Brief Announcement: Faster CONGEST Approximation Algorithms for Maximum Weighted Independent Set in Sparse Graphs

Salwa Faour
University of Freiburg, Germany

Fabian Kuhn
University of Freiburg, Germany

Abstract

The maximum independent set problem is a classic optimization problem in graph theory that has also been studied quite intensively in the distributed setting. Although the problem is hard to approximate within reasonable factors in general, there are good approximation algorithms known for several sparse graph families. In the present paper, we consider deterministic distributed CONGEST algorithms for the weighted version of the problem in trees and graphs of bounded arboricity (i.e., hereditary sparse graphs).

For trees, we prove that the task of deterministically computing a $(1-\varepsilon)$-approximate solution to the maximum weight independent set (MWIS) problem has a tight $\Theta(\log^*(n)/\varepsilon)$ complexity. The lower bound already holds on unweighted oriented paths. On the upper bound side, we show that the bound can be achieved even in unrooted trees.

For graphs $G=(V,E)$ of arboricity $\beta>1$, we give two algorithms. If the sum of all node weights is $w(V)$, we show that for any $\varepsilon>0$, an independent set of weight at least $(1-\varepsilon)\cdot\frac{w(V)}{4\beta}$ can be computed in $O(\log^2(\beta/\varepsilon)/\varepsilon+\log^* n)$ rounds. This result is obtained by a direct application of the local rounding framework of Faour, Ghaffari, Grunau, Kuhn, and Rozhoň [SODA '23]. We further show that for any $\varepsilon>0$, an independent set of weight at least $(1-\varepsilon)\cdot\frac{w(V)}{2\beta+1}$ can be computed in $O(\log^3(\beta)\cdot\log(1/\varepsilon)/\varepsilon^2\cdot\log n)$ rounds. For $\varepsilon=\omega(1/\sqrt{\beta})$, this significantly improves on a recent result of Gil [OPODIS '23], who showed that a $1/\lfloor(2+\varepsilon)\beta\rfloor$-approximation to the MWIS problem can be computed in $O(\beta/\varepsilon\cdot\log n)$ rounds. As an intermediate step to our result, we design an algorithm to compute an independent set of total weight at least $(1-\varepsilon)\cdot\sum_{v\in V}\frac{w(v)}{\deg(v)+1}$ in time $O(\log^3(\Delta)\cdot\log(1/\varepsilon)/\varepsilon+\log^* n)$, where Δ is the maximum degree of the graph.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases CONGEST model, weighted independent set, approximation, trees, arboricity

Digital Object Identifier 10.4230/LIPIcs.DISC.2025.54

Related Version *Full Version*: `https://arxiv.org/abs/2506.10845`

1 Introduction and Related Work

Given a graph $G=(V,E)$ with node weights $w: V\to\mathbb{R}_{\geq 0}$, the *maximum weight independent set (MWIS)* problem asks for an independent set $I\subseteq V$ (i.e., a set I of pairwise non-adjacent nodes) such that the total weight $w(I):=\sum_{v\in V}w(v)$ is maximized. The MWIS problem is a classic optimization problem on graphs, which has also been studied quite extensively in the distributed setting. In the present paper, we focus on sparse families of graphs for which good MWIS approximations are known in the centralized and also in the distributed setting. We start by briefly summarizing the most relevant existing literature.

In the standard setting, the graph $G=(V,E)$ on which we intend to solve some given graph problem (e.g., MWIS) is also the communication graph. There are $n=|V|$ nodes and each node is equipped with an $O(\log n)$-bit unique identifier. The nodes V communicate over the edges in synchronous rounds. In the LOCAL model [18], the nodes can exchange

39th International Symposium on Distributed Computing (DISC 2025).
Editor: Dariusz R. Kowalski; Article No. 54; pp. 54:1–54:7
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

arbitrarily large messages. In the more restricted CONGEST model, in each round, each node can send a (possibly different) message of $O(\log n)$ bits to each neighbor. When considering CONGEST algorithms for the MWIS problem, we further assume that a single node weight can be communicated with a single message. The internal computation at the nodes is not restricted. Initially, the nodes do not know anything about the topology of the graph G. For simplicity, we however assume that the nodes do know the values of the relevant parameters of G. At the end of an independent set algorithm, each node must know if it is in the independent set or not.

1.1 Distributed MWIS Algorithms in General Graphs

The most widely studied distributed independent set problem is the problem of computing a maximal independent set (MIS), i.e., an independent set that cannot be extended. The distributed complexity of computing an MIS has been studied intensively since the 1980s, e.g., [1, 5, 10, 9, 17, 11, 15]. The current best randomized algorithms have complexity $O(\log\Delta)+\tilde{O}(\log^{5/3}\log n)$ in LOCAL and $O(\log\Delta)+\tilde{O}(\log^{3}\log n)$ in CONGEST [10, 9, 11].[1] The best known deterministic MIS algorithms require $\tilde{O}(\log^{5/3} n)$ rounds in LOCAL and $\tilde{O}(\log^2\Delta\cdot\log n)$ rounds in CONGEST [9, 11]. For the *maximum cardinality independent set (MCIS)* problem, i.e., for the unweighted version of the MWIS problem, an MIS directly gives a $1/\Delta$-approximation (where Δ denotes the maximum degree of the graph). In [3], it is shown that in CONGEST, a $1/\Delta$-approximation for MWIS can be computed in time $O(\log W\cdot T_{\text{MIS}})$, where W denotes the ratio between the largest and smallest node weight and T_{MIS} is the time to compute an MIS. Subsequently, Kawarabayashi, Khoury, Schild, and Schwartzman [14] give an algorithm to compute an independent set of weight at least $(1-\varepsilon)\cdot\frac{w(V)}{\Delta+1}$ in time $O(1/\varepsilon)$ times the time to compute an independent set of weight at least $\frac{w(V)}{c(\Delta+1)}$ for some constant $c>0$. For the latter problem, they give two algorithms. One of them uses the local-ratio technique [2] to reduce the problem to computing a single MIS of a (locally computable) subgraph of the input graph. The other one is a randomized algorithm that requires $\text{poly}(\log\log n)$ rounds. In [9], it is shown that it suffices to run $O(\log(1/\varepsilon))$ instances of computing an independent set of weight $\frac{w(V)}{c(\Delta+1)}$. Further, the paper gives a deterministic $O(\log^2\Delta+\log^* n)$-time CONGEST algorithm to compute an independent set of total weight at least $\frac{w(V)}{4(\Delta+1)}$.

1.2 Distributed MWIS Algorithms for Sparse Graphs

We now get to the case of sparse graphs, which are the focus of this paper. It has been shown by Turán [19] that every graph of average degree $\overline{d}$ has an independent set of size at least $n/(\overline{d}+1)$. A natural generalization of this is the so-called Caro-Wei bound [7, 20], which states that every graph G has an independent set of size

$$\alpha(G)\geq\mathsf{CaroWei}(G):=\sum_{v\in V}\frac{1}{\deg(v)+1}\geq\frac{n}{\overline{d}+1}. \tag{1}$$

The second inequality follows from an application of the Cauchy-Schwarz inequality. As observed by Boppana, the Caro-Wei bound can be obtained in expectation by a simple random process that can be implemented in a single round in the CONGEST model [6, 13]. Every node picks an independent random number from a sufficiently large domain and a

[1] The notation $\tilde{O}(\cdot)$ hides polylogarithmic factors in the argument, i.e., $\tilde{O}(x)=x\cdot\text{poly}\log x$.

node joins the independent set if and only if it picked a smaller number than all its neighbors. This algorithm is equivalent to running a single phase of Luby's classic MIS algorithm [17]. In [6], it is shown that the Caro-Wei bound (1) is a $\frac{4\cdot(\sqrt{2}-1)}{\overline{d}+2} \approx \frac{1.657}{\overline{d}+2}$-approximation of the MCIS problem. For sparse graphs, we often want to express the approximation factor as a function of the arboricity β of the graph. The arboricity of a graph is defined as the number of forests into which the edges of a graph can be partitioned. Since we always have $\beta \geq 2\overline{d}$, the result of [6] implies that (1) is a $\frac{2\cdot(\sqrt{2}-1)}{\beta+1} \approx \frac{0.828}{\beta+1}$-approximation of the MCIS problem. In [9], it is shown that there is a deterministic CONGEST algorithm with round complexity $O\big(\frac{\log^2(\Delta/\varepsilon)}{\varepsilon} + \log^* n\big)$ that comes within a factor $(1/2-\varepsilon)$ of the bound of (1), as well as of a natural generalization of (1) in the weighted case. As one of the technical results of the present paper, we show that one can improve the $(1/2-\varepsilon)$ factor to a $(1-\varepsilon)$ factor at the cost of an additional $\log(\Delta)/\varepsilon$ factor in the round complexity (cf. Theorem 4).

The first paper that achieves an efficient $1/O(\beta)$-approximation for the MWIS problem in the CONGEST model is [14]. They achieve a $(1-\varepsilon)/(8\beta)$-approximation by first decomposing the graph into $O(\log n)$ layers of degree at most $4\beta - 1$ and by running an algorithm to obtain an independent set of weight $(1-\varepsilon)\cdot w_i(V_i)/(4\beta)$ within each layer V_i (and where w_i is an appropriate weight function that is used in layer V_i). The overall round complexity is $O(\log n)$ times the time for computing the independent set of weight $(1-\varepsilon)\cdot w_i(V_i)/(4\beta)$ in each layer. By using the randomized algorithm of [14], the overall complexity of the resulting randomized $(1-\varepsilon)/(8\beta)$-approximation for the MWIS problem is $O\big(\frac{\log n \cdot \text{poly}\log\log n}{\varepsilon}\big)$ and by using the deterministic algorithm of [9], the overall complexity of the resulting deterministic $(1-\varepsilon)/(8\beta)$-approximation for the MWIS problem is $O\big((\log^2 \beta \cdot \log(1/\varepsilon)\cdot \log n\big)$. In fact, at the cost of an additional $1/\varepsilon$-factor in the round complexity, the approximation quality of the algorithm of [14] can be improved to $(1-\varepsilon)/(4\beta)$. Those results were then improved by Gil in [12], who in particular provides a deterministic CONGEST algorithm to compute a $1/\lfloor(2+\varepsilon)\beta\rfloor$-approximation in time $O(\beta\cdot\log(n)/\varepsilon)$.

More special classes of sparse graphs are trees and more generally minor-closes families of graphs (e.g., planar graphs). In [16, 8], it was shown that on cycles (and paths), deterministically computing a constant approximation for the maximum cardinality independent set problem requires $\Omega(\log^* n)$ rounds. Note that in all graphs of arboricity $\beta = O(1)$, computing a constant MCIS approximation is trivial. Half the nodes have degree at most 4β and an independent set of size $\Omega(n)$ among the nodes of degree at most 4β can be computed by computing an MIS in time $O(\beta + \log^* n)$ [5]. For planar graphs (and thus also for trees), it is shown in [8] that in the LOCAL model, even for the MWIS problem, a $(1-\varepsilon)$-approximation can be computed in time $\text{poly}(\varepsilon^{-1})\cdot O(\log^* n)$.

1.3 Our Contributions

In our paper, we focus on deterministic approximation algorithms for the MWIS problem in the CONGEST model in sparse families of graphs. In the following, we list our technical contributions in detail and we also give an overview over the most important ideas that are needed to prove the stated theorems.

1.3.1 Approximating Maximum Weight Independent Set in Trees

We start by establishing a tight bound for computing a $(1-\varepsilon)$-approximation in trees.

▶ **Theorem 1.** *Let $\varepsilon > \log(n)/n$ be a parameter. The deterministic* CONGEST *model complexity of computing a $(1-\varepsilon)$-approximate solution for MWIS in tree networks is $\Theta\big(\frac{\log^* n}{\varepsilon}\big)$. The upper bound holds in general unrooted trees. The lower bound even holds for the unweighted version of the problem in oriented paths and in the* LOCAL *model.*

The lower bound is based on a relatively simple reduction from the MCIS problem in paths [8, 16]. The condition $\varepsilon > \log(n)/n$ is a technical condition because for arbitrary $\varepsilon > 0$, the lower bound becomes $\Omega(\log^*(\varepsilon n)/\varepsilon)$.

The upper bound is based on ideas that were developed for planar graphs in [8] and which are based the following core idea. First, one defines a weight function on edges, where for $\{u, v\} \in E$, we set $w(\{u, v\}) := \min\{w(u), w(v)\}$. The algorithm of [8] then computes a clustering of the nodes of V such that the total weight of the edges connecting nodes in different clusters is at most an $\varepsilon/2$-fraction of the total weight of all the edges. By computing an optimal weighted independent set within each cluster, one can then obtain a $(1 - \varepsilon)$-approximate MWIS solution of the given instance. The independent sets of the clusters are combined by taking the union and removing the smaller weight node of every intercluster edge. The clustering algorithm of [8] obtains clusters of diameter $\text{poly}(\varepsilon^{-1})$ in time $\text{poly}(\varepsilon^{-1}) \cdot O(\log^* n)$ in the LOCAL model. In trees, the clustering algorithm of [8] can be implemented in the CONGEST model. Further, in trees, an optimal MWIS solution inside each cluster can be computed in time linear in cluster diameter by using a straightforward dynamic programming algorithm. The core challenge for proving the upper bound of Theorem 1 is to obtain a clustering algorithm in which the maximum cluster diameter is only $O(1/\varepsilon)$. The algorithm of [8] consists of basic steps in which the cluster diameter grows by constant factor and the total weight of the intercluster edges shrinks by a constant factor. The cluster diameter however grows by a factor that is larger than the factor by which the total weight of the intercluster edges decreases. In our algorithm, we show that in trees, one can interleave those basic steps of [8] with steps in which the clusters are split into clusters of smaller diameter, without increasing the overall weight of the intercluster edges by too much.

1.3.2 Approximation of MWIS as a Function of the Arboricity

For the remainder of this section, we assume that we are given a graph $G = (V, E)$ of arboricity β. In [14], it is shown that there is a deterministic CONGEST algorithm to compute an independent set of weight at least $\frac{w(V)}{(4+\varepsilon)\beta}$ in randomized time $O\big(\frac{\log n \cdot \text{poly} \log \log n}{\varepsilon}\big)$ and in combination with an algorithm from [9] in deterministic time $O\big(\frac{\log^2 \beta \cdot \log(1/\varepsilon) \cdot \log n}{\varepsilon}\big)$. In the following, we show that for moderately small β both those bounds can be improved significantly.

▶ **Theorem 2.** *For any $\varepsilon > 0$, there is a $O\big(\frac{\log^2(\beta/\varepsilon)}{\varepsilon} + \log^* n\big)$-round deterministic* CONGEST *algorithm to compute an independent set I of weight $w(I) \geq (1 - \varepsilon) \cdot \frac{w(V)}{4 \cdot \beta}$ in any graph $G = (V, E)$ of arboricity $\beta \geq 1$ and with node weights $w : V \to \mathbb{R}_{\geq 0}$.*

The theorem can be proven by a relatively straightforward blackbox application of the local rounding framework of [9]. More concretely, we can show that Theorem 2 follows almost directly from Lemma 4.2 in [9]. Theorem 2 implies a $\frac{1-\varepsilon}{4\beta}$-approximation of MWIS. In [12], it is shown that a $\frac{1}{\lfloor (2+\varepsilon)\beta \rfloor}$-approximation can be computed deterministically in $O\big(\frac{\beta \cdot \log n}{\varepsilon}\big)$ CONGEST rounds. For $\varepsilon = 1/o(\sqrt{\beta})$, we significantly improve on this bound as stated by the following theorem.

▶ **Theorem 3.** *For any $\varepsilon > 0$, there is a $O\big(\frac{\log^3(\beta) \cdot \log(1/\varepsilon) \cdot \log n}{\varepsilon^2}\big)$-round deterministic* CONGEST *algorithm to compute an independent set I of weight $w(I) \geq (1 - \varepsilon) \cdot \frac{w(V)}{2\beta+1}$ in any graph $G = (V, E)$ of arboricity $\beta \geq 1$ and with node weights $w : V \to \mathbb{R}_{\geq 0}$.*

A major step towards proving Theorem 3 is an efficient algorithm to compute an independent set of weight arbitrarily close to a natural weighted generalization of the Caro-Wei bound (1). Specifically, we prove the following theorem.

▶ **Theorem 4.** *For any $\varepsilon > 0$, there is a $O\big(\frac{\log^3(\Delta)\cdot\log(1/\varepsilon)}{\varepsilon} + \log^* n\big)$-round deterministic* CONGEST *algorithm to compute an independent set I of weight $w(I) \geq (1-\varepsilon)\cdot\sum_{v\in V}\frac{w(v)}{\deg(v)+1}$ in any graph $G = (V, E)$ of maximum degree $\Delta \geq 1$ and with node weights $w : V \to \mathbb{R}_{\geq 0}$.*

Before discussing the main ideas needed to prove, Theorem 4, we discuss its direct application to approximating the maximum cardinality independent set problem in graphs of arboricity β. As proven in [6] and discussed in Section 1.2, the Caro-Wei bound (1) provides a $\frac{2\cdot(\sqrt{2}-1)}{\beta+1}$-approximation to the maximum cardinality independent set problem in graphs of arboricity β. Theorem 4 thus implies that such an approximation can be computed in $O\big(\frac{\log^3(\Delta)\cdot\log(1/\varepsilon)}{\varepsilon} + \log^* n\big)$ rounds in the CONGEST model. This can be further improved slightly. Because in a graph G of arboricity β, the average degree is at most 2β, the number of nodes of degree $> c\cdot\beta^2/\varepsilon$ is at most $2\varepsilon n/(c\beta)$, which can be shown to be at most $4\varepsilon/c\cdot\alpha(G)$, where $\alpha(G)$ denotes the independence number of G. It therefore suffices to apply Theorem 4 to the subgraph induced by the nodes of degree at most $c\beta^2/\varepsilon$ for a sufficiently large constant c. We thus obtain the following corollary.

▶ **Corollary 5.** *For any $\varepsilon > 0$, there is a $O\big(\frac{\log^3(\beta/\varepsilon)\cdot\log(1/\varepsilon)}{\varepsilon} + \log^* n\big)$-round deterministic* CONGEST *algorithm to compute a $(1-\varepsilon)\cdot\frac{2\cdot(\sqrt{2}-1)}{\beta+1}$-approximation to the maximum cardinality independent set problem in graphs of arboricity $\beta \geq 1$.*

We next sketch how Theorem 4 is proven and how it is used to prove Theorem 3. The starting point to proving Theorem 4 is an $O\big(\frac{\log^2\Delta\cdot\log(1/\varepsilon)}{\varepsilon} + \log^* n\big)$-round algorithm from [9] to compute an independent set of weight at least $(1-\varepsilon)\cdot\frac{w(V)}{\Delta+1}$ in a graph of maximum degree Δ. The idea now is to divide the nodes of G into $O(\log(\Delta)/\varepsilon)$ degree classes so that the degrees within one class differ by at most a factor $1+\varepsilon$. We then want to essentially apply the algorithm from [9] separately for each degree class, starting from the small degrees. This would work almost directly for unweighted graphs. To make it work, one just has to include all the still available nodes from the lower degree classes when computing the independent set of some degree class.

For weighted graphs, sequential composition of independent set is a bit more tricky. We can however utilize a technique developed in the context of the local-ratio method [2] and used in the context of distributed MWIS approximation in [3, 14, 9]. Given a graph G with node weights $w(v)$ and an independent set I_0, one can define a new weight function $w'(u) := \max\{0, w(u) - w(I_0 \cap N^+(u))\}$, where $N^+(u)$ denotes the inclusive neigborhood of u. If one then computes a second independent set I' consisting only of nodes v for which $w'(v) > 0$, the combined independent set $I := I' \cup (I_0 \setminus N^+(I'))$ is guaranteed to have a total weight of at least $w(I_0) + w'(I')$ (w.r.t. to the original weight function $w(\cdot)$). In this way, we can iterate through the $O(\log n/\varepsilon)$ degree classes in a similar way as in the unweighted case to obtain an independent set satisfying the requirements for Theorem 4.

We next discuss our algorithm to obtain the bound claimed by Theorem 3. Note that since the average degree of a graph of arboricity β can be almost 2β, we might have to compute an independent set of weight close to $w(V)$ divided by the average degree. There are graphs of arboricity β for which this is best possible (e.g., unweighted cliques) and we therefore in some cases cannot afford to allow too much "slack" when iteratively computing an independent set. A standard tool to algorithmically deal with graphs of bounded arboricity is the so-called H-decomposition as introduced in [4]. Because the average degree of a graph with arboricity

β is at most 2β, the number of nodes of degree more than $(2+\varepsilon)\beta$ is at most $(1-\Theta(\varepsilon))\cdot n$. When peeling off all nodes of degree at most $(2+\varepsilon)\beta$, we therefore get rid of a $\Theta(\varepsilon)$-fraction of all the nodes. Repeating this idea gives a decomposition of the graph into $O(\log(n)/\varepsilon)$ layers so that the nodes in each layer have at most $(2+\varepsilon)\beta$ neighbors in the same layer and higher layers. We are not aware of an algorithm to compute an independent set of weight close to $w(V)/(2\beta)$ that does not at least implicitly use an H-decomposition (which in most cases leads to a time complexity that is at least linear in the number of layers $\Theta(\log(n)/\varepsilon)$).

The highlevel idea of our approach is the following. We again use the framework of [2, 3, 14] to iteratively construct an independent set in the weighted setting. Because the nodes in the lowest layer are always the ones of bounded degree, we start our iterative construction at the lowest layer and work our way up the decomposition. Some of the main challenges with this approach already appear for the unweighted case and in the following highlevel discussion, we therefore focus on this case. Let V_0 be the nodes of the lowest layer of our decomposition. All nodes in V_0 have degree at most $(2+\varepsilon)\beta$ and we can therefore compute an independent set $I_0 \subseteq V_0$ of size close to $|V_0|/(2\beta)$ (or even close to the Caro-Wei bound (1) within V_0). However, even if the independent set I_0 is relatively large, it could consist of nodes that mostly have neighbors in $V \setminus V_0$ so that a lot of nodes in V_0 have no neighbor in I_0. When moving to the next layer, we can however not include any remaining nodes in V_0 because this might create some high degree nodes in the layer, we are considering. Our goal therefore should be to select an independent set I_0 in V_0 such that it removes most of the nodes in V_0. As an estimate of the number of removed nodes in V_0, we use $\sum_{v\in I_0}(\deg_0(v)+1)$, where $\deg_0(v)$ is the degree of v within V_0. If this sum is close to $|V_0|$, we are guaranteed to make sufficient progress (even if many nodes of V_0 remain because other nodes in V_0 have multiple neighbors in I_0). We can compute such an independent set I_0 as follows. We define a weight $\overline{w}(v) := \deg_0(v)+1$ for each node $v \in V_0$. Note that for an independent set I_0, we have $\overline{w}(I_0) = \sum_{v\in I_0}(\deg_0(v)+1)$. By using Theorem 4, we can compute an independent set I_0 of total weight $\overline{w}(I_0) \approx \sum_{v\in V_0} \frac{\overline{w}(v)}{\deg_0(v)+1} = |V_0|$, which is exactly what we need.

References

1 N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986. doi:10.1016/0196-6774(86)90019-2.

2 A. Bar-Noy, R. Bar-Yehuda, A. Freund, J. Naor, and B. Schieber. A unified approach to approximating resource allocation and scheduling. *J. of the ACM (JACM)*, 48:1069–1090, 2001. doi:10.1145/502102.502107.

3 Reuven Bar-Yehuda, Keren Censor-Hillel, Mohsen Ghaffari, and Gregory Schwartzman. Distributed approximation of maximum independent set and maximum matching. In *Proc. of 36th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 165–174, 2017. doi:10.1145/3087801.3087806.

4 L. Barenboim and M. Elkin. Sublogarithmic distributed mis algorithm for sparse graphs using Nash-Williams decomposition. In *Proc. 27th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 25–34, 2008.

5 L. Barenboim, M. Elkin, and F. Kuhn. Distributed $(\Delta+1)$-coloring in linear (in Δ) time. *SIAM J. Computing*, 43(1):72–95, 2015.

6 Ravi B. Boppana, Magnús M. Halldórsson, and Dror Rawitz. Simple and local independent set approximation. In *Proc. 25th Int. Coll. on Structural Information and Communication Complexity (SIROCCO)*, pages 88–101, 2018. doi:10.1007/978-3-030-01325-7_12.

7 Y. Caro. New results on the independence number. Technical report, Tel Aviv University, 1979.

8 A. Czygrinow, M. Hańćkowiak, and W. Wawrzyniak. Fast distributed approximations in planar graphs. In *DISC*, volume 5218 of *Lecture Notes in Computer Science*, pages 78–92. Springer, 2008. doi:10.1007/978-3-540-87779-0_6.

9 S. Faour, M. Ghaffari, C. Grunau, F. Kuhn, and V. Rozhon. Local distributed rounding: Generalized to mis, matching, set cover, and beyond. In *Proc. of the 2023 ACM-SIAM Symposium on Discrete Algorithms, (SODA 2023), Florence, Italy, January 22-25, 2023*, pages 4409–4447, SIAM, 2023.

10 M. Ghaffari. An improved distributed algorithm for maximal independent set. In *Proc. 27th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 270–277, 2016.

11 Mohsen Ghaffari and Christoph Grunau. Near-optimal deterministic network decomposition and ruling set, and improved MIS. In *Proc. 65th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 2148–2179, 2024. doi:10.1109/FOCS61266.2024.00007.

12 Y. Gil. Improved Deterministic Distributed Maximum Weight Independent Set Approximation in Sparse Graphs. In *27th International Conference on Principles of Distributed Systems (OPODIS 2023)*, pages 16:1–16:20, 2024.

13 Magnús M. Halldórsson and Christian Konrad. Computing large independent sets in a single round. *Distributed Computing*, 31:69–82, 2018. doi:10.1007/S00446-017-0298-Y.

14 K. Kawarabayashi, S. Khoury, A. Schild, and G. Schwartzman. Improved distributed approximations for maximum independent set. In *Proc. 34th Symp. on Distributed Computing (DISC)*, pages 35:1–35:16, 2020.

15 F. Kuhn, T. Moscibroda, and R. Wattenhofer. Local computation: Lower and upper bounds. *J. of the ACM*, 63(2), 2016. doi:10.1145/2742012.

16 C. Lenzen and R. Wattenhofer. Leveraging linial's locality limit. In *Proceedings of 22nd Symp. on Distributed Computing (DISC)*, pages 394–407, 2008.

17 M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15:1036–1053, 1986. doi:10.1137/0215074.

18 D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.

19 P. Turán. On an extremal problem in graph theory (in Hungarian). *Mat. Fiz. Lapok*, 48:436–452, 1941.

20 V. K. Wei. A lower bound on the stability number of a simple graph. Technical Report No. 81-11217-9, Bell Laboratories Technical Memorandum, 1981.