

Computing Large Independent Sets in a Single Round*

Magnús M. Halldórsson¹ and Christian Konrad²

¹ICE-TCS, School of Computer Science, Reykjavik University, Menntavegur 1, 101
Reykjavik, Iceland, mmh@ru.is

²Department of Computer Science, Centre for Discrete Mathematics and its Applications
(DIMAP), University of Warwick, Coventry, UK, c.konrad@warwick.ac.uk

Abstract

Independent sets play a central role in distributed algorithmics. We examine here the minimal requirements for computing non-trivial independent sets. In particular, we focus on algorithms that operate in a single communication round. A classic result of Linial shows that a constant number of rounds does not suffice to compute a *maximal* independent set. We are therefore interested in the *size* of the solution that can be computed, especially in comparison to the optimal.

Our main result is a randomized one-round algorithm that achieves poly-logarithmic approximation on graphs of *polynomially bounded-independence*. Specifically, we show that the algorithm achieves the Caro-Wei bound (an extension of the Turán bound for independent sets) in general graphs up to a constant factor, and that the Caro-Wei bound yields a poly-logarithmic approximation on bounded-independence graphs. The algorithm uses only a single bit message and operates in a beeping model, where a node receives only the disjunction of the bits transmitted by its neighbors.

We give limitation results that show that these are the minimal requirements for obtaining non-trivial solutions. In particular, a sublinear approximation cannot be obtained in a single round on general graphs, nor when nodes cannot both transmit and receive messages. We also show that our analysis of the Caro-Wei bound on polynomially bounded-independence graphs is tight, and that the poly-logarithmic approximation factor does not extend to $O(1)$ -claw free graphs.

1 Introduction

Something For Almost Nothing. When designing approximation algorithms, the usual goal is to find desirable trade-offs between approximation guarantee and the resources required by the algorithm, such as computation time, memory consumption, or, in the area of distributed computing, message size and the number of communication rounds. If only very limited access to computational resources is available, it is often asked how much effort it takes to obtain at least something from the given problem instance. In distributed computing, those limits are explored for example with regards to communication patterns and the total number of communication rounds. It has been shown that non-trivial computation is possible even when the communication pattern of nodes is restricted to beeps [10]. Moreover, highly non-trivial local algorithms [31, 35, 24] that employ only a constant number of communication rounds have been obtained (e.g. even some NP-hard problems can be solved by local algorithms [5, 6, 7]).

In this paper, we ask whether non-trivial computation is possible if we grant a distributed algorithm only a single communication round. Specifically, we ask whether reasonable approximations to the maximum independent set problem can be computed in this harsh setting.

*Supported by Icelandic Research Fund grants 120032011 and 152679-051. C. Konrad is also supported by the Centre for Discrete Mathematics and its Applications (DIMAP) at Warwick University and by EPSRC award EP/N011163/1.

Computational Model. We consider a network of computational units V modelled by a graph $G = (V, E)$, which also constitutes the problem input. Algorithms run in a single round: They first compute locally, then send a message to their neighbors, and after receiving they compute locally again and declare their output.

Our algorithm works in a restricted beeping model, where the message transmitted is a single bit sent to all neighbors. A node (whether transmitting or not) receives only the disjunction of the bits sent by the neighbors (or, equivalently, it learns whether *some* neighbor beeped or not). Additionally, the algorithm operates anonymously, without information about identifiers, port labels or orientations. We assume that, initially, each node $v \in V$ knows its degree $d_G(v)$. This is a non-standard assumption in beeping models, but is necessary for one-round algorithms with non-trivial approximation guarantees. Indeed, we give a proof in the appendix showing that if degree information is not provided, then every algorithm has an approximation ratio of $\Omega(\frac{n}{\log n})$.

For lower bounds, we assume the more powerful \mathcal{LOCAL} model, where nodes can send (receive) individual messages of unbounded sizes to (resp. from) their neighbors.

Independent Sets. An independent set I in a graph $G = (V, E)$ is a subset of pairwise non-adjacent vertices. An independent set I is maximal if it is inclusion-wise maximal, i.e., $I \cup \{v\}$ is not an independent set for any $v \in V \setminus I$. A maximum independent set is one of maximum cardinality. The independence number of graph G is the size of a maximum independent set in G and is denoted by $\alpha(G)$. Computing maximum independent sets is NP-hard on general graphs [22] and is even hard to approximate within factor $n^{1-\epsilon}$, for any $\epsilon > 0$ [18, 38].

Main Result. Our main result concerns graphs of *polynomially bounded-independence*, a graph class that includes unit disc graphs and similar graph classes that are used for modelling wireless networks (for a precise definition see the next paragraph). We show that in the harsh setting of a single communication round, a poly-logarithmic approximation ratio can be achieved in polynomially bounded-independence graphs. Furthermore, we show that not only the number of communication rounds but also message size can be reduced to an absolute minimum, i.e., to a single bit message.

Bounded-independence Graphs. Graphs of bounded-independence capture many intersection graphs of geometrical objects which are used for modelling conflict graphs of wireless networks. Given a collection $X = \{X_1, \dots, X_n\}$ of geometrical objects, the corresponding intersection graph on vertex set X is obtained by introducing an edge between two vertices X_i, X_j iff the objects X_i and X_j intersect. In the literature, conflict graphs of wireless networks are often modelled by unit disc graphs [33, 15], the intersection graph of discs with equal radii, where the radii of the discs correspond to the transmission range of the wireless transmitters. Unit disc graphs have many beneficial properties that allow for the design of efficient distributed algorithms, but the assumption of identical transmission radii for all wireless transmitters is often too restrictive. Consequently, the unit disc graphs model has been extended to more elaborate models such as quasi unit disc graphs [26] or general disc graphs. In a general disc graph, no restriction on the radii of the discs are imposed, but the parameter $\delta = r_{\max}/r_{\min}$ is introduced into the analysis of algorithms, where r_{\max} and r_{\min} denote the maximum and minimum radii of a disc, respectively.

All graphs of the graph classes mentioned above are of bounded-independence, a property that restricts the size of a maximum independent set within the set of nodes at a given maximal distance from any node. The (inclusive) r -neighborhood of a node v is the set of nodes at distance at most r from v (including v).

Definition 1 (Bounded-independence). *Graph $G = (V, E)$ is of bounded-independence if there is a bounding function $f(r)$ so that for each node $v \in V$, the size of a maximum independent set in the inclusive r -neighborhood of v is at most $f(r)$, $\forall r \geq 1$. We say that G is of polynomially bounded-independence if $f(r)$ is a polynomial.*

If G is of bounded-independence, then we say that G is a BI-graph, and if G is of polynomially bounded-independence, then we say that G is a PBI-graph.

It is easily verified that unit disc graphs are BI-graphs with respect to a bounding function in $O(r^2)$, and (general) disc graphs are BI-graphs with respect to a bounding function in $O((r\delta)^2)$. Many important problems such as the maximal independent set problem, or the $(\Delta + 1)$ -coloring problem can be solved on

BI-graphs in $O(\log^* n)$ rounds by an algorithm of Schneider and Wattenhofer [34], underlining the usefulness of this graph class for distributed computation.

Turán’s Bound and a One-round Algorithm. A starting point of our work is an extension of a celebrated theorem by Paul Turán. Turán showed that every graph $G = (V, E)$ contains an independent set of size at least $n/(\bar{d} + 1)$, where \bar{d} is the average degree of G [36]. This was extended by Caro [8] and Wei [37] who showed that G contains an independent set of size at least

$$\beta(G) := \sum_{v \in V} \frac{1}{d_G(v) + 1},$$

where $d_G(v)$ denotes the degree of vertex v in G . An independent set of expected size $\beta(G)$ can be found by a (folklore) simple linear time randomized algorithm that follows from the analysis of the Caro-Wei bound given by Alon and Spencer in [3]. This algorithm works as follows: Each node v chooses a random real value between 0 and 1 and adds itself to the independent set I if none of its neighbors have chosen a larger real value than v . Then, the probability that v is added to the independent set is $\frac{1}{d_G(v)+1}$, and, hence, by linearity of expectation, $\mathbb{E}|I| = \sum_{v \in V} \frac{1}{d_G(v)+1} = \beta(G)$.

This algorithm can also be implemented distributively in a single communication round. Instead of choosing a random real value, every node chooses a random value from a large enough ordered set (e.g. $\{1, 2, \dots, n^3\}$ suffices) so that neighboring nodes choose different values with large enough probability. In order to be able to determine such a number, nodes require knowledge of n , i.e., the order of the input graph. Furthermore, communicating the chosen value to neighboring nodes requires messages of size $O(\log n)$. In the following, we will refer to this algorithm as ALON-SPENCER-IS.

It is easy to see that in general graphs, an independent set of size $\beta(G)$ may be a factor $\Theta(n)$ smaller than the independence number $\alpha(G)$ ¹. This raises the following questions:

1. Are there interesting graph classes for which $\beta(G)$ is a non-trivial approximation to the independence number $\alpha(G)$?
2. What are the minimum communication requirements for achieving the $\beta(G)$ bound?
3. Is there a one-round independent set algorithm with approximation factor $o(n)$ on general graphs?

Our Results in Detail. Concerning Question 1, we prove that an independent set of size $\beta(G)$ is a poly-logarithmic approximation to a maximum independent set in PBI-graphs. For instance on unit disc graphs, an independent set of size $\beta(G)$ is an $O((\frac{\log n}{\log \log n})^2)$ -approximation to a maximum independent set. Moreover, we prove that our analysis is tight up to a constant factor by constructing families of d -dimensional unit sphere graphs, for any constant integer d . We also show that on the more general class of k -claw-free graphs², for $k \geq 3$, the Caro-Wei bound constitutes a $O(\sqrt{nk})$ -approximation, and this bound is tight.

With regards to Question 2, we show that the communication requirements can be reduced to an absolute minimum, at the price of a constant factor in the bound. We present a different and even simpler one-round algorithm that computes an independent set of expected size at least $0.267\beta(G)$ using a single bit message, thus decreasing the message sizes from $O(\log n)$ to 1. This algorithm has the additional advantage that it does not require knowledge of n . The latter property and the low communication requirements make an implementation in a beeping model possible (see Section 2 for a detailed discussion of the model requirements of our algorithm). Note that our main result, a poly-logarithmic approximation one-round single-bit message algorithm for the maximum independent set problem in PBI-graphs, follows from the previous two results.

Last, we answer Question 3 in the negative. We provide a lower bound that shows that any (possibly randomized) one-round algorithm has approximation ratio $\Omega(n)$, even on disc graphs.

¹Consider for instance a complete bipartite graph G with equally sized bipartitions on $2n$ vertices where edges are added that turn one bipartition into a clique. Then, $\alpha(G) = n$ and $\beta(G) = O(1)$.

²A graph is k -claw-free, if it does not contain the complete bipartite graph $K_{1,k}$ as an induced subgraph.

Further Related Work. Independent sets are among the most studied problems in distributed computing. However, most works consider the maximal independent set problem while this paper addresses the maximum independent set problem. It is known that in general graphs, computing a maximal independent set requires $\Omega\left(\min\left\{\sqrt{\frac{\log n}{\log \log n}}, \frac{\log \Delta}{\log \log \Delta}\right\}\right)$ rounds of communication [24], and even on a ring, $\Omega(\log^* n)$ rounds are necessary [27]. The currently fastest maximal independent set algorithm for general graphs of Ghaffari [12] runs in $O(\log \Delta) + 2^{O(\sqrt{\log \log n})}$ rounds. It is known that the $2^{O(\sqrt{\log \log n})}$ term is best possible unless the $2^{O(\sqrt{\log n})}$ deterministic maximal independent set algorithm of Panconesi and Srinivasan [32] can be improved [9].

Concerning approximations to the maximum independent set problem, a $O(n^\epsilon)$ -approximation can be computed in $O(\frac{1}{\epsilon})$ rounds in general graphs, and that is best possible [7]. In planar graphs, a $(1 + \epsilon)$ -approximation can be computed in $O(\log^* n)$ [11]. The $O(\log^* n)$ -round algorithm of [34] gives a constant-factor approximation in BI-graphs, since in this graph class, any maximal independent set is a constant approximation of a maximum independent set.

The study of constant-round distributed algorithms was proposed in [4, 31, 27], and today a multitude of such algorithms are known, as evidence by the survey of Suomela [35]. Few non-trivial one-round distributed algorithms are known. For example, Linial showed that a vertex coloring with $O(\Delta^2 \log n)$ colors can be computed in a single communication round [27]. Kuhn and Wattenhofer [25] proved that every coloring computed in one round uses $\Omega(\Delta^2 / \log \Delta + \log \log n)$ colors.

In recent years, numerous works have studied the maximal independent set problem in *beeping* models [10, 1, 20, 29, 19], but all those algorithms require a (poly-)logarithmic number of rounds.

Last, we note that the Caro-Wei and Turán bounds have previously been used as quality guarantees for independent set approximation (e.g., [13, 14]).

Notations. Throughout the paper, we use the following notations. Let $G = (V, E)$ be a graph with $n = |V|$. For a node $v \in V$, let $\Gamma_G(v)$ denotes the neighborhood of v and $d_G(v) = |\Gamma_G(v)|$ its degree. If the graph G in which vertex v appears is clear from the context, then we may also write $d(v)$ or $\Gamma(v)$ to denote v 's degree or neighborhood, respectively. The d -neighborhood of v , denoted $\Gamma_G^d(v)$, is the set of nodes of distance at most d from v excluding v , while the set of nodes at distance exactly d from v is denoted by $\Gamma_G^{(d)}(v)$. Let $\Gamma_G^d[v] := \Gamma_G^d(v) \cup \{v\}$ (and $\Gamma_G[v] = \Gamma_G(v) \cup \{v\}$). For a subset of vertices $U \subseteq V$, the graph $G[U]$ is the subgraph of G induced by the vertices U .

Outline. An algorithm with single-bit messages achieving the Caro-Wei bound up to a constant factor is presented in Section 2. It is then shown in Section 3 that the Caro-Wei bound is a poly-logarithmic approximation to the independence number in PBI-graphs. In Section 4, these results are shown to be the strongest such results possible in several different ways. We conclude in Section 5 and point out interesting research directions.

2 One-Round Algorithm With Single Bit Message

In this section, we give a randomized one-round algorithm that achieves the Caro-Wei bound up to a constant factor and is even simpler than the ALON-SPENCER-IS algorithm.

We will consider the one-round algorithm, Algorithm 1, which can be seen as a simplified version of a well-known distributed maximal independent set algorithm commonly referred to as Luby's algorithm [28] (a similar version of the algorithm was independently discovered at the same time by Alon, Babai and Itai [2]). In each round of Luby's algorithm, nodes of a general graph $G = (V, E)$ are added to an initially empty independent set. One round consists of two phases: First, each node $v \in V$ pre-selects itself with probability $\Theta(\frac{1}{d_G(v)})$ as a candidate to join the independent set. Then, in the second phase, ties are broken among the pre-selected nodes so that nodes with larger degree are preferred. Finally, selected nodes and their neighbors are removed from G , and the round is completed. The algorithm terminates when G is empty.

In our version of the algorithm, a simplified method for breaking ties is used. Instead of preferring nodes with larger degree, we only add a pre-selected node to the independent set if none of its neighbors have been pre-selected. This method of breaking ties has previously been used, e.g., in [16, 23, 17].

Algorithm 1 One-round independent set algorithm

Require: $G = (V, E)$ {Input graph}

- 1: $I \leftarrow \emptyset$ {the independent set to be computed}
 - 2: $p_i \leftarrow \frac{1}{d_G(v)+1}$
 - 3: $T_v \leftarrow \text{coin}(p_i)$ {Pre-selection step: If $T_v = \text{true}$ then v is a candidate to join I }
 - 4: **for all** $v \in V$ with $T_v = \text{true}$ **do**
 - 5: **if** $\bigvee_{u \in \Gamma_G(v)} T_u = \text{false}$ **then** {Check whether a neighbor of v has been pre-selected}
 - 6: $I \leftarrow I \cup \{v\}$ { v is selected into the IS}
 - 7: **end if**
 - 8: **end for**
-

In a distributed implementation of this algorithm, a node v *beeps* (broadcasts the same message 1 to all its neighbors) if $T_v = \text{true}$ and remains silent otherwise. In Line 5 of the algorithm, it is required that every candidate node v learns whether at least one neighbor emitted a beep signal. It is hence enough if nodes have the ability to learn the disjunction of incoming messages.

We will first prove that the algorithm achieves the Caro-Wei bound up to a constant factor and then discuss the precise model requirements of the algorithm.

Our main theorem relies on a technical lemma, which bounds a certain quantity away from zero and is proved first.

Lemma 1. *Let $G = (V, E)$ be any graph. Then:*

$$S(G) := \sum_{v \in V} \frac{1}{d(v)+1} \left(1 - \sum_{u \in \Gamma(v)} \frac{1}{d(u)+1} \right) > 0 .$$

Proof. First, observe that

$$1 - \sum_{u \in \Gamma(v)} \frac{1}{d(u)+1} > \sum_{u \in \Gamma(v)} \left(\frac{1}{d(v)+1} - \frac{1}{d(u)+1} \right),$$

which implies

$$S(G) > \sum_{v \in V} \frac{1}{d(v)+1} \sum_{u \in \Gamma(v)} \left(\frac{1}{d(v)+1} - \frac{1}{d(u)+1} \right) .$$

The contribution of each edge $(v, u) \in E$ to the right side of the previous inequality is

$$\begin{aligned} & \frac{1}{d(v)+1} \left(\frac{1}{d(v)+1} - \frac{1}{d(u)+1} \right) \\ & + \frac{1}{d(u)+1} \left(\frac{1}{d(u)+1} - \frac{1}{d(v)+1} \right) \\ & = \left(\frac{1}{d(v)+1} - \frac{1}{d(u)+1} \right)^2 \geq 0 , \end{aligned}$$

and hence, $S(G) > 0$. □

Theorem 1. *Algorithm 1 is a randomized distributed one-round algorithm using single bit messages that finds on each input graph G an independent set of expected size at least $0.224\beta(G)$.*

Proof. Let v be a node in G . Algorithm 1 adds v to the independent set if two independent events happen: v is pre-selected in Line 3 while none of its neighbors are pre-selected. Then, by linearity of expectation,

$$\begin{aligned} \mathbb{E}|I| &= \sum_{v \in V} \mathbb{P}[v \text{ pre-selected}] \cdot \mathbb{P}[v \in I | v \text{ pre-selected}] \\ &= \sum_{v \in V} \frac{1}{d_G(v) + 1} \underbrace{\left(\prod_{u \in \Gamma_G(v)} \left(1 - \frac{1}{d_G(u) + 1} \right) \right)}_K. \end{aligned} \quad (1)$$

We first bound the quantity K from the previous inequality. To this end, fix a vertex $v \in V$. Use the notation $D_v := \sum_{w \in \Gamma_G(v)} \frac{1}{d_G(w) + 1}$. Using that $1 + x \leq e^x$, for any real value x , and $1 - x \geq e^{-1.39x}$, for $0 \leq x \leq 1/2$, we have that for a node v ,

$$\begin{aligned} K &\geq \exp \left(-1.39 \sum_{w \in \Gamma_G(v)} \frac{1}{d_G(w) + 1} \right) \\ &= \frac{1}{e} \exp(1 - 1.39D_v) \geq \frac{1}{e} (2 - 1.39D_v) \\ &= 0.224 + \frac{1.39}{e} (1 - D_v). \end{aligned}$$

Plugging K back into Equality 1 gives

$$\mathbb{P}[v \in I] \geq 0.224 \cdot \frac{1}{d_G(v) + 1} + \frac{1.39}{e} \cdot \frac{1}{d_G(v) + 1} (1 - D_v).$$

Summing up and applying the linearity of expectation,

$$\begin{aligned} \mathbb{E}[|I|] &= \sum_{v \in V} \mathbb{P}[v \in I] \geq 0.224\beta(G) + \frac{1.39}{e} S(G) \\ &\geq 0.224\beta(G), \end{aligned}$$

by Lemma 1. □

Model Requirements. In Algorithm 1, nodes do not require information about global properties, such as the total number of nodes or a polynomial upper bound thereof. They also do not need to know their neighbors, but only an estimate of their degree. In beeping, radio or sensor network models, degree information is not generally provided. For the maximal independent set problem, algorithms with degree information generally outperform algorithms that operate without such information: While for many models that do not provide degree information, algorithms typically use $\Omega(\log^2 n)$ rounds [30, 10, 1, 20, 29, 19] (see also the $\Omega(\log^2 / \log \log n)$ lower bound of [21]), degree information as employed for example in Luby's algorithm [28, 2] allows for $O(\log n)$ rounds. In our one-round setting, degree information is however crucial for obtaining non-trivial approximation guarantees. We prove in the appendix that if degree information is not provided then every one-round randomized algorithm has an approximation ratio of $\Omega(\frac{n}{\log n})$.

The amount of local computation is proportional to the length of the bit representation of the degree, while the amount of space needed (in bits) is constant, besides access to the node's (approximate) degree.

Synchronization is not really needed: Nodes can converge to a solution once they have heard from all transmitting neighbors.

In terms of transmission capabilities, one bit messages are enough, and a broadcast transmission (the same message goes to all neighbors) is sufficient. The algorithm requires that transmitting nodes are able to hear if at least one neighbor is also transmitting. This ability is known as *sender-side collision detection*.

These requirements are matched by the model $B_{cd}L$ (*Broadcasting nodes possess collision detection*) in the spectrum of beeping models of [29]. It has previously been used for maximal independent set computation in [1, 20].

Sender-side collision detection is the strongest requirement that we impose on the underlying model. Interestingly, it can be avoided if we equipped listening nodes with a stronger reception ability that we denote by *full reception*, i.e., the ability to detect whether *all* neighbors of a node transmitted.

Lemma 2. *Suppose nodes have the ability to detect if all neighbors transmitted or not. Then, Algorithm 1 can be implemented without sender-side collision detection.*

Proof. We invert the meaning of transmit and no transmit, i.e., nodes v with $T_v = \text{true}$ remain silent and nodes with $T_v = \text{false}$ transmit a beep signal. The decision rule of the algorithm is now equivalent to checking if all neighbors transmit while the node listens. \square

Full reception is also a strong assumption, and one may wonder whether a weaker requirement would suffice. In Section 4.2, we show however that either sender-side collision detection or full reception is required for computing non-trivial independent sets in one round, even in almost 3-regular unit interval graphs.

3 Poly-logarithmic Approximation On Bounded-independence Graphs

In this section, we show that an independent set of size $\beta(G)$ is a poly-logarithmic approximation of a maximum independent set in PBI-graphs.

We first show that in any graph $G = (V, E)$, for any node $v \in V$, the sum of the inverted degrees of the nodes in the $2 \frac{\log n}{\log \log n}$ -neighborhood of v is $\Omega(1)$ (Lemma 3). In BI-graphs, the size of an independent set in the subgraph induced by such a $2 \frac{\log n}{\log \log n}$ -neighborhood is at most $f(2 \frac{\log n}{\log \log n})$, by definition. Hence, within the $2 \frac{\log n}{\log \log n}$ -neighborhood of any node $v \in V$, the ratio between the size of a maximum independent set and the sum of inverted degrees is $O(f(\frac{\log n}{\log \log n}))$. This argument is then extended to hold for the entire graph (Theorem 2), which implies our main result.

Lemma 3. *Let $G = (V, E)$ be an arbitrary graph with maximum degree Δ , and let $m = \min\{\Delta, 2 \frac{\log n}{\log \log n}\}$. Then:*

$$\sum_{u \in \Gamma_G^m[v]} \frac{1}{d_G(u)} = \Omega(1) .$$

Proof. Let $v \in V$ be a node, and let $d_0 = d_G(v)$. For abbreviation, let $s_j = |\Gamma_G^{(j)}(v)|$ for $j \geq 1$. We set $s_0 = 1$, and we have $s_1 = d_0$. Furthermore, we define $d_i = \frac{1}{s_i} \sum_{u \in \Gamma_G^{(i)}(v)} d_G(u)$ to be the average degree of the nodes in $\Gamma_G^{(i)}(v)$. Then, the inverted degree sum of the nodes in the m -neighborhood can be bounded as follows:

$$\begin{aligned} \sum_{u \in \Gamma_G^m[v]} \frac{1}{d_G(u)} &= \frac{1}{d_0} + \sum_{j=1}^m \sum_{u \in \Gamma_G^{(j)}(v)} \frac{1}{d_G(u)} \\ &\geq \frac{1}{d_0} + \sum_{j=1}^m \sum_{u \in \Gamma_G^{(j)}(v)} \frac{1}{d_j} = \frac{1}{s_1} + \frac{s_1}{d_1} + \sum_{j=2}^m \frac{s_j}{d_j}, \end{aligned} \tag{2}$$

where the first inequality follows from the relationship between the harmonic mean and the arithmetic mean. For $i \geq 2$, consider a node $u \in \Gamma_G^{(i)}(v)$ of degree at least d_i . Then, $\Gamma_G(u) \subseteq \Gamma_G^{(i-1)}(v) \cup (\Gamma_G^{(i)}(v) \setminus \{u\}) \cup \Gamma_G^{(i+1)}(v)$. Hence, $d_G(u) \leq s_{i-1} + s_i - 1 + s_{i+1}$, and since $d_i \leq d_G(u)$, we also have $d_i \leq s_{i-1} + s_i + s_{i+1}$. Similarly, for




Figure 1: Left: Sequence s_i is not strictly increasing. This implies that there exists an index j such that $s_{j-1} > s_{j-2}$ and $s_{j-1} \geq s_j$. Thus, the expression $\frac{s_{j-1}}{s_{j-2}+s_{j-1}+s_j}$ is bounded from below by $1/3$. Right: Sequence s_i is strictly increasing. This implies that the number of indices $j \in J$ such that $\frac{s_j}{2s_j+s_{j+1}} \geq \frac{\log \log n}{\log n}$ is at most $\frac{\log n}{\log \log n}(1+o(1))$. Thus, $\Theta(\frac{\log n}{\log \log n})$ indices outside J are enough such that $\sum_{i \notin J} \frac{s_j}{2s_j+s_{j+1}} = \Omega(1)$.

d_1 we obtain the inequality $d_1 \leq s_1 + s_2$. Using this in Inequality 2, we obtain:

$$\begin{aligned} \sum_{u \in \Gamma_G^m[v]} \frac{1}{d_G(u)} &\geq \frac{1}{s_1} + \frac{s_1}{d_1} + \sum_{j=2}^m \frac{s_j}{d_j} \\ &\geq \frac{1}{s_1} + \frac{s_1}{s_1 + s_2} + \sum_{j=2}^m \frac{s_j}{s_{j-1} + s_j + s_{j+1}}. \end{aligned} \quad (3)$$

In order to bound the right side of Inequality 3, we treat the cases when the sequence s_i is strictly increasing and when it is not strictly increasing separately. Both cases are illustrated and summarized in Figure 1.

Suppose that the sequence $(s_i)_{1 \leq i \leq m}$ is not strictly increasing. Let j be the smallest index so that $s_j \leq s_{j-1}$. If $j = 2$, then the term $\frac{s_1}{s_1+s_2}$ of Inequality 3 can be bounded by $\frac{s_1}{s_1+s_2} \geq \frac{s_1}{s_1+s_1} = 1/2$, and thus, $\sum_{u \in \Gamma_G^m[v]} \frac{1}{d_G(u)} > \frac{1}{2} = \Omega(1)$. Suppose that $j > 2$. Then, since j is the smallest index, we have $s_{j-2} < s_{j-1}$. Therefore, the addend with index $j-1$ of the sum in the right side in Inequality 3 can be bounded as follows:

$$\frac{s_{j-1}}{s_{j-2} + s_{j-1} + s_j} > \frac{s_{j-1}}{3 \cdot s_{j-1}} = 1/3,$$

which implies $\sum_{u \in \Gamma_G^m[v]} \frac{1}{d_G(u)} > \frac{1}{3} = \Omega(1)$.

Assume now that the sequence $(s_i)_i$ is strictly increasing. We bound the right side of Inequality 3 as follows:

$$\begin{aligned} \frac{1}{s_1} + \frac{s_1}{s_1 + s_2} + \sum_{j=2}^m \frac{s_j}{s_{j-1} + s_j + s_{j+1}} \\ \geq \frac{1}{s_1} + \frac{s_1}{s_1 + s_2} + \sum_{j=2}^m \frac{s_j}{2 \cdot s_j + s_{j+1}}. \end{aligned} \quad (4)$$

Let $J \subseteq \{2, \dots, m\}$ be the subset of indices so that for each $j \in J$: $\frac{s_j}{2s_j+s_{j+1}} \leq \frac{\log \log n}{\log n}$. This implies that we have $s_{j+1} \geq s_j \left(\frac{\log n}{\log \log n} - 2 \right)$, for $j \in J$. Since the sequence $(s_i)_i$ is strictly increasing, we can bound the size of the set J as follows:

$$\left(\frac{\log n}{\log \log n} - 2 \right)^{|J|} \leq n,$$

and thus $|J| \leq \frac{\log n}{\log \log n}(1+o(1))$. Since $m = 2 \frac{\log n}{\log \log n}$, there are $\Theta(\frac{\log n}{\log \log n})$ indices i with $i \notin J$ and $\frac{s_i}{2s_i+s_{i+1}} \geq \frac{\log \log n}{\log n}$. Then, the addends in the right side of Inequality 4 that correspond to those indices

$i \notin J$ sum up to a constant bounded away from 0, for sufficiently large values of n . This proves part 1 of the lemma.

We derive now a bound on m in terms of the maximum degree Δ . To this end, we depart from Inequality 4. Notice that the bound on Δ implies $s_j \leq s_{j-1}\Delta$. Thus, for any j , the addend in Inequality 4 that corresponds to j is bounded as follows: $\frac{s_j}{2s_j s_{j-1}} \geq \frac{s_j}{2s_j + \Delta s_j} = \frac{1}{2 + \Delta}$. Since $m = \Theta(\Delta)$, the right side of Inequality 4 sums up to a constant. \square

Theorem 2. *Let $G = (V, E)$ be a PBI-graph with maximum degree Δ and bounding function f . Then:*

$$\alpha(G) = O\left(\beta(G) \cdot f(\min\{\Delta, \frac{\log n}{\log \log n}\})\right).$$

Proof. Let $m = \min\{\Delta, 2\frac{\log n}{\log \log n}\}$. Let S be a maximal $(2m + 1)$ -independent set in G , i.e., a maximal set of vertices of mutual distance at least $2m + 1$. Let I^* denote a maximum independent set in G . Since S is maximal, every vertex in I^* is at a distance at most $2m$ from a vertex in S , and thus $|I^*| \leq |S| \cdot f(2m)$. Since S is $(2m + 1)$ -independent, the m -neighborhoods around nodes in S are disjoint. Thus, using Lemma 3, we have

$$\begin{aligned} \beta(G) &= \sum_{v \in V} \frac{1}{d_G(v) + 1} \geq \sum_{s \in S} \sum_{v \in \Gamma_G^m(s)} \frac{1}{d_G(v) + 1} \\ &= \Omega(|S|). \end{aligned}$$

Thus,

$$\alpha(G) \leq |S| \cdot f(2m) = O(\beta(G) \cdot f(2m)) = O(\beta(G)f(m)),$$

where we used that $f(2m) = O(f(m))$ holds since f is bounded by a polynomial. \square

4 Limitation Results

We present here several results that indicate that our algorithmic result cannot be improved on. We first see in Section 4.1 that one-round algorithms only yield a $\Omega(n)$ -approximation on (general) disc graphs. In Section 4.2, we prove that either sender-side collision detection or full reception is necessary for one round algorithms. Then, we prove in Section 4.3 that our analysis for d -dimensional unit sphere graphs is tight. We see in Section 4.5 that going beyond BI-graphs is hard, in particular, there are claw-free graphs for which the Caro-Wei bound only yields polynomial approximation factors. Finally, we show in Section 4.4 that multiple iterations of our algorithm do not substantially improve its approximation factor.

4.1 Lower Bound for One-round Algorithms on General Graphs

In this subsection, we prove that every possibly randomized distributed one-round algorithm computes an independent set of size at most $n/\omega(G)$ on any regular graph G , where $\omega(G)$ denotes the *clique number*, i.e., the size of a largest clique. This implies that every one-round algorithm has an approximation factor of at least $\frac{\alpha(G)\omega(G)}{n}$. We then show that there is a disc graph with $\omega(G) = \Omega(n)$ and $\alpha(G) = \Omega(n)$, implying that even on disc graphs, no non-trivial approximation is possible in one round.

To this end, let $G = (V, E)$ be a d -regular graph, for an integer d . We assume that each node has a unique label chosen from $\mathcal{U} = \{1, \dots, m\}$, where $m \geq n$. Let \mathcal{L} denote the set of all possible labellings of V . In order to prove our lower bound, we exploit the fact that all nodes in V have the same local *views*, i.e., in one round, all nodes can only learn the d labels and random coin flips of their adjacent nodes. Since all nodes run the same algorithm, in average over all possible labellings \mathcal{L} , the probabilities for all nodes to end up in I are equal. This fact is used in the following theorem:

Theorem 3. *Every possibly randomized one-round distributed algorithm for maximum independent set on a d -regular graph $G = (V, E)$ that outputs a correct solution with probability at least $1 - 1/n$ has an expected approximation factor of at least $\frac{\omega(G)\alpha(G)}{2n}$.*

Proof. Consider a possibly randomized one-round algorithm for maximum independent set. Then, as previously argued, for all $u, v \in V$, we have $\mathbb{P}[u \in I] = \mathbb{P}[v \in I]$, where the probabilities are taken over all possible labellings \mathcal{L} and the random coin flips of the algorithm. Let $p = \mathbb{P}[v \in I]$, for any node v . Then,

$$\mathbb{E}|I| = \sum_{u \in V} \mathbb{P}[u \in I] = np.$$

Let C be a clique of G of size $\omega(G)$. Then, since the error probability of the algorithm is at most $1/n$, we have

$$\begin{aligned} \omega(G)p &= \mathbb{E}|C \cap I| = \sum_i \mathbb{P}[|C \cap I| = i] \\ &\leq \mathbb{P}[|C \cap I| \leq 1] + \mathbb{P}[|C \cap I| \geq 2] \omega(G) \\ &\leq 1 + \frac{1}{n} \cdot n = 2, \end{aligned}$$

and hence, $p \leq \frac{2}{\omega(G)}$. Therefore, $\mathbb{E}|I| = np \leq \frac{2n}{\omega(G)}$. The expected approximation ratio is hence at least $\frac{\omega(G)\alpha(G)}{2n}$. \square

Consider now the following disc graph $G_r = (C_1 \cup C_2 \cup I, E)$, parametrized by an integer $r \geq 2$, as in Figure 2. Set I consists of r non-overlapping unit discs arranged on a line. Set C_1 (C_2) contains $r - 1$ identical discs of large radii (radius $O(n^3)$ suffices) arranged such that they overlap with all vertices of I from above (below, respectively). Graph G_r has $n = 3r - 2$ vertices, is $(2r - 2)$ -regular, and $\alpha(G_r) = \omega(G_r) = r$ holds.




Figure 2: A $(2r - 2)$ -regular disc graph G_r with $3r - 2$ vertices and $\omega(G_r) = \alpha(G_r) = r$.

Plugging G_r into Theorem 3, we obtain the following corollary:

Corollary 1. *Every possibly randomized one-round distributed algorithm for maximum independent set that is correct with probability at least $1 - 1/n$ has an expected approximation factor of at least $\frac{1}{18}n$ on disc graphs.*

4.2 Model Aspects

In this subsection, we show that every randomized one-round algorithm that computes a $o(n)$ -approximation to the independent set problem either uses sender-side collision detection (a transmitting node can detect whether at least one of its neighbors also transmitted) or full reception (receiving nodes can detect whether all of its neighbors transmitted).

Theorem 4. *Every randomized one-round algorithm that is correct with probability $1 - 1/n$ without sender-side collision detection or full reception has an expected approximation ratio of $\Omega(n)$, even on almost 3-regular unit-interval graphs.*

Proof. We consider one-round algorithms that learn nothing when transmitting. When listening, they have collision detection, i.e., can distinguish between 0, 1, or at least 2 neighbors transmitting. The basic requirement is correctness: The set computed must be independent with probability at least $1 - 1/n$.

Given number n , which we assume to be a multiple of 4, consider the following almost-cubic graph $G = (V, E)$, where $V = \{0, 1, \dots, n-1\}$ and $E = \{(i, i+1) : i = 0, 1, 2, \dots, n-2\} \cup \{(i, i+2) : i \bmod 4 \leq 1\}$. All nodes except nodes 0 and $n-1$ are of degree three; nodes 0 and $n-1$ are of degree 2. Graph G can be represented as a proper interval graph. A portion of the graph is illustrated in Figure 3.

We allow the graph to be labeled. In this case, we assume a uniform distribution over the possible labeling from a finite set set of labels $[m]$, where $m \geq n$. Thus, each node receives a unique label uniformly at random from $[m]$.




Figure 3: Graph $G = (\{0, 1, \dots, n-1\}, E)$ used in the lower bound construction of Theorem 4. G consists of $n/4$ repetitions of the boxed pattern. Nodes 0 and $n-1$ are of degree 2; all other nodes are of degree 3. A key property of G is that $\Gamma_G[4j+1] = \Gamma_G[4j+2]$, for every $0 \leq j \leq n/4 - 1$.

When listening, a node can only learn the label of their transmitting neighbor when hearing from a single neighbor. Each node makes two decisions: whether to transmit in the round and whether to join the independent set, where the latter can depend on how many (and which) neighbors it heard from if it was listening. Both of these actions can be probabilistic. We show that the expected number of nodes joining the independent set must be constant, in order for the solution to be correct with probability $1 - 1/n$.

We first treat the case of nodes joining after transmitting. Let p denote the probability that a node transmits, averaged over the labelings of the node. Let q_T denote the probability that a node joins the independent set in the case that it transmits, again averaged over the labelings of the node. Consider the event A_j , for $j = 1, 2, \dots, n/2$, that both i and $i+1$ transmit and join the independent set. None of these independent events can take place for the solution to be correct. The probability of each event A_j is $(pq_T)^2$, so the probability that none occurs is $(1 - (pq_T)^2)^{n/2} \leq e^{-(pq_T)^2 \cdot n/2} \leq 1 - (pq_T)^2 n/4$, using that $1 - x \leq e^{-x} \leq 1 - x/2$, for $0 \leq x \leq 1$. Since the probability of correctness is assumed to be at least $1 - 1/n$, it holds that $1 - 1/n \leq 1 - (pq_T)^2 n/4$ or $1/n \geq (pq_T)^2 n/4$ or $pq_T \leq 2/n$. On the other hand, the probability that a given node joins the independent set after transmitting is pq_T , so the expected number of nodes that join this way is at most $npq_T \leq n \cdot 2/n = 2$.

Consider now the case of nodes joining after listening. There are four cases, depending on how many of their neighbors transmitted. Let q_i denote the probability that a node joins the solution after i neighbor transmitted, averaged over the labels of the node and its transmitting neighbors. Fix $i = 0, 1, 2$. Consider the event $B_{j,i}$, for $j = 0, 1, \dots, n/4 - 1$ and $i = 0, 1, 2$, that nodes $4j+1$ and $4j+2$ both join after hearing from exactly i neighbors (adjacent nodes cannot both hear from three neighbors). Observe that nodes $4j+1$ and $4j+2$ have two common neighbors and no other neighbors. Thus, the probability of $B_{j,i}$ occurring is $\mathbb{P}[B_{j,i}] = p^i(1-p)^{4-i}q_j^2$, and the probability that none of them occurs (for fixed i) is

$$\begin{aligned} \mathbb{P} \left[\bigwedge_j \overline{B_{j,i}} \right] &\leq (1 - p^i(1-p)^{4-i}q_j^2)^{n/4} \\ &\leq 1 - p^i(1-p)^{4-i}q_j^2 \frac{n}{8}, \end{aligned}$$

where we used $1 - x \leq e^{-x} \leq 1 - x/2$, for $0 \leq x \leq 1$, again. Since the solution is correct with probability at

least $1 - 1/n$, it follows that $1/n \geq p^i(1-p)^{4-i}q_i^2 \frac{n}{8}$, or

$$p^{i/2}(1-p)^{2-i/2}q_i \leq \sqrt{8}/n \leq 3/n. \quad (5)$$

On the other hand, the probability that any node different from 0 and $n-1$ joins the independent set after listening and hearing from exactly i nodes is $\binom{3}{i}(1-p)^{4-i}p^i q_i$, so the expected number E_i of nodes that join this way is at most $(n-2)\binom{3}{i}(1-p)^{4-i}p^i q_i$. Applying Inequality 5, we find that expected number of nodes that join this way is

$$E_i \leq (n-2)\binom{3}{i}(1-p)^{4-i}p^i q_i \leq \binom{3}{i}.$$

In total, the expected number of nodes that join the independent set is then bounded by: two nodes that join after transmitting, the two nodes 0 and $n-1$, and $\binom{3}{0} + \binom{3}{1} + \binom{3}{2} = 1 + 3 + 3 = 7$ of the nodes $V \setminus \{0, n-1\}$ that join after listening. \square

The argument can be extended to the intermediate setting where a node can count up to k transmitting neighbors, but cannot distinguish how many beyond k are transmitting, for a fixed value k .

4.3 Lower Bound for d -dimensional Unit Sphere Graphs

In this section, we design d -dimensional unit sphere graphs G such that $\alpha(G) = \Omega\left(\beta(G)f(\min\{\Delta, \frac{\log n}{\log \log n}\})\right)$ holds, rendering the analysis of Theorem 2 tight.

A d -dimensional *unit sphere graph* $G = (S, E)$ is the intersection graph of d -dimensional unit spheres $S = \{s_1, \dots, s_n\}$ (all spheres have the same radius): Each sphere s_i constitutes a vertex in G and two spheres are adjacent iff they intersect. For $d = 1$, a unit sphere graph is a unit interval graph, and for $d = 2$, a unit sphere graph is a unit disc graph.

Let $d > 0$ be some fixed dimension. We will denote our hard instance graph with $H_k = (V_H, E_H)$ where k is a parameter which we define later. We start our construction of H_k with a *grid graph* $G_k = (V_G, E_G)$ that is parametrized by an integer $k \geq 1$. The vertex set of G_k is defined as $V_G = \{v_x \mid x \in \{0, 1, \dots, k-1\}^d\}$. Let v_x, v_y with $x, y \in \{0, \dots, k-1\}^d$ be two vertices of V_G . Then v_x and v_y are adjacent iff $|x - y| = 1$, where $|x| = \sum_{1 \leq i \leq d} |x_i|$.

The hard instance graph H_k is obtained from G_k as follows: For every vertex $v_x \in V_G$, a clique C_x of size $s(|x|)$ is introduced in H_k , where

$$s(i) = d^i k^{di} \log^i n. \quad (6)$$

Suppose that v_x and v_y are adjacent nodes in G_k . Then all nodes of C_x are connected to all nodes of C_y in H_k , or, in other words, $C_x \cup C_y$ also forms a clique in H_k . We say that a node v_x or a clique C_x is in layer i , if its distance from the vertices of clique $C_{(0, \dots, 0)}$ is i , or, in other words, $|x| = i$.

First, notice that H_k is in fact a d -dimensional unit sphere graph. Each vertex $v \in C_x \subseteq V_H$ with $x \in \{0, \dots, k-1\}^d$ corresponds to a sphere centered at position x with radius $1/2$ (for convenience, in this construction we suppose that all spheres have the radius $1/2$ instead of 1). An example is provided in Figure 4.

We first observe that graph H_k contains an independent set of size $(\frac{k}{2})^d$.

Lemma 4. *Consider graph H_k . Then $\alpha(H_k) \geq (\frac{k}{2})^d$.*

Proof. Let $X = \{2i : i \in \mathbb{N} \cup \{0\} \text{ and } 2i \leq k-1\}$ be the even numbers up to $k-1$ including 0, and let I be a set that contains exactly one vertex from each clique C_x with $x \in X^d$. Then I is independent and $|I| = |X^d| = |X|^d \geq (\frac{k}{2})^d$. \square

We prove now that H_k is a PBI-graph with respect to the bounding function $f(r) = (2r+1)^d$.

Lemma 5. *The d -dimensional unit sphere graph H_k is of bounded independence with respect to the bounding function $f(r) = (2r+1)^d$.*




Figure 4: Illustration of the two-dimensional case: On the left, the grid graph G_4 is illustrated. On the right, the hard instance unit disc graph H_4 is shown. H_4 is obtained from G_4 by replacing each node at position (i, j) with a clique of size $s(i + j)$.

Proof. For some $x \in \{0, \dots, k-1\}^d$, the size of a maximum independent set in the k -neighborhood of a node $v \in C_x \subseteq V_H$ is the same as the size of a maximum independent set of the node $v_x \in V_G$ in the corresponding grid graph. Therefore, the r -neighborhood of an arbitrary node $v_x \in V_G$ with $x \in \{0, \dots, k-1\}^d$ is a subset of the nodes with indices $j \in \{x_1 - r, \dots, x_1 + r\} \times \dots \times \{x_d - r, \dots, x_d + r\}$. Thus, $|\{x_1 - r, \dots, x_1 + r\} \times \dots \times \{x_d - r, \dots, x_d + r\}| = (2r + 1)^d$ is a bound on the size of an independent set in the r -neighborhood of v . \square

Next, we identify the correct value for k so that H_k has $O(n)$ vertices, and we show that $\beta(H_k) = O(1)$.

Lemma 6. Consider graph $H_k = (V_H, E_H)$ with $k = C \cdot \frac{\log n}{d^2 \log \log n}$, for a small enough constant C . Then, H_k has $O(n)$ vertices, and $\beta(H_k) = \Theta(1)$.

Proof. Let $V_i := \{v \in C_x : |x| = i\}$ be the set of nodes in layer i , and denote by n_i the number of cliques in layer i . Then, $|V_i| = n_i \cdot s(i)$. First, note that by construction of H_k we have $n_i \leq n_{i+1}d$. This allows us to establish a relation between $|V_i|$ and $|V_{i+1}|$:

$$\begin{aligned} |V_i| &= n_i \cdot s(i) \leq n_{i+1}d \cdot (d^i k^{di} \log^i n) \leq \frac{n_{i+1} \cdot s(i+1)}{k^d \log n} \\ &= \frac{|V_{i+1}|}{k^d \log n}. \end{aligned}$$

Then, since $|V_H| = \sum_{i=0}^{d(k-1)} |V_i|$ and by the previous inequality, we obtain $|V_H| = O(|V_{d(k-1)}|)$. We compute:

$$|V_{d(k-1)}| = s(d(k-1)) = O\left(d^{kd} k^{d^2 k} \log^{kd} n\right) = O(n),$$

where the last equality can be verified using the definition $k = C \frac{\log n}{d^2 \log \log n}$. We thus established $|V_H| = O(n)$.

Next, in order to prove $\beta(H_k) = \Theta(1)$, notice that for $i < (k-1)d$, the degree of every node of V_i is at least $s(i+1)$, the size of a clique in layer $i+1$. Furthermore, the degrees of the nodes of clique $C_{(k-1, \dots, k-1)}$ are at least $s(d(k-1))$. Thus,

$$\begin{aligned} \beta(H_k) &= \sum_{i=0}^{d(k-1)-1} \frac{n_i \cdot s(i)}{s(i+1)} + \frac{n_{d(k-1)} s_{d(k-1)}}{s_{d(k-1)}} \\ &\leq \sum_{i=0}^{d(k-1)-1} \frac{k^d}{dk^d \log n} + 1 = \frac{k-1}{\log n} + 1 = \Theta(1), \end{aligned}$$

where we used the rough estimate $n_i \leq k^d$. \square

Using the previous lemma, we show that the analysis of Theorem 2 is best possible.

Theorem 5. *There is an infinite family of PBI-graphs \mathcal{G} such that for every $G \in \mathcal{G}$ with bounding function f :*

$$\alpha(G) = \Omega\left(\beta(G) \cdot f\left(\min\left\{\Delta, \frac{\log n}{\log \log n}\right\}\right)\right).$$

Proof. Let n be an arbitrary large integer and let $d \geq 1$ be a constant integer. Furthermore, let $k = C \cdot \frac{\log n}{\log \log(n)^{d^2}}$, for a small enough constant C , as in Lemma 6. By Lemma 6, graph G_k has $O(n)$ vertices and $\beta(G_k) = \Theta(1)$ holds. Furthermore, by Lemma 5, $f(r) = (2r + 1)^d$ is a bounding function of G_k . By Lemma 4, G_k contains an independent set of size $\Omega\left(\left(\frac{k}{2}\right)^d\right)$, and thus:

$$\begin{aligned} \beta(G_k) \cdot f\left(\frac{\log n}{\log \log n}\right) &= O\left(\left(\frac{\log n}{\log \log n}\right)^d\right), \text{ and} \\ \alpha(G_k) &= \Omega\left(\left(\frac{k}{2}\right)^d\right) = \Omega\left(\left(\frac{\log n}{2d^2 \log \log n}\right)^d\right). \end{aligned}$$

The theorem thus holds for any constant d . □

Since the expected size of an independent set computed by Algorithm 1 is $\Theta(\beta(G))$ (and $\beta(G)$ if computed by ALON-SPENCER-IS), we obtain the following corollary, which shows that the analysis of Theorem 1 is tight.

Corollary 2. *There is an infinite family of PBI-graphs \mathcal{G} such that for every $G \in \mathcal{G}$ with bounding function f , the expected approximation ratios of Algorithm 1 and ALON-SPENCER-IS are $\Omega\left(f\left(\min\left\{\Delta, \frac{\log n}{\log \log n}\right\}\right)\right)$ on input G .*

4.4 Lower Bound for Multiple Iterations

Instead of running Algorithm 1 or ALON-SPENCER-IS once, one may wonder whether applying these algorithms repeatedly leads to an improved approximation ratio. In this section, we show that this is not the case: We show that running constant number of iterations improves the approximation factor at most by a constant factor.

We consider the algorithm as depicted in Algorithm 2 (ONE-ROUND-IS denotes either Algorithm 1 or ALON-SPENCER-IS):

Algorithm 2 Multiple iterations independent set alg.

Require: $G = (V, E)$ {Input graph}, r {number of rounds}

- 1: $V' \leftarrow V$ {active nodes}, $I \leftarrow \emptyset$ {the independent set to be computed}
 - 2: **for** $i = 1 \dots r$ **do**
 - 3: $I \leftarrow I \cup \text{ONE-ROUND-IS}(G[V'])$
 - 4: $V' \leftarrow V' \setminus (I \cup \Gamma_G(I))$
 - 5: **end for**
 - 6: **return** I
-

We will show that in r iterations, Algorithm 2 computes an independent set of size at most r^d with probability $(1 - O(\frac{1}{d \log n}))^i$ on graph H_k , which was defined in the previous section. Since H_k contains an independent set of size $\Omega\left(\left(\frac{k}{2}\right)^d\right)$ (Lemma 4), this proves that the approximation ratio of Algorithm 2 is $\Omega\left(\left(\frac{k}{2r}\right)^d\right)$.

Consider the situation of Algorithm 2 at the end of the i th iteration of the for-loop. We will prove that at this moment, all cliques of layers at most $(k - 1)^d - 2i$ are contained in V' with high probability. This proves that with high probability, in each iteration the algorithm selects vertices of the outermost layer, which then leads to the removal of a subset of the vertices of the two outermost layers.

In the following, denote by V'_i the set V' of Algorithm 2 after the i th iteration of the for-loop, and denote by I_i the set I after the i th iteration. The index $i = 0$ describes the situation before the first execution of the for-loop, i.e., $V'_0 = V$ and $I_0 = \emptyset$.

Lemma 7. *In Algorithm 2, all cliques of layers $\leq (k-1)^d - 2i$ are contained in V'_i , with probability $(1 - O(\frac{1}{d \log n}))^i$.*

Proof. The proof is by induction on i . As a base case, it is easily verified that V'_0 and I_0 fulfill the statement of the lemma. Suppose now the lemma holds for iteration i . We will prove that it still holds after iteration $i+1$.

Denote by E the event that all cliques of layers at most $(k-1)^d - 2i$ are included in V'_i , and assume that E holds. Let C_x be a clique with $|x| \leq (k-1)^d - 2i - 1$ (recall that $|x| = \sum_{1 \leq i \leq d} |x_i|$). Then, the probability that a node v of C_x is chosen into the independent set by either Algorithm 2 or ALON-SPENCER-IS in iteration $i+1$ is at most $\frac{1}{s(|x|+1)}$, since, conditioned on E , there is at least one clique $C_y \in V'_i$ of layer $|x|+1$ incident to v . Thus, by the union bound, the probability that at least one node of C_x is chosen is at most

$$\frac{|C_x|}{s(|x|+1)} = \frac{s(|x|)}{s(|x|+1)} = O\left(\frac{1}{dk^d \log n}\right).$$

Applying the union bound again, the probability that a node of any of the cliques of layers at most $(k-1)^d - 2i - 1$ is chosen in round $i+1$ is $O(\frac{1}{d \log n})$, since the total number of cliques in H_k is k^d .

This implies that, conditioned on E , with probability $O(\frac{1}{d \log n})$, all cliques of layers at most $(k-1)^d - 2i - 2$ are included in V'_{i+1} . Note that even though we proved that, with high probability, nodes of layer $(k-1)^d - 2i - 1$ are not chosen by the algorithm, some of their neighbors of layer $(k-1)^d - 2i$ might be chosen, which would lead to their removal in Line 4 of the algorithm.

By the induction hypothesis, $\mathbb{P}[E] \geq (1 - O(\frac{1}{d \log n}))^i$, and, therefore, with probability at least $(1 - O(\frac{1}{d \log n}))^{i+1}$, all cliques C_x of layers at most $(k-1)^d - 2i - 2 = (k-1)^d - 2(i-1)$ are included in V'_{i+1} , which completes the lemma. \square

Theorem 6. *There is an infinite family of PBI-graphs \mathcal{G} such that for every $G \in \mathcal{G}$ with bounding function f , running r iterations of Algorithm 2 obtains an*

$$\Omega\left(f\left(\frac{1}{r} \cdot \min\left\{\Delta, \frac{\log n}{\log \log n}\right\}\right)\right)$$

approximation on input G with probability $(1 - O(\frac{1}{d \log n}))^i$.

Proof. Let n be an arbitrary large integer and let $d \geq 1$ be a constant integer. Furthermore, let $k = C \cdot \frac{\log n}{\log \log(n)^{d^2}}$, for a small enough constant C , as in Lemma 6. By Lemma 7, all cliques C_x of layers at most $(k-1)^d - 2r$ are included in V'_r with probability $\Omega((1 - O(\frac{1}{d \log n}))^i)$. Hence, among the eliminated nodes, every independent set is of size $O(r^d)$. Since H_k contains an independent set of size $\Omega((\frac{\log n}{\log \log n})^d)$ (Lemma 4), the result follows. \square

4.5 The Caro-Wei Bound in Claw-Free-Graphs

Every PBI-graph is $O(f(1)+1) = O(1)$ -claw free. Claw-free graphs are a natural superclass of PBI graphs, and one may wonder how well the Caro-Wei bound behaves on this graph class. We find that $\beta(G)$ is a $\Theta(\sqrt{bn})$ -approximation to $\alpha(G)$ on b -claw-free graphs.

For integer $b \geq 1$ with $b = O(\sqrt{n})$, we construct a $(b+2)$ -claw-free graph G on $O(n)$ vertices that contains an independent set of size $\Omega(\sqrt{nb})$ while $\beta(G) = O(1)$. This construction shows that in $(b+2)$ -claw-free graphs, $\beta(G)$ may only be a polynomial approximation of the size of a maximum independent set.

To this end, let G' be the graph that consists of k copies of $K_{a,b}$, the complete bipartite graph with bipartitions of sizes a and b . Then, G is obtained from G' by joining all left sides of the copies of $K_{a,b}$

(the bipartitions of size a) into a clique. In other words, $G = (U \cup V, E)$ is a split graph with a clique $U = \{u_1, \dots, u_{ka}\}$ and independent set $V = \{v_1, \dots, v_{kb}\}$, where $u_i v_j \in E \Leftrightarrow \lfloor \frac{i}{a} \rfloor = \lfloor \frac{j}{b} \rfloor$. Notice that G is $(b+2)$ -claw-free.

Theorem 7. *Let n, b be integers with $b = O(\sqrt{n})$. Then there exists a $(b+2)$ -claw-free graph G on $O(n)$ vertices such that $\alpha(G) = \Omega(\sqrt{nb})\beta(G)$.*

Proof. Consider the graph G as defined above with $a = \Theta(\sqrt{nb})$ and $k = \Theta(\sqrt{n}/\sqrt{b})$. Observe that U has ka nodes in a clique, each adjacent to b nodes in V , while nodes in V are independent and of degree a . Then,

$$\begin{aligned}\alpha(G) &= kb = \Theta(\sqrt{nb}), \\ \beta(G) &= \frac{ka}{ka-1+b} + \frac{kb}{a+1} \leq 1 + \frac{\Theta(\sqrt{nb})}{\Theta(\sqrt{nb})+1} = O(1),\end{aligned}$$

and thus $\frac{\alpha(G)}{\beta(G)} = \Omega(\sqrt{nb})$. Graph G has $k(a+b) = O(n + \sqrt{nb}) = O(n)$ vertices, since $b = O(\sqrt{n})$. \square

We now prove a matching upper bound.

Theorem 8. *Let $b \geq 2$ be an integer, and let G be a $(b+1)$ -claw-free graph. Then, $\alpha(G) \leq \sqrt{nb} \cdot \beta(G)$.*

Proof. Let I^* denote an independent set of size $\alpha(G)$. Let $d = \frac{1}{\alpha(G)} \sum_{v \in I^*} d_G(v)$ be the average degree of the nodes in I^* . Notice that since G is $(b+1)$ -claw-free, each vertex outside I^* is incident to at most b nodes of I^* . Counting edges incident on I^* , using that I^* is independent, we get that $\sum_{v \in I^*} d_G(v) \leq (n - \alpha(G))b$. That implies that $d \leq \frac{(n - \alpha(G))b}{\alpha(G)}$. Then, since $\beta(G) \geq \sum_{v \in I^*} \frac{1}{d_G(v)+1} \geq \frac{|I^*|}{d+1}$ (using the relationship between the harmonic and arithmetic mean),

$$\beta(G) \geq \frac{\alpha(G)}{d+1} \geq \frac{\alpha(G)^2}{(n - \alpha(G))b + \alpha(G)} > \frac{\alpha(G)^2}{nb},$$

since $\alpha(G)(1-b) < 0$. Since $\beta(G) \geq 1$, we obtain that $\frac{\alpha(G)}{\beta(G)} \leq \min\{\alpha(G), \frac{nb}{\alpha(G)}\} \leq \sqrt{nb}$. \square

5 Conclusion

In this paper, we gave a one-round, single-bit messages randomized algorithm, which computes an independent set of expected size $\Theta(\beta(G))$, where $\beta(G)$ is the Caro-Wei bound. We proved that the Caro-Wei bound approximates the size of a maximum independent set in polynomially bounded-independence graphs within a poly-logarithmic factor, which implies that the approximation factor of our algorithm is poly-logarithmic for this graph class. We complemented our results by showing that no one-round algorithm can achieve an $o(n)$ -approximation factor on general graphs.

A natural question to examine in the future is whether using larger (but still constant) number of rounds gives markedly better results, either by extending the graph class or resulting in asymptotically better approximations. A pertinent question might then be whether there is a property of a larger neighborhood that materially improves on just knowing the degree.

References

- [1] Yehuda Afek, Noga Alon, Ziv Bar-Joseph, Alejandro Cornejo, Bernhard Haeupler, and Fabian Kuhn. Beeping a maximal independent set. *Distributed computing*, 26(4):195–208, 2013.
- [2] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986.

- [3] Noga Alon and Joel H Spencer. *The probabilistic method*. John Wiley & Sons, 2004.
- [4] Dana Angluin. Local and global properties in networks of processors (extended abstract). In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing, April 28-30, 1980, Los Angeles, California, USA*, pages 82–93, 1980.
- [5] Leonid Barenboim. On the locality of some NP-complete problems. In *Proc. of International Colloquium on Automata, Languages, and Programming, ICALP'12 (2)*, pages 403–415, 2012.
- [6] Leonid Barenboim, Michael Elkin, and Cyril Gavoille. A fast network-decomposition algorithm and its applications to constant-time distributed computation - (extended abstract). In *Structural Information and Communication Complexity - 22nd International Colloquium, SIROCCO 2015, Montserrat, Spain, July 14-16, 2015, Post-Proceedings*, pages 209–223, 2015.
- [7] Marijke H.L. Bodlaender, Magnús M. Halldórsson, Christian Konrad, and Fabian Kuhn. Brief announcement: Local independent set approximation. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC '16, New York, NY, USA, 2016. ACM.
- [8] Y. Caro. New results on the independence number. Technical report, Tel Aviv University, 1979.
- [9] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the local model. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2016)*, 2016.
- [10] Alejandro Cornejo and Fabian Kuhn. Deploying wireless networks with beeps. In *Proceedings of the 24th International Conference on Distributed Computing*, DISC'10, pages 148–162, Berlin, Heidelberg, 2010. Springer-Verlag.
- [11] Andrzej Czygrinow, Michal Hańćkowiak, and Wojciech Wawrzyniak. Fast distributed approximations in planar graphs. In *Proceedings of the 22nd International Symposium on Distributed Computing*, DISC '08, pages 78–92, Berlin, Heidelberg, 2008. Springer-Verlag.
- [12] Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 270–277, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics.
- [13] Mark Goldberg and Thomas H. Spencer. An efficient parallel algorithm that finds independent sets of guaranteed size. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 219–225, January 1990.
- [14] Bjarni V. Halldórsson, Magnús M. Halldórsson, Elena Losievskaja, and Mario Szegedy. Streaming algorithms for independent sets in sparse hypergraphs. *Algorithmica*, 76:490–501, 2016.
- [15] Magnús M. Halldórsson. Wireless scheduling with power control. *ACM Trans. Algorithms*, 9(1):7:1–7:20, December 2012.
- [16] Magnús M. Halldórsson and Christian Konrad. Distributed algorithms for coloring interval graphs. In *Proceedings of the 28th International Conference on Distributed Computing*, DISC'14, 2014.
- [17] Magnús M. Halldórsson and Pradipta Mitra. Nearly optimal bounds for distributed wireless scheduling in the sinr model. *Distrib. Comput.*, 29(2):77–88, April 2016.
- [18] Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105–142, 1999.
- [19] Stephan Holzer and Nancy Lynch. Brief announcement: Beeping a maximal independent set fast. In *30th International Symposium on Distributed Computing (DISC)*, 2016.

- [20] Peter Jeavons, Alex Scott, and Lei Xu. Feedback from nature: simple randomised distributed algorithms for maximal independent set selection and greedy colouring. *Distributed Computing*, 29(5):377–393, 2016.
- [21] Tomasz Jurdzinski and Grzegorz Stachowiak. Probabilistic algorithms for the wakeup problem in single-hop radio networks. In *Proceedings of the 13th International Symposium on Algorithms and Computation*, ISAAC '02, pages 535–549, London, UK, UK, 2002. Springer-Verlag.
- [22] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [23] Thomas Kesselheim and Berthold Vöcking. Distributed contention resolution in wireless networks. In *Proceedings of the 24th International Conference on Distributed Computing*, DISC'10, pages 163–178, Berlin, Heidelberg, 2010. Springer-Verlag.
- [24] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. *J. ACM*, 63(2):17:1–17:44, March 2016.
- [25] Fabian Kuhn and Roger Wattenhofer. On the complexity of distributed graph coloring. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Principles of Distributed Computing*, PODC '06, pages 7–15, New York, NY, USA, 2006. ACM.
- [26] Fabian Kuhn, Roger Wattenhofer, and Aaron Zollinger. Ad-hoc networks beyond unit disk graphs. In *Proceedings of the 2003 Joint Workshop on Foundations of Mobile Computing*, DIALM-POMC '03, pages 69–78, New York, NY, USA, 2003. ACM.
- [27] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [28] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.
- [29] Yves Métivier, John Michael Robson, and Akka Zemmari. On distributed computing with beeps. *CoRR*, abs/1507.02721, 2015.
- [30] Thomas Moscibroda and Roger Wattenhofer. Maximal independent sets in radio networks. In *Proceedings of the Twenty-fourth Annual ACM Symposium on Principles of Distributed Computing*, PODC '05, pages 148–157, New York, NY, USA, 2005. ACM.
- [31] Moni Naor and Larry Stockmeyer. What can be computed locally? *SIAM J. Comput.*, 24(6):1259–1277, December 1995.
- [32] Alessandro Panconesi and Aravind Srinivasan. On the complexity of distributed network decomposition. *J. Algorithms*, 20(2):356–374, March 1996.
- [33] Stefan Schmid and Roger Wattenhofer. Algorithmic models for sensor networks. In *Proceedings of the 20th International Conference on Parallel and Distributed Processing*, IPDPS'06, pages 177–177, Washington, DC, USA, 2006. IEEE Computer Society.
- [34] Johannes Schneider and Roger Wattenhofer. An Optimal Maximal Independent Set Algorithm for Bounded-Independence Graphs. *Distributed Computing*, 22, March 2010.
- [35] Jukka Suomela. Survey of local algorithms. *ACM Comput. Surv.*, 45(2):24:1–24:40, March 2013.
- [36] P. Turán. On an extremal problem in graph theory (in Hungarian). *Mat. Fiz. Lapok*, 48:436–452, 1941.
- [37] V.K. Wei. A lower bound on the stability number of a simple graph. Technical report, Bell Laboratories, 1981.

[38] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007.

A Impossibility Result for One-round Algorithms Without Degree Information

We show in this section that degree information is required in order to obtain non-trivial approximation guarantees in one round.

We assume the following setting. Let n be an integer. We will consider labelled graphs on n vertices, where the labels of the nodes are chosen from the set $\mathcal{L} = [4n + 3]$. Initially, nodes only know their labels and the cardinality n of the graph; in particular, the nodes do not know their degrees. Nodes then either transmit a signal or remain silent. Each node v subsequently receives a bit b_v indicating whether at least one of v 's neighbors transmitted; this is full-duplex, i.e., independent of whether v transmitted or not. Last, based on this information each node decides whether to join the output independent set. The output must be correct with probability 1. Let \mathcal{A} be a randomized distributed one-round maximum independent set algorithm operating in this way.

We show that, for each \mathcal{A} , there is an input graph labeled with labels of \mathcal{L} for which the approximation factor of \mathcal{A} is $\Omega(\frac{n}{\log n})$. Specifically, we use the clique graph K_n when the algorithm is aggressive, opting to transmit with relatively high probability, and use the path P_n , when transmissions are unlikely.

Theorem 9. *The approximation factor of \mathcal{A} is $\Omega(\frac{n}{\log n})$.*

Proof. For a label $l \in \mathcal{L}$ let $p(l)$ denote the probability that a node with label l transmits. Furthermore, for $t, b \in \{0, 1\}$ let $p_{tb}(l)$ be the probability that a node with label l joins the independent set conditioned on transmitting/not transmitting (indicated by t) and receiving/not receiving (indicated by b).

First, we observe that if the algorithm is to achieve any approximation at all, the algorithm can decide to transmit deterministically on only few labels. Specifically, for at most $2n$ labels l does it hold that $p(l) \in \{0, 1\}$. Namely, if all the nodes are deterministic in the same way, then no information is passed between them. Thus, they must make their choice of joining the solution depending only on their label. Since adjacent nodes cannot both join the solution with positive probability, there can be at most one label that allows the node to join. Thus, if there are $n + 1$ labels for which $p(l) = 0$ (or symmetrically, $p(l) = 1$), then for n of these labels the nodes will never join the solution; the output of the algorithm is then zero, for an infinite approximation ratio.

We focus therefore on the set \mathcal{L}' of at least $2n + 3$ labels for which transmission is not deterministic, i.e., labels l such that $p(l) \in (0, 1)$. We argue that \mathcal{L}' contains at most one label l_1 with $p_{01}(l_1) > 0$, at most one label l_2 with $p_{00}(l_2) > 0$, and at most one label l_3 with $p_{11}(l_3) > 0$. We only give the argument for l_1 , the arguments for l_2 and l_3 are similar: Suppose that there were two labels $l_1, l'_1 \in \mathcal{L}'$ with $p_{01}(l_1) > 0$ and $p_{01}(l'_1) > 0$. Then, suppose there are nodes u, v with those labels that form a clique together with a third node w with a label of \mathcal{L}' . Since there is a non-zero probability that w transmits while neither u or v transmit, both u and v may join the solution with positive probability, which contradicts the correctness of the algorithm.

Thus, for all $l \in \mathcal{L}'' = \mathcal{L}' \setminus \{l_1, l_2, l_3\}$, we have $p_{00}(l) = p_{01}(l) = p_{11}(l) = 0$. The only safe configuration for a node with a label from \mathcal{L}'' to join the independent set is then when it transmits and does not receive a message from its neighbors. If there are n labels $l \in \mathcal{L}''$ with $p(l) \geq 24 \log(n)/n$, then consider the n -clique graph with nodes assigned such labels. The random variable X counting the number of transmissions has expectation $\mathbb{E}X \geq 24 \log(n)$ and by a Chernoff bound,

$$\mathbb{P}[X \leq \mathbb{E}X/2] \leq \exp\left(-\frac{1}{8} \cdot 24 \log(n)\right) \leq n^{-2}.$$

Thus, with high probability, more than one node of G transmits, and since only transmitting nodes whose neighbors do not transmit can join the independent set, none of these nodes join the independent set. The approximation factor of the algorithm is thus at least $\frac{1}{n^{-2}} = n^2$.

On the other hand, if there are n labels $l \in \mathcal{L}''$ with $p(l) \leq 24 \log(n)/n$, then consider the path graph P_n with nodes assigned such labels. The number X of transmitting nodes has expectation $\mathbb{E}X \leq 24 \log(n)$ and by a Chernoff bound,

$$\mathbb{P}[X \geq 2\mathbb{E}X] \leq \exp(-24 \log n/3) \leq n^{-2} .$$

Thus, with high probability, at most $48 \log n$ nodes transmit, and by the discussion above, the independent set computed is thus of size at most $48 \log n$. The approximation factor of the algorithm in this case is $\Omega(\frac{n}{\log n})$.

In all cases, the approximation factor is $\Omega(\frac{n}{\log n})$, giving the result. □