Topic 9: Topological Methods for Distributed Graph Algorithms

Michael Kerber and Yannic Maus

February 3, 2024

Disclaimer: We do not expect PhD candidates to necessarily already have a background in topological methods and/or distributed graph algorithms. If a candidate has excellent mathematical skills or a solid foundation in theoretical computer science, then the required knowledge can be obtained during the initial part of the 4 year period

Scientific background Distributed systems and algorithms are omnipresent in our daily life and there is no doubt that the future will be even more *distributed*. More than two decades ago, topological methods (originating from Mathematics) have been used for analyzing distributed algorithms in asynchronous shared memory systems. Due to the large interest of this prime application of topology to a research area within computer science, many of these results have been summarized in a book by Herlihy, Feichtner-Kozloz, and Rajsbaum [11].

Another branch of distributed computing deals with distributed graph algorithms in synchronous message passing systems. There are many models to study these, with the LOCAL model, introduced by Linial more than 30 years ago [12, 14], being one of the central models. In the model, a communication network is abstracted as a graph G with nodes being computing entities and edges forming communication links. Then nodes communicate with each other in synchronous rounds in order to solve some graph problem on G in as few communication rounds as possible. One central problem of the area with many unresolved open questions concerns the graph coloring problem. Here, the objective is to assign a color to each node of the graph such that adjacent vertices receive different colors [2]. At the end of the computation each node of the network has to output its own color. The problem has many variants, depending on the number of colors, the type of input graphs, and whether we color nodes or edges. Many of these coloring problems are prototypical for a broader class of local constraint satisfaction problems and are at the heart of distributed complexity theory. In the last decade, the area of distributed graph algorithms has undergone a significant development resulting in many faster algorithms and several new algorithmic lower bounds for such problems, e.g., [9, 7, 15, 8, 1]. Still, there is a gap between upper and lower bounds that seemingly cannot be bridged with current techniques. The early topological techniques do not reason about models like the LOCAL model. However, for selected problems, recent achievements, e.g., [3, 5, 4, 6], have been successful in applying topological methods to such and similar settings and there is ample potential for extending this line of research.

Aims. The objective of this dissertation is to use topological methods to obtain new results in the area of distributed graph algorithms. We aim at devising new algorithms and new impossibility results. The focus of our study are *local constraint satisfaction problems* like graph coloring problems. One main focus lies on improving lower bounds for graph coloring problems, where (despite significant progress on related problems) Linial's 30 year old $\Omega(\log^* n)$-round lower bound is still the state of the art for coloring graphs with a number of colors that exceeds the maximum degree of the graph [12]. Another focus lies on understanding local constraint satisfaction problems in which the nodes of the input graph are equipped with additional inputs [13]. This setting is particularly suited for the investigation with topological methods and provides ample low-risk starting problems.

Approaches/Methods. One integral part of this dissertation project is to develop new topological tools to reason on distributed algorithms. The high level approach is to formalize the input space as well as the (feasible) output space of a distributed algorithm as topological objects $\mathcal{I}$ and $\mathcal{O}$, respectively. A (correct) distributed algorithm is a topological transformation of $\mathcal{I}$ to $\mathcal{O}$ where the allowed transformations depend on the computational model. Then, topological properties of $\mathcal{I}$, $\mathcal{O}$, and the allowed transformations imply the (non)-existence of efficient algorithms. The initial starting point is to understand $\mathcal{I}$, $\mathcal{O}$ for concrete problems such as the list coloring problem. In *list coloring problems*, each node is equipped with a list of available colors and has to output one of these colors. Almost all state of the art graph coloring algorithms actually solve the (possibly harder) list coloring problem. The problem is approachable, as currently, the already developed topological methods are most powerful in proving impossibilities when nodes are equipped with some input and the nodes' lists serve as such an input.

Another orthogonal line of attack is given through so called neighborhood graphs that date back to Linial's first paper in the area [12, 10]. They are at the center of proving lower bounds for the graph coloring problem. It is known that lower bounds on the chromatic number of these graphs provide lower bounds on the complexity of the distributed graph coloring problem. The Borsuk-Ulam theorem, a famous result in topology, relates the connectedness of certain simplicial complexes and the chromatic number of these graphs. We will use these topological tools with the goal to determine the complexity of the distributed graph coloring problem.

In general, the research theme for this dissertation is timely and builds up on a new line of research of using topological methods for analyzing synchronous distributed graph algorithms [3, 5, 4, 6]. Introducing new tools is essential for the area to narrow the gap between existing upper and lower bounds and topological tools will surely provide ample new insights.

Time Frame. This dissertation is planned for 4 years. The vast number of precise local constraint satisfaction problems limits the risk for doctoral student. As an additional risk mitigation strategy each graph problem can be restricted to a special graph class, often still yielding far reaching results but being more approachable. Further, the topic offers ample possibilities for an academic career of the PhD student afterwards. It is even likely, that the topic will be established as its own sub community.

Primary researchers involved. This project is supervised by Michael Kerber and Yannic Maus.

Michael Kerber is a full professor at TU Graz and an internationally recognized expert in computational topology and geometry. He works mostly on fast algorithms for computing and comparing persistence diagrams, and on multi-parameter persistent homology. Kerber was program committee chair of the *2022 Symposium on Computational Geometry*, the main conference in computational geometry and topology. Four PhD students have completed their thesis under his supervision. His group currently consists of one postdoc and 4 PhD students.

Yannic Maus is an assistant professor (tenure-track) in the computer science department at TU Graz. He is an expert for distributed graph algorithms and spent two years at the Technion. For his PhD thesis, that laid the foundation for a generic and complexity theoretic treatment of distributed graph algorithms, he received several prices including the *ACM/EATCS Principles of Computing Doctoral Dissertation Award* and the *GI Dissertation Award*. He is co-supervising one PhD student and by the end of 2023 his group will consist of two (additional) PhD students.

References

[1] A. Balliu, S. Brandt, J. Hirvonen, D. Olivetti, M. Rabie, and J. Suomela. Lower bounds for maximal matchings and maximal independent sets. *J. ACM*, 68(5):39:1–39:30, 2021.

[2] L. Barenboim and M. Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments.* Morgan & Claypool Publishers, 2013.

[3] A. Castañeda, P. Fraigniaud, A. Paz, S. Rajsbaum, M. Roy, and C. Travers. A topological perspective on distributed network algorithms. In *Structural Information and Communication Complexity - 26th International Colloquium, SIROCCO*, pages 3–18, 2019.

[4] P. Fraigniaud, R. Gelles, and Z. Lotker. The topology of randomized symmetry-breaking distributed computing. In *PODC '21: ACM Symposium on Principles of Distributed Computing*, pages 415–425. ACM, 2021.

[5] P. Fraigniaud and A. Paz. The topology of local computing in networks. In *47th Intern. Colloquium on Automata, Languages, and Programming, ICALP*, pages 128:1–128:18, 2020.

[6] P. Fraigniaud, A. Paz, and S. Rajsbaum. A speedup theorem for asynchronous computation with applications to consensus and approximate agreement. In *PODC '22: ACM Symposium on Principles of Distributed Computing*, pages 460–470. ACM, 2022.

[7] M. Ghaffari, D. G. Harris, and F. Kuhn. On derandomizing local distributed algorithms. pages 662–673, 2018.

[8] M. Ghaffari and F. Kuhn. Deterministic distributed vertex coloring: Simpler, faster, and without network decomposition. pages 1009–1020, 2021.

[9] M. Ghaffari, F. Kuhn, and Y. Maus. On the complexity of local distributed graph problems. pages 784–797. ACM, 2017.

[10] D. Hefetz, Y. Maus, F. Kuhn, and A. Steger. A polynomial lower bound for distributed graph coloring in a weak LOCAL model. pages 99–113, 2016.

[11] M. Herlihy, D. Feichtner-Kozlov, and S. Rajsbaum. *Distributed Computing Through Combinatorial Topology.* Morgan Kaufmann, Amsterdam, 2014.

[12] N. Linial. Locality in distributed graph algorithms. *SIAM J. on Computing*, 21(1):193–201, 1992.

[13] M. Naor and L. Stockmeyer. What can be computed locally? pages 184–193, 1993.

[14] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach.* SIAM, 2000.

[15] V. Rozhon and M. Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. pages 350–363. ACM, 2020.