2022 IEEE 42nd International Conference on Distributed Computing Systems (ICDCS) | 978-1-6654-7177-0/22/$31.00 ©2022 IEEE | DOI: 10.1109/ICDCS54860.2022.00153

Awake-Efficient Distributed Algorithms for Maximal Independent Set

Khalid Hourani
Department of Computer Science
University of Houston
kmhourani@uh.edu

Gopal Pandurangan
Department of Computer Science
University of Houston
gopalpandurangan@gmail.com

Peter Robinson
School of Computer & Cyber Sciences
Augusta University
perobinson@augusta.edu

***Abstract*—We present a simple algorithmic framework for designing efficient distributed algorithms for the fundamental symmetry breaking problem of Maximal Independent Set (MIS) in the *sleeping model* [Chatterjee et al, PODC 2020]. In the sleeping model, only the rounds in which a node is *awake* are counted for the *awake complexity*, while *sleeping* rounds are ignored. This is motivated by the fact that a node spends resources only in its awake rounds and hence the goal is to minimize the awake complexity.**

Our framework allows us to design distributed MIS algorithms that have $\mathcal{O}(\text{polyloglog}\, n)$ (worst-case) awake complexity in certain important graph classes which satisfy the so-called *adjacency property*. Informally, the adjacency property guarantees that the graph can be partitioned into an appropriate number of classes so that each node has at least one neighbor belonging to every class. Graphs that can satisfy the adjacency property are random graphs with large clustering coefficient such as *random geometric graphs* as well as *line graphs* of regular (or near regular) graphs.

We first apply our framework to design two randomized distributed MIS algorithms for random geometric graphs of arbitrary dimension d (even non-constant). The first algorithm has $\mathcal{O}(\text{polyloglog}\, n)$ (worst-case) awake complexity with high probability, where n is the number of nodes in the graph.[1] This means that any node in the network spends only $\mathcal{O}(\text{polyloglog}\, n)$ awake rounds; this is almost exponentially better than the (traditional) time complexity of $\mathcal{O}(\log n)$ rounds (where there is no distinction between awake and sleeping rounds) known for distributed MIS algorithms on general graphs or even the faster $\mathcal{O}\left(\sqrt{\frac{\log n}{\log\log n}}\right)$ rounds known for Erdos-Renyi random graphs. However, the (traditional) time complexity of our first algorithm is quite large—essentially proportional to the degree of the graph. Our second algorithm has a slightly worse awake complexity of $\mathcal{O}(d\, \text{polyloglog}\, n)$, but achieves a significantly better time complexity of $\mathcal{O}(d \log n\, \text{polyloglog}\, n)$ rounds whp.

We also show that our framework can be used to design $\mathcal{O}(\text{polyloglog}\, n)$ awake complexity MIS algorithms in other types of random graphs, namely an augmented Erdos-Renyi random graph that has a large clustering coefficient.

I. Introduction

Ad hoc wireless, sensor, and IoT networks are inherently resource-constrained and hence it is important to design distributed algorithms that are resource-efficient. Towards this goal, Chatterjee et al. [1] proposed the *sleeping* model for distributed algorithms, where nodes can operate in two modes: *awake* and *sleeping*. Each node can choose to enter the awake or asleep state at the start of any specified round. In the sleeping mode, a node cannot send, receive, or perform any local computation; messages sent to it are also lost. However, resources (such as energy) utilized in sleeping mode are negligible and thus only awake rounds are counted. The goal in the sleeping model is to design distributed algorithms that solve problems in a small number of awake rounds, i.e., have small *awake complexity*, which is the (worst-case) number of awake rounds needed by any node until it terminates.

While the main goal is to minimize awake complexity, we would also like to minimize the (traditional) *time complexity* of the algorithm, which counts the (worst-case) total number of rounds taken by any node, including both awake and sleeping rounds.

Our main goal in this paper is to design distributed algorithms that have small awake complexity for the *Maximal Independent Set (MIS) problem*, a fundamental problem in distributed computing with various applications.

II. Our Main Results

We present a simple algorithmic framework for designing efficient distributed algorithms for the fundamental symmetry breaking problem of Maximal Independent Set (MIS) in the *sleeping model*.

We then apply our framework to design two randomized distributed MIS algorithms for random geometric graphs of arbitrary dimension d (even non-constant). Both of our MIS algorithms for random geometric graphs have significantly lower awake complexity while the second algorithm also has a time complexity that is comparable to the best known algorithms. The first algorithm has $\mathcal{O}(\text{polyloglog}\, n)$ (worst-case) awake complexity with high probability, where n is the number of nodes in the graph. This means that any node in the network spends only $\mathcal{O}(\text{polyloglog}\, n)$ awake rounds; this is almost exponentially better than the (traditional) time complexity of $\mathcal{O}(\log n)$ rounds (where there is no distinction between awake and sleeping rounds) known for distributed MIS algorithms on general graphs (e.g., [2], [3], [6], [7]) or even the faster $\mathcal{O}\left(\sqrt{\frac{\log n}{\log\log n}}\right)$ rounds known for Erdos-Renyi random graphs [5].

We note that the (traditional) time complexity of our MIS algorithm for random geometric graphs is quite large—

Supported, in part, by NSF grants IIS-1633720, CCF-1717075, and BSF award 2016419.

[1]With high probability (whp), means "with probability at least $1 - 1/n^c$, for some constant $c > 0$."

2575-8411/22/$31.00 ©2022 IEEE
DOI 10.1109/ICDCS54860.2022.00153

essentially proportional to the degree of the graph. We then use the same framework to present a second algorithm (a modification of the first) that has a slightly worse awake complexity of $\mathcal{O}(d \operatorname{polyloglog} n)$, but achieves a significantly better time complexity of $\mathcal{O}(d \log n \operatorname{polyloglog} n)$ rounds whp, where d is the dimension of the random geometric graph.

We note that random geometric graphs of *constant* dimension are growth-bounded graphs and hence the algorithm of Schneider and Wattenhofer [8] can be applied to obtain an $\mathcal{O}(\log^* n)$ round deterministic algorithm in the traditional model. However, if the dimension is superconstant, say if the dimension is a slow growing function of the network size n then the running time of the algorithm of Schneider and Wattenhofer becomes large, e.g., if the dimension is $\mathcal{O}(\log \log n)$ then the running time becomes $\mathcal{O}\big((3 + 2^{\log \log n})^{\log \log n}\big) = \mathcal{O}\big((\log n)^{\log \log n}\big)$. On the other hand, our algorithm has $\mathcal{O}(\operatorname{polyloglog} n)$ awake complexity *regardless of the dimension*.

We also apply our framework to design $\mathcal{O}(\operatorname{polyloglog} n)$ (worst-case) awake complexity in random graphs with large clustering coefficient. These random graphs are constructed by taking Erdos-Renyi random graphs and adding additional random edges between neighbors of nodes to increase the clustering coefficient. These graphs are similar to Watts-Strogatz type random graphs [9] that are used to model social networks.

III. High-level Overview of our Algorithms

Our algorithm uses a simple framework to achieve very low (worst-case) awake complexity for MIS. We partition the nodes into a *small* number of disjoint sets called *classes*, say $C_1, C_2, \ldots, C_\ell$ and then solve MIS *sequentially* in each of these classes. When the MIS algorithm runs in class C_j ($1 \leq j \leq \ell$), all classes C_k ($k > j$) are asleep. Also all classes C_i ($i < j - 1$) are asleep. Note that class C_{j-1} is awake (for $j > 1$). A crucial property is that each node in C_j is able to obtain the MIS status of its (respective) neighbors belonging to lower indexed classes (i.e., $C_1, \ldots, C_{j-1}$) using the so-called *adjacency property* of random geometric graphs. The adjacency property holds because random geometric graphs have large *clustering coefficient*; informally, this means that many neighbors of a node v are themselves mutually neighbors with high probability. This property allows efficient "hand over" of information: a node in C_j will get to know the status of neighbors belonging to smaller indexed classes $C_1 \ldots C_{j-2}$ (nodes in these classes have already finished and are asleep) from neighbors in C_{j-1}. Thus only classes C_j and C_{j-1} are awake at a time. If we choose the number of disjoint sets, ℓ, to be large, then the maximum degree of the subgraph induced by each class is small, which allows us to use the MIS algorithm due to Rozhon and Ghaffari [7] that runs in $\mathcal{O}(\operatorname{polyloglog} n)$ rounds in graphs of $\operatorname{polylog} n$ degree. This technique guarantees small *awake complexity*, but if the number of disjoint sets is large, then, since the sets are sequentially processed, the overall running time (including sleeping rounds) is large, i.e., at least $\Omega(\ell)$, the number of disjoint sets.

In our second algorithm, we aim to significantly reduce the overall running time, and hence we partition the vertex set into only $\mathcal{O}(\log n)$ disjoint classes of *geometrically* increasing sizes. This creates a key technical challenge of finding MIS in large-sized graphs (which may have large maximum degree). To overcome this, we use a second main idea. We show that the *residual* subgraph induced by these nodes (which excludes neighbors of MIS nodes in prior classes) has $\mathcal{O}\big(4^d \log n\big)$ degree, for dimension d. Showing this bound involves a careful argument, since the MIS nodes that cover the classes $C_1, \ldots, C_{i-1}$ depend on the nodes already chosen in these classes and on the algorithm used to determine the MIS (we use the MIS algorithm of [7]). We apply the principle of deferred decisions to overcome the dependencies and show that no matter what MIS nodes are chosen to cover the classes $C_1, \ldots, C_{i-1}$, the residual graph induced in the class C_i is small. The above degree bound allows us to use the MIS algorithm due to Rozhon and Ghaffari [7] to obtain an algorithm that has $\mathcal{O}(d \operatorname{polyloglog} n)$ awake complexity and small overall running time of $\mathcal{O}(d \log n \operatorname{polyloglog} n)$.

Our framework also works for other types of random graphs that have large clustering coefficient which ensures that the adjacency property is satisfied.

Our algorithms work in the LOCAL model in general, since the amount of information that needs to be transferred depends on the size of the MIS in a node's neighborhood. If this size is not too large, then our algorithm can work also in the CONGEST model. For example, for random geometric graphs of constant dimension, the number of MIS nodes in the neighborhood of a node is constant and our algorithm works in the CONGEST model for such graphs.

We refer to the full paper [4] for the full description of the algorithms, their correctness and analysis proofs, and other details.

References

[1] Soumyottam Chatterjee, Robert Gmyr, and Gopal Pandurangan. Sleeping is efficient: MIS in $O(1)$-rounds node-averaged awake complexity. In *PODC '20*, pages 99–108. ACM, 2020.

[2] Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *ACM-SIAM SODA*, pages 270–277, 2016.

[3] Mohsen Ghaffari, Christoph Grunau, and Václav Rozhon. Improved deterministic network decomposition. In *ACM-SIAM SODA*, pages 2904–2923, 2021.

[4] Khalid Hourani, Gopal Pandurangan, and Peter Robinson. Awake-efficient distributed algorithms for maximal independent set, 2021. Manuscript, to appear in arxiv.

[5] K. Krzywdziński and K. Rybarczyk. Distributed algorithms for random graphs. *Theoretical Computer Science*, 605:95–105, 2015.

[6] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.

[7] Václav Rozhon and Mohsen Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *STOC*, pages 350–363, 2020.

[8] Johannes Schneider and Roger Wattenhofer. A log-star distributed maximal independent set algorithm for growth-bounded graphs. In *ACM PODC*, pages 35–44, 2008.

[9] Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *Nature*, 393(6684):440–442, 1998.