Brief Announcement: Distributed Minimum Vertex Coloring and Maximum Independent Set in Chordal Graphs*

Christian Konrad†
Department of Computer Science, University of Bristol
Bristol, United Kingdom
Christian.Konrad@bristol.ac.uk

Viktor Zamaraev‡
Department of Computer Science, Durham University
Durham, United Kingdom
Viktor.Zamaraev@durham.ac.uk

ABSTRACT

We give deterministic distributed $(1+\epsilon)$-approximation algorithms for Minimum Vertex Coloring and Maximum Independent Set on chordal graphs in the LOCAL model. Our coloring algorithm runs in $O(\frac{1}{\epsilon}\log n)$ rounds, and our independent set algorithm has a runtime of $O(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log^* n)$ rounds. For coloring, existing lower bounds imply that the dependencies on $\frac{1}{\epsilon}$ and $\log n$ are best possible. For independent set, we prove that $\Omega(\frac{1}{\epsilon})$ rounds are necessary.

Both our algorithms make use of the tree decomposition of the input chordal graph. They iteratively peel off interval subgraphs, which are identified via the tree decomposition of the input graph, thereby partitioning the vertex set into $O(\log n)$ layers. For coloring, each interval graph is colored independently, which results in various coloring conflicts between the layers. These conflicts are then resolved in a separate phase, using the particular structure of our partitioning. For independent set, only the first $O(\log \frac{1}{\epsilon})$ layers are required as they already contain a large enough independent set. We develop a $(1+\epsilon)$-approximation maximum independent set algorithm for interval graphs, which we then apply to those layers.

This work raises the question as to how useful tree decompositions are for distributed computing.

KEYWORDS

local model, approximation algorithms, minimum vertex coloring, maximum independent set, chordal graphs

ACM Reference Format:
Christian Konrad and Viktor Zamaraev. 2018. Brief Announcement: Distributed Minimum Vertex Coloring and Maximum Independent Set in Chordal Graphs. In *PODC '18: ACM Symposium on Principles of Distributed Computing, July 23–27, 2018, Egham, United Kingdom.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3212734.3212787

1 INTRODUCTION

The LOCAL Model. In the LOCAL model of distributed computation, the input graph $G = (V, E)$ represents a communication network, where every network node hosts a computational entity. Nodes have unique IDs. A distributed algorithm is executed on all network nodes simultaneously and proceeds in discrete rounds. Initially, besides their own IDs, nodes only know their neighbors. Each round consists of a computation and a communication phase. In the computation phase, nodes are allowed to perform unlimited computations. In the communication phase, nodes can exchange individual messages of unbounded sizes with each of their neighbors. The runtime of the algorithm is the total number of communication rounds, and the objective is to design algorithms that run in as few rounds as possible. The output is typically distributed: For vertex colorings, it is required that upon termination of the algorithm, every node knows its own color, and for independent sets, every node knows whether it participates in the independent set.

Distributed Vertex Coloring. Vertex coloring problems have been studied in distributed computational models since more than 30 years. Given a graph $G = (V, E)$, a *(legal) c-coloring* of G is an assignment $\gamma : V \to [c]$ of at most c colors to the nodes of G such that every pair of adjacent nodes receives different colors. The algorithmic challenge lies in computing colorings with few colors. The *chromatic number* $\chi(G)$ is the smallest c such that there is a c-coloring of G. The Minimum Vertex Coloring problem (MVC), one of the problems studied in this paper, asks to find a $\chi(G)$-coloring. This is a difficult task, even in the centralized setting: In general graphs, MVC is hard to approximate within a factor of $n^{1-\epsilon}$, for any $\epsilon > 0$ [22].

Most research papers on distributed vertex coloring address the problem of computing a $(\Delta+1)$-coloring, where Δ is the maximum degree of the input graph. In the distributed setting, this is a non-trivial task, and a long line of research has culminated in the randomized algorithm of Chang et al. [7], which runs in $O(\log^* n + 2^{\sqrt{\log\log n}})$ rounds, and the deterministic algorithm of Fraigniaud et al. [10], which runs in $O(\sqrt{\Delta}\log^{2.5}\Delta + \log^* n)$ rounds.

Only very few research papers address the MVC problem in a distributed model itself. On general graphs, the best distributed algorithm computes a $O(\log n)$-approximation in $O(\log^2 n)$ rounds [2]. This algorithm uses exponential time computations, which due to the computational hardness of MVC is necessary unless $P = NP$. Barenboim et al. [4] gave a $O(n^\epsilon)$-approximation algorithm that runs in $\exp O(1/\epsilon)$ rounds. Both the exponential time computations and the relatively large best known approximation factor of $O(\log n)$ on general graphs motivate the study of special graph classes. Besides results on graph classes with bounded chromatic number [3, 13, 14], the only graph class with unbounded chromatic number that has been addressed in the literature are *interval graphs*,

*A full version of this paper appears in [17].
†C. Konrad was supported by DIMAP and by EPSRC award EP/N011163/1.
‡V. Zamaraev was supported by EPSRC award EP/P020372/1.

PODC '18, July 23–27, 2018, Egham, United Kingdom

ACM ISBN 978-1-4503-5795-1/18/07.
https://doi.org/10.1145/3212734.3212787

which are the intersection graphs of intervals on the line. Halldórsson and Konrad gave a $(1+\epsilon)$-approximation algorithm for MVC on interval graphs that runs in $O(\frac{1}{\epsilon}\log^* n)$ rounds [16] (see also [15]). This work is the most relevant related work to our results.

Distributed Independent Sets. An *independent set* in a graph is a subset of pairwise non-adjacent nodes. In this paper we study the Maximum Independent Set problem (MIS) that asks to find an independent set of maximum size. Similar to MVC, the MIS problem is hard to approximate within a factor of $n^{1-\epsilon}$, for every $\epsilon > 0$ [22]. In the distributed setting, Luby [20] and independently Alon et al. [1] gave distributed $O(\log n)$ rounds maximal independent set algorithms more than 30 years ago. Improved results are possible for graphs with bounded maximum degree ([5, 11]) or on specific graph classes (e.g. [8, 21]). Using exponential time computations, a $(1+\epsilon)$-approximation to MIS can be computed in general graphs in $O(\frac{1}{\epsilon}\log n)$ rounds [6] (see also [12]). Deterministic distributed MIS algorithms may be inferior to randomized ones: It is known that every deterministic MIS O(1)-approximation algorithm on a path requires $\Omega(\log^* n)$ rounds [9, 18], while a simple randomized O(1) rounds O(1)-approximation algorithm exists [9].

Chordal Graphs. In this paper, we study MVC and MIS on *chordal graphs*. A graph is chordal, if every cycle on at least four nodes contains a *chord*, i.e., an edge connecting two non-consecutive nodes of the cycle. Chordal graphs play an important role in graph theory and have many applications, for example in belief propagation in machine learning. They constitute a superclass of interval graphs and trees and an inportant subclass of perfect graphs. The key motivations for our work are as follows:

1. Minimum Vertex Colorings. Since the best known distributed MVC algorithm only gives a $O(\log n)$-approximation, we are interested in pinpointing graph structures that are difficult to handle. In this paper, we show that MVC and MIS can both be well solved on chordal graphs. A defining feature of a chordal graph is that it does not contain any induced cycles of lengths at least 4. This in turn implies that difficult instances for distributed coloring necessarily contain induced cycles of length at least 4.

Furthermore, as previously mentioned, MVC can be solved well on interval graphs in the distributed setting [16]. We are therefore interested in identifying more general graph classes that admit distributed $(1+\epsilon)$-approximation algorithms for MVC. Since trees are chordal, Linial's lower bound for coloring trees applies [19]. Linial proved that coloring trees with a constant number of colors requires $\Omega(\log n)$ rounds, which gives a $\Omega(\log n)$ lower bound for any constant factor approximation to MVC on chordal graphs. This separates the difficulties of MVC on chordal and interval graphs. Furthermore, it is proved in [16] that a $(1+\epsilon)$-approximation to MVC on interval graphs requires $\Omega(\frac{1}{\epsilon})$ rounds. Combined, we obtain a $\Omega(\frac{1}{\epsilon}+\log n)$ lower bound on the round complexity for approximating MVC on chordal graphs within a factor of $1+\epsilon$.

2. Tree Decompositions. Tree decompositions are a powerful algorithmic tool that allow for the design of (centralized) linear time algorithms for NP-hard problems on graphs of bounded treewidth. They have played however only a minor role in the design of distributed algorithms. A *tree decomposition* of a graph $G=(V,E)$ is identified by a set of bags $S_1, S_2, \ldots \subseteq V$ that are arranged in a forest $\mathcal{T}$ such that every adjacent pair of nodes $uv \in E$ is contained in at least one bag, and, for any $v \in V$, the set of bags that contain v induces a non-empty tree in $\mathcal{T}$, which we denote $\mathcal{T}(v)$. It is well known that a graph G is chordal if and only if it has a tree decomposition $\mathcal{T}=(\mathcal{C},\mathcal{E})$ whose set of vertices (i.e., bags) $\mathcal{C}$ is the family of maximal cliques of G. We call such a tree decomposition a *clique forest* of chordal graph G.

One potential reason for the limited success of tree decompositions in distributed computing is that even simple graphs, such as a ring on n nodes, require that many bags of their tree decompositions consist of nodes that are at distance $\Omega(n)$ in the original graph. For these graphs, it is thus impossible that nodes obtain coherent local views of a tree decomposition in $o(n)$ rounds.

Chordal graphs are well-suited for studying distributed algorithms that exploit the input graph's tree decomposition, since in a chordal graph each bag is a clique. Thus, every bag that contains a node $v \in V$ further only contains nodes that lie in v's neighborhood. We exploit this locality property and show that in the LOCAL model, nodes can indeed obtain coherent local views of a global tree decomposition.

2 NEW RESULTS

We give deterministic distributed $(1+\epsilon)$-approximation algorithms for MVC and MIS on chordal graphs in the LOCAL model. Our algorithm for MVC runs in $O(\frac{1}{\epsilon}\log n)$ rounds, and our algorithm for MIS has a runtime of $O(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log^* n)$ rounds. For MVC, the dependencies of the runtime on $\log n$ and $\frac{1}{\epsilon}$ are best possible (though the existence of an algorithm with runtime $O(\frac{1}{\epsilon}+\log n)$ is not ruled out). For MIS, we prove that any, possibly randomized, $(1+\epsilon)$-approximation algorithm requires $\Omega(\frac{1}{\epsilon})$ rounds, even on paths.

We say that a path in a graph is *binary*, if all its vertices have degree at most 2 in the graph. We say that a binary path is *pendant*, if at least one of its end vertices has degree 1; and we say that it is *internal*, if all its vertices have degree 2.

2.1 Minimum Vertex Coloring

Our algorithm consists of the pruning, the coloring, and the color correction phases: In the pruning phase, the node set V is partitioned into at most $\lceil\log n\rceil$ *layers* $V_1,\ldots,V_{\lceil\log n\rceil}$ such that, for every $i \in [\lceil\log n\rceil]$, $G[V_i]$ constitutes an interval graph. In each step of the pruning phase, we remove every node $v \in U_i$ from the current graph $G[U_i]$ (we set $U_1=V$ and hence $G[U_1]=G$) whose corresponding subtree $\mathcal{T}(v)$ in the clique forest $\mathcal{T}_i$ of $G[U_i]$ is a subpath of a pendant path or an internal path of diameter at least $3k$. The set of removed nodes, which we denoted by V_i, induces an interval graph. We prove that the clique forest $\mathcal{T}_{i+1}$ of the resulting graph $G[U_{i+1}]$, where $U_{i+1}=U_i\setminus V_i$, can be obtained by removing all pendant paths and all internal paths of diameter at least $3k$ from $\mathcal{T}_i$. We also show that the pruning process ends after at most $\lceil\log n\rceil$ iterations and thus creates at most $\lceil\log n\rceil$ layers.

In the coloring phase, each interval graph $G[V_i]$ is colored with at most $(1+1/k)\chi(G[V_i])+1$ colors. In the centralized setting, it would be easy to color these interval graphs optimally. However, since we implement the algorithm in the distributed setting, and an optimal coloring on interval graphs cannot be computed distributively in few rounds, we impose a weaker quality guarantee that can be achieved distributively. The colorings of different layers

are computed independently from each other and do not give a coherent coloring of the entire input graph.

In the color correction phase, these incoherencies are corrected. To this end, the colors of $V_{\lceil \log n \rceil}$ remain unchanged and we correct the layers iteratively, starting with layer $\lceil \log n \rceil - 1$ and proceeding downwards to layer V_1. In a general step, for every path $\mathcal{P} \in \mathcal{L}_i$, we show that the nodes $W \subseteq V_i$ whose subtrees are subpaths of $\mathcal{P}$ form an interval graph together with those nodes in $\bigcup_{j \geq i+1} V_j$ that have coloring conflicts towards W. Notice that each path $\mathcal{P}$ connects to at most two bags (i.e., maximal cliques) in $\mathcal{T}_i$. The neighborhood of W thus consists of subsets of these (at most two) cliques, which implies that all conflicting nodes in $\bigcup_{j \geq i+1} V_j$ are included in these cliques as well. We then reuse a recoloring result previously proved by Halldórsson and Konrad [16], which shows that we can resolve all conflicts by changing the colors of those nodes in W that are at distance at most $k + 3$ from the (at most) two conflicting cliques.

Theorem 2.1. *For every $\epsilon \geq \frac{2}{\chi(G)}$, there is a deterministic $(1+\epsilon)$-approximation algorithm for MVC on chordal graphs that runs in $O(\frac{1}{\epsilon} \log n)$ rounds in the* LOCAL *model.*

2.2 Maximum Independent Set

Our distributed MIS algorithm uses an adapted version of the peeling process used in our coloring algorithm. The key observation that allows us to obtain a runtime of $o(\log n)$ is the fact that the first $O(\log \frac{1}{\epsilon})$ layers computed by our peeling process already contain an independent set of large enough size. Our algorithm proceeds as follows: In each iteration $i = 1, \ldots, O(\frac{1}{\epsilon})$ of the peeling process, we remove all pendant paths and all internal paths of large enough diameter[1] (let $\mathcal{L}_i$ be the set of removed paths) from the clique forest $\mathcal{T}_i$ of the graph induced by the remaining nodes. Next, we compute large independent sets among the nodes whose trees are included in each path $\mathcal{P} \in \mathcal{L}_i$. If $\mathcal{P}$ has a large independence number[2] then we run our $(1+\epsilon)$-approximation algorithm for interval graphs in $O(\frac{1}{\epsilon} \log^* n)$ rounds (which we design as an auxiliary tool for our main algorithm for chordal graphs). If $\mathcal{P}$ has small independence number we need to compute an optimal independent set in order to locally stay within a $(1+\epsilon)$-approximation guarantee. This can be achieved using only $O(\frac{1}{\epsilon})$ rounds, since paths with small independence number necessarily have small diameter. The runtime is thus dominated by the product of the number of iterations $O(\log \frac{1}{\epsilon})$ and the $O(\frac{1}{\epsilon} \log^* n)$ runtime of our MIS algorithm for interval graphs.

Theorem 2.2. *For any $\epsilon \in (0, 1/2)$, there is a deterministic $(1+\epsilon)$-approximation algorithm for MIS on chordal graphs that runs in $O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}) \log^* n)$ rounds in the* LOCAL *model.*

2.3 Lower Bound on the Complexity for MIS

Using an indistinguishability argument, we obtain a lower bound on the round complexity for MIS on paths.

Theorem 2.3. *For every $\epsilon > 0$ and n large enough, every randomized algorithm in the* LOCAL *model with expected approximation factor at most $1 + \epsilon$ for MIS requires $\Omega(\frac{1}{\epsilon})$ rounds.*

[1] The *diameter* of a path in the clique forest is the maximum distance in the graph between the nodes belonging to the bags of the path.

[2] The *independence number* of a path in the clique forest is the maximum number of pairwise non-adjacent nodes that belongs to the bags of the path.

REFERENCES

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

[2] Leonid Barenboim. 2012. On the Locality of Some NP-complete Problems. In *Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming - Volume Part II (ICALP'12)*. Springer-Verlag, Berlin, Heidelberg, 403–415. https://doi.org/10.1007/978-3-642-31585-5_37

[3] Leonid Barenboim and Michael Elkin. 2010. Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition. *Distributed Computing* 22, 5 (2010), 363–379.

[4] Leonid Barenboim, Michael Elkin, and Cyril Gavoille. 2015. A Fast Network-Decomposition Algorithm and Its Applications to Constant-Time Distributed Computation. In *Post-Proceedings of the 22Nd International Colloquium on Structural Information and Communication Complexity - Volume 9439 (SIROCCO 2015)*. Springer-Verlag New York, Inc., New York, NY, USA, 209–223. https://doi.org/10.1007/978-3-319-25258-2_15

[5] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. 2016. The Locality of Distributed Symmetry Breaking. *J. ACM* 63, 3, Article 20 (June 2016), 45 pages. https://doi.org/10.1145/2903137

[6] Marijke H.L. Bodlaender, Magnús M. Halldórsson, Christian Konrad, and Fabian Kuhn. 2016. Brief Announcement: Local Independent Set Approximation. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC '16)*. ACM, New York, NY, USA, 93–95. https://doi.org/10.1145/2933057.2933068

[7] Yi-Jun Chang, Wenzheng Li, and Seth Pettie. 2018. An Optimal Distributed $(\Delta + 1)$-Coloring Algorithm?, In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing. *arXiv preprint https://arxiv.org/abs/1711.01361*.

[8] Richard Cole and Uzi Vishkin. 1986. Deterministic Coin Tossing with Applications to Optimal Parallel List Ranking. *Inf. Control* 70, 1 (July 1986), 32–53. https://doi.org/10.1016/S0019-9958(86)80023-7

[9] Andrzej Czygrinow, Michal Hańćkowiak, and Wojciech Wawrzyniak. 2008. Fast Distributed Approximations in Planar Graphs. In *Distributed Computing*, Gadi Taubenfeld (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 78–92.

[10] Pierre Fraigniaud, Marc Heinrich, and Adrian Kosowski. 2016. Local Conflict Coloring. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*. 625–634.

[11] Mohsen Ghaffari. 2016. An Improved Distributed Algorithm for Maximal Independent Set. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '16)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 270–277. http://dl.acm.org/citation.cfm?id=2884435.2884455

[12] Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. 2017. On the Complexity of Local Distributed Graph Problems. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*. ACM, New York, NY, USA, 784–797. https://doi.org/10.1145/3055399.3055471

[13] Mohsen Ghaffari and Christina Lymouri. 2017. Simple and Near-Optimal Distributed Coloring for Sparse Graphs. In *Distributed Computing: 31th International Symposium, DISC 2017, Vienna, Austria, October 16-20, 2017. Proceedings*.

[14] Andrew V. Goldberg, Serge A. Plotkin, and Gregory E. Shannon. 1988. Parallel Symmetry-Breaking in Sparse Graphs. *SIAM J. Discrete Math.* 1, 4 (1988), 434–446. https://doi.org/10.1137/0401044

[15] Magnús M. Halldórsson and Christian Konrad. 2014. *Distributed Algorithms for Coloring Interval Graphs*. Springer, 454–468. https://doi.org/10.1007/978-3-662-45174-8_31

[16] Magnús M. Halldórsson and Christian Konrad. 2017. Improved Distributed Algorithms for Coloring Interval Graphs with Application to Multicoloring Trees. In *Post-Proceedings of the 24th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2017)*.

[17] Christian Konrad and Viktor Zamaraev. 2018. Distributed Coloring and Independent Set in Chordal Graphs. *arXiv preprint https://arxiv.org/abs/1805.04544* (2018).

[18] Christoph Lenzen and Roger Wattenhofer. 2008. Leveraging Linial's Locality Limit. In *Distributed Computing*, Gadi Taubenfeld (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 394–407.

[19] Nathan Linial. 1992. Locality in Distributed Graph Algorithms. *SIAM J. Comput.* 21, 1 (Feb. 1992), 193–201. https://doi.org/10.1137/0221015

[20] Michael Luby. 1986. A Simple Parallel Algorithm for the Maximal Independent Set Problem. *SIAM J. Comput.* 15, 4 (1986), 1036–1053. https://doi.org/10.1137/0215074

[21] Johannes Schneider and Roger Wattenhofer. 2010. An optimal maximal independent set algorithm for bounded-independence graphs. *Distributed Computing* 22, 5 (01 Aug 2010), 349–361. https://doi.org/10.1007/s00446-010-0097-1

[22] David Zuckerman. 2007. Linear Degree Extractors and the Inapproximability of Max Clique and Chromatic Number. *Theory of Computing* 3, 1 (2007), 103–128.