $\mathcal{G}$. Since any schedule is either proper or non-proper, the assertion follows. ∎

Hence, we have shown that the reduction has the claimed gap, and have proved our main result.

Theorem 5.9 *Unless* $NP \subseteq BPTIME(n^{O(\log \log n)})$*, for some universal constant* c *there is no additive* $(c \cdot \log^2 n)$*-approximation for the radio broadcast problem.*

Acknowledgments

The first-named author wishes to thank Oded Regev for helpful discussions.

References

[1] 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.

[2] S. Arora, L. Babai, J. Stern, Z. Sweedyk. The hardness of approximate optima in lattice codes. In *Proc. Symp. on Foundations of Comp. Science*, pp. 724-733, 1993.

[3] S. Arora and K. Lund. In *Approximation Algorithms for NP-hard Problems*, D. Hochbaum, ed., PWS Publishing, 1996.

[4] 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.

[5] B. S. Chlebus, L. Gasieniec, A. Gibbons, A. Pelc, W. Rytter. Deterministic broadcasting in ad hoc radio networks. Distributed Computing 15(1): 27-38 (2002)

[6] M. Chrobak, L. Gasieniec, W. Rytter. Fast Broadcasting and Gossiping in Radio Networks. In *Proc. 41st Symp. on Foundations of Computer Science*, pp. 575-581, Redondo Beach, CA, Nov. 2000.

[7] I. Chlamtac and S. Kutten. A spatial-reuse TDMA/FDMA for mobile multihop radio networks. In *Proc. IEEE INFOCOM*, pp. 389-94, 1985.

[8] A. F. Clementi, A. Monti, R. Silvestri. Distributed broadcast in radio networks of unknown topology. TCS 1-3(302): 337-364 (2003), Preliminary version appeared in *Proc. of the 14th ACM-SIAM Symp. on Discr. Algorithms*, pp. 709-718, Washington, DC, Jan. 2001.

[9] A. Czumaj and W. Rytter. Broadcasting Algorithms in Radio Networks with Unknown Topology. In *Proc. 44th Symp. on Foundations of Computer Science*, pp. 492-501, 2003.

[10] Chlamtac and Weinstein The wave expansion approach to broadcasting in multihop radio networks, In *Proc. IEEE INFOCOM*, pp. 874-881, 1987.

[11] M. Elkin and G. Kortsarz. A logarithmic lower bound for radio broadcast. J. Algorithms, vol 52, num 1, 8-25, 2004.

[12] M. Elkin and G. Kortsarz. An improved algorithm for radio broadcast, in *Proc. of ACM-SIAM Symposium on Discrete Algorithms*, pp. 222-231, Vancouver, British Columbia, Canada, 2005.

[13] U. Feige. A threshold of $\ln n$ for approximating set cover, Journal of ACM, 45(4), 634-652, July 1998.

[14] N. Garg, G. Konjevod, R. Ravi. A Polylogarithmic Approximation Algorithm for the Group Steiner Tree Problem. J. Algorithms 37(1), pp. 66-84, 2000.

[15] I. Gaber and Y. Mansour. Broadcast in radio networks. In *Proc. of the 6th ACM-SIAM Symp. on Discr. Algorithms*, pp. 577–585, 1995.

[16] R. Gandhi, S. Parthasarathy and A. Mishra. Minimizing Broadcast Latency and Redundancy in Ad Hoc Networks. In Proc. of the Fourth ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC'03), pages 222-232, Jun. 2003.

[17] E. Halperin and R. Krauthgamer. Polylogarithmic inapproximability, in *Proc. of the 35th Symp. on Theory of Computing*, pp. 585-594, 2003.

[18] P. Indyk. Explicit constructions of selectors and related combinatorial structures, with applications. In *Proc. 15th Symp. on Discr. Algorithms*, pp. 697-704, 2002.

[19] Kushilevitz and Mansour. An $\Omega(D \log(n/D))$ lower bound for broadcast in radio networks. In *SIAM J. of Computing*, Vol. 27, No. 3, pp. 702-712, June 1998.

[20] D. Kowalski and A. Pelc, Deterministic broadcasting time in radio networks of unknown topology, In *Proc. 43rd Ann. IEEE Symposium on Foundations of Computer Science*, pp. 63-72, 2002.

[21] D. Kowalski and A. Pelc, Broadcasting in undirected ad hoc radio networks In *Proc. 16th ACM Symp. on Principles of Distr. Comp.*, pp. 73 - 82, 2003.

[22] D. Kowalski, A. Pelc., Centralized Deterministic Broadcasting in Undirected Multi-hop Radio Networks. APPROX-RANDOM, 171-182, 2004.

[23] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *J. Assoc. Comput. Mach.*, 41(5):960–981, 1994.

[24] R. Raz. A Parallel Repetition Theorem, *SIAM J. of Computing*, 27(3) pp. 763-803, 1998.