

Graph Algorithms: Distributed Meets Dynamic

Keren Censor-Hillel^{*1}, Yasamin Nazari^{*2}, Eva Rotenberg^{*3},
Thatchaphol Saranurak^{*4}, and Martín Costa^{†5}

1 Technion – Haifa, IL. ckeren@cs.technion.ac.il

2 VU Amsterdam, NL. y.nazari@vu.nl

3 Technical University of Denmark – Lyngby, DK. eva@rotenberg.dk

4 University of Michigan – Ann Arbor, US. thsa@umich.edu

5 University of Warwick – Coventry, GB. martin.costa@warwick.ac.uk

Abstract

This report contains the program and outcomes of the Dagstuhl Seminar “Graph Algorithms: Distributed Meets Dynamic”. The field of dynamic graph algorithms address recomputation following edge/vertex insertions/deletions in the input graph. Distributed graph algorithms focus on computing when the input resides across multiple machines, that need to communicate for their joint computation. The seminar brought together researchers from the two communities of dynamic graph algorithms and distributed computing. The aim was to transfer knowledge and techniques between the dynamic and distributed settings, build new collaborations and to explore research directions on computational models of the combined distributed dynamic setting.

Seminar November 17–22, 2024 – <https://www.dagstuhl.de/24471>

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms; Theory of computation → Distributed algorithms; Theory of computation → Dynamic graph algorithms

Keywords and phrases distributed algorithms, dynamic algorithms

Digital Object Identifier 10.4230/DagRep.14.11.92


1 Executive Summary

Thatchaphol Saranurak (University of Michigan – Ann Arbor, US)

Keren Censor-Hillel (Technion – Haifa, IL)

Yasamin Nazari (VU Amsterdam, NL)

Eva Rotenberg (Technical University of Denmark – Lyngby, DK)

License  Creative Commons BY 4.0 International license

© Thatchaphol Saranurak, Keren Censor-Hillel, Yasamin Nazari, and Eva Rotenberg

This seminar brought together researchers from two research communities on distributed and dynamic graph algorithms. In dynamic settings, algorithms receive updates to the graph, and should update the solution efficiently. In many distributed settings, computation takes place in a decentralized manner over a network or multiple machines. The input is distributed among multiple computing devices that need to communicate to find a solution to a given problem. The first connection between these models that was explored in the seminar is the technical toolbox. In recent years, there has been a growing number of results in both areas that use similar algorithmic tools and technical ideas. Examples of these common

* Editor / Organizer

† Editorial Assistant / Collector



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Graph Algorithms: Distributed Meets Dynamic, *Dagstuhl Reports*, Vol. 14, Issue 11, pp. 92–107

Editors: Keren Censor-Hillel, Yasamin Nazari, Eva Rotenberg, and Thatchaphol Saranurak



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

technical tools include expander-based techniques, truncating Vizing chains for improved edge colouring, algebraic techniques based on matrix multiplication, spanners and hopsets, clustering coresets, etc.

While some of these techniques have found applications in one or both of these models, there are often technical challenges in applying them in different settings and knowledge of models of specific tools is needed. This is why the seminar aimed to increase collaborations and discussions between the two communities. The second type of connection is from a modeling point of view, motivated by practice. In many practical big data settings the input is both distributed and dynamic at the same time. First, it is stored in a decentralized way over many machines, and second, the input changes over time. Few talks focused on the models that combine multiple computational settings.

The format of the seminar was as follows: the first day included an introduction session in addition to an ice-breaker “research speed-dating” event. In this ice-breaker event, participants were paired together in several short sessions with an attempt to pair participants who are not already familiar with each other’s work. We had 4 plenary overview/survey talks, and the rest of the presentations were 20 minute talks and 10 minutes of questions and discussions. We also had two engaging open problem sessions in addition to various allocated times for collaboration.

We aimed to invite a diverse group of participants, both based on demographic factors and geographic affiliation, as well as seniority level. We had participants from Europe, North America, and Asia, and with experiences ranging from PhD students and postdocs to academic and industry research with varied level of seniority.

2 Table of Contents

Executive Summary

Thatchaphol Saranurak, Keren Censor-Hillel, Yasamin Nazari, and Eva Rotenberg 92

Overview of Talks

Fully Dynamic k-Median with Near-Optimal Update Time and Recourse <i>Sayan Bhattacharya</i>	96
Deterministic expander routing: faster and more versatile <i>Yi-Jun Chang</i>	97
Fully Dynamic k-Clustering with Fast Update Time and Small Recourse <i>Martín Costa</i>	97
Tree-Packing Revisited: Faster Fully Dynamic Min-Cut and Arboricity <i>Tijn de Vos</i>	98
Improved All-Pairs Approximate Shortest Paths in Congested Clique <i>Michal Dory</i>	98
Distributed Graph Algorithms with Predictions <i>Faith Ellen</i>	99
Expanders via Local Edge Flips <i>George Giakkoupis</i>	99
Dynamic Implicit Coloring Algorithms For Low-Arboricity Graphs <i>Christoph Grunau</i>	99
Towards Scalable and Practical Batch-Dynamic Connectivity <i>Jakub Lacki and Adam Karczmarz</i>	100
One Round Distributed Clique Listing <i>Quanquan C. Liu</i>	100
Smoothed Analysis of Dynamic Networks <i>Ami Paz</i>	101
Distributed Replacement Paths and Distance Sensitivity Oracles <i>Vijaya Ramachandran</i>	101
Vizing's Theorem in Near-Linear Time <i>Shay Solomon</i>	102
Locality in distributed, dynamic, online, and quantum settings <i>Jukka Suomela</i>	102
Fully Dynamic Exact SSSP <i>Jan van den Brand</i>	103
Dynamic Rounding: Success for Matching; Open Questions Beyond <i>David Wajc</i>	103
Faster $(\Delta + 1)$ -Edge Coloring: Breaking the $m\sqrt{n}$ Time Barrier <i>Tianyi Zhang</i>	104

Open problems

An Open Problem in Dynamic Edge Coloring
Martín Costa 104

Open Problem: Low-Degree Spanning Tree
Thatchaphol Saranurak 105

Open Problem: Dynamic Brooks’ Theorem
Tianyi Zhang 105

Participants 107

3 Overview of Talks

3.1 Fully Dynamic k -Median with Near-Optimal Update Time and Recourse

Sayan Bhattacharya (University of Warwick – Coventry, GB)

License © Creative Commons BY 4.0 International license
© Sayan Bhattacharya

Joint work of Sayan Bhattacharya, Martín Costa, Ermiya Farokhnejad

Main reference Sayan Bhattacharya, Martín Costa, Ermiya Farokhnejad: “Fully Dynamic k -Median with Near-Optimal Update Time and Recourse”, CoRR, Vol. abs/2411.03121, 2024.

URL <https://doi.org/10.48550/ARXIV.2411.03121>

In metric k -clustering, we are given as input a set of n points in a general metric space, and we have to pick k centers and cluster the input points around these chosen centers, so as to minimize an appropriate objective function. In recent years, significant effort has been devoted to the study of metric k -clustering problems in a *dynamic setting*, where the input keeps changing via updates (point insertions/deletions), and we have to maintain a good clustering throughout these updates [Fichtenberger, Lattanzi, Norouzi-Fard and Svensson, SODA’21; Bateni, Esfandiari, Fichtenberger, Henzinger, Jayaram, Mirrokni and Weise, SODA’23; Lacki, Haeupler, Grunau, Rozhon and Jayaram, SODA’24; Bhattacharya, Costa, Garg, Lattanzi and Parotsidis, FOCS’24; Forster and Skarlatos, SODA’25]. The performance of such a dynamic algorithm is measured in terms of three parameters: (i) Approximation ratio, which signifies the quality of the maintained solution, (ii) Recourse, which signifies how stable the maintained solution is, and (iii) Update time, which signifies the efficiency of the algorithm.

We consider a textbook metric k -clustering problem, metric k -median, where the objective is the sum of the distances of the points to their nearest centers. We design the first dynamic algorithm for this problem with near-optimal guarantees across all three performance measures (up to a constant factor in approximation ratio, and polylogarithmic factors in recourse and update time). Specifically, we obtain a $O(1)$ -approximation algorithm for dynamic metric k -median with $\tilde{O}(1)$ recourse and $\tilde{O}(k)$ update time. Prior to our work, the state-of-the-art here was the recent result of [Bhattacharya, Costa, Garg, Lattanzi and Parotsidis, FOCS’24], who obtained $O(\epsilon^{-1})$ -approximation ratio with $\tilde{O}(k^\epsilon)$ recourse and $\tilde{O}(k^{1+\epsilon})$ update time.

We achieve our results by carefully synthesizing the concept of robust centers introduced in [Fichtenberger, Lattanzi, Norouzi-Fard and Svensson, SODA’21] along with the randomized local search subroutine from [Bhattacharya, Costa, Garg, Lattanzi and Parotsidis, FOCS’24], in addition to several key technical insights of our own.

3.2 Deterministic expander routing: faster and more versatile

Yi-Jun Chang (National University of Singapore, SG)

License © Creative Commons BY 4.0 International license
© Yi-Jun Chang

Joint work of Yi-Jun Chang, Shang-En Huang, Hsin-Hao Su

Main reference Yi-Jun Chang, Shang-En Huang, Hsin-Hao Su: “Deterministic Expander Routing: Faster and More Versatile”, in Proc. of the 43rd ACM Symposium on Principles of Distributed Computing, PODC 2024, Nantes, France, June 17-21, 2024, pp. 194–204, ACM, 2024.

URL <https://doi.org/10.1145/3662158.3662797>

In this talk, I will present two improvements in deterministic expander routing: (1) an improved round complexity upper bound that nearly matches the state-of-the-art for randomized algorithms, and (2) a smooth tradeoff between preprocessing and routing time, leading to improved upper bounds in distributed subgraph finding. Additionally, as a side result of independent interest, we establish the equivalence between expander routing and sorting, showing that they are reducible to each other up to a polylogarithmic factor in round complexity within the CONGEST model.

3.3 Fully Dynamic k -Clustering with Fast Update Time and Small Recourse

Martín Costa (University of Warwick – Coventry, GB)

License © Creative Commons BY 4.0 International license
© Martín Costa

Joint work of Sayan Bhattacharya, Martín Costa, Naveen Garg, Silvio Lattanzi, Nikos Parotsidis

Main reference Sayan Bhattacharya, Martín Costa, Naveen Garg, Silvio Lattanzi, Nikos Parotsidis: “Fully Dynamic k -Clustering with Fast Update Time and Small Recourse”, in Proc. of the 65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024, pp. 216–227, IEEE, 2024.


URL <https://doi.org/10.1109/FOCS61266.2024.00023>

In the dynamic metric k -median problem, we wish to maintain a set of k centers $S \subseteq V$ in an input metric space (V, d) that gets updated via point insertions/deletions, so as to minimize the objective $\sum_{x \in V} \min_{y \in S} d(x, y)$. The quality of a dynamic algorithm is measured in terms of its approximation ratio, “recourse” (the number of changes in S per update) and “update time” (the time it takes to handle an update). The ultimate goal in this line of research is to obtain a dynamic $O(1)$ approximation algorithm with $\tilde{O}(1)$ recourse and $\tilde{O}(k)$ update time.

We come arbitrarily close to resolving the main open question on this topic by developing a new framework of randomized local search that is suitable for adaptation in a dynamic setting. For every $\epsilon > 0$, this gives us a dynamic k -median algorithm with $O(1/\epsilon)$ approximation ratio, $\tilde{O}(k^\epsilon)$ recourse and $\tilde{O}(k^{1+\epsilon})$ update time. This framework also generalizes to dynamic k -clustering with ℓ^p -norm objectives, giving similar bounds for the dynamic k -means and a new trade-off for dynamic k -center.

3.4 Tree-Packing Revisited: Faster Fully Dynamic Min-Cut and Arboricity

Tijn de Vos (Paris Lodron Universität Salzburg, AT)

License  Creative Commons BY 4.0 International license

© Tijn de Vos

Joint work of Tijn de Vos, Aleksander B.G. Christiansen

Main reference Tijn de Vos, Aleksander B. G. Christiansen: “Tree-Packing Revisited: Faster Fully Dynamic Min-Cut and Arboricity”, in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12–15, 2025, pp. 700–749, SIAM, 2025.

URL <https://doi.org/10.1137/1.9781611978322.21>

A tree-packing is a collection of spanning trees of a graph. It has been a useful tool for computing the minimum cut in static, dynamic, and distributed settings. In particular, [Thorup, Comb. 2007] used them to obtain his dynamic min-cut algorithm with $\tilde{O}(\lambda^{14.5}\sqrt{n})$ worst-case update time. We reexamine this relationship, showing that we need to maintain fewer spanning trees for such a result; we show that we only need to pack $\Theta(\lambda^3 \log m)$ greedy trees to guarantee a 1-respecting cut or a trivial cut in some contracted graph.


Based on this structural result, we then provide a deterministic algorithm for fully dynamic exact min-cut, that has $\tilde{O}(\lambda^{5.5}\sqrt{n})$ worst-case update time, for min-cut value bounded by λ . In particular, this also leads to an algorithm for general fully dynamic exact min-cut with $\tilde{O}(m^{1-1/12})$ amortized update time, improving upon $\tilde{O}(m^{1-1/31})$ [Goranci et al., SODA 2023].

We also give the first fully dynamic algorithm that maintains a $(1 + \varepsilon)$ -approximation of the fractional arboricity – which is strictly harder than the integral arboricity. Our algorithm is deterministic and has $O(\alpha \log^6 m / \varepsilon^4)$ amortized update time, for arboricity at most α . We extend these results to a Monte Carlo algorithm with $O(\text{poly}(\log m, \varepsilon^{-1}))$ amortized update time against an adaptive adversary. Our algorithms work on multi-graphs as well.

Both result are obtained by exploring the connection between the min-cut/arboricity and (greedy) tree-packing. We investigate tree-packing in a broader sense; including a lower bound for greedy tree-packing, which – to the best of our knowledge – is the first progress on this topic since [Thorup, Comb. 2007].

3.5 Improved All-Pairs Approximate Shortest Paths in Congested Clique

Michal Dory (University of Haifa, IL)

License  Creative Commons BY 4.0 International license

© Michal Dory

Joint work of Michal Dory, Hong Duc Bui, Shashwat Chandra, Yi-Jun Chang, Dean Leitersdorf

Main reference Hong Duc Bui, Shashwat Chandra, Yi-Jun Chang, Michal Dory, Dean Leitersdorf: “Improved All-Pairs Approximate Shortest Paths in Congested Clique”, in Proc. of the 43rd ACM Symposium on Principles of Distributed Computing, PODC 2024, Nantes, France, June 17–21, 2024, pp. 391–400, ACM, 2024.

URL <https://doi.org/10.1145/3662158.3662804>

I will discuss a recent work where we show an $O(1)$ -approximation algorithm for All-Pairs Shortest Paths, that takes just $O(\log \log \log n)$ rounds in the distributed Congested Clique model, improving exponentially over the state-of-the art. I will describe the main ingredients of the algorithm, that can find applications in other settings, such as dynamic algorithms.

3.6 Distributed Graph Algorithms with Predictions

Faith Ellen (University of Toronto, CA)

License © Creative Commons BY 4.0 International license
© Faith Ellen

Algorithms with predictions are algorithms that are given extra information about a problem, which may be incorrect. The better the prediction, the more it should help for solving the problem. I will present a framework for evaluating deterministic distributed graph algorithms with predictions, some templates for transforming algorithms in synchronous message passing systems to effectively make use of predictions, and apply it to obtain good distributed algorithms with predictions for the Maximal Independent Set problem.

3.7 Expanders via Local Edge Flips

George Giakkoupis (INRIA – Rennes, FR)

License © Creative Commons BY 4.0 International license
© George Giakkoupis
Main reference George Giakkoupis: “Expanders via local edge flips in quasilinear time”, in Proc. of the STOC ’22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 – 24, 2022, pp. 64–76, ACM, 2022.
URL <https://doi.org/10.1145/3519935.3520022>

Mahlmann and Schindelhauer (2005) proposed the following simple process, called flip-chain, for transforming any given connected d -regular graph into a d -regular expander. In each step, a random 3-path $abcd$ is selected, and edges ab and cd are replaced by two new edges ac and bd , provided that ac and bd do not exist already. A motivation for the study of the flip-chain arises in the design of overlay networks, where it is common practice that adjacent nodes periodically exchange random neighbors, to maintain good connectivity properties. It is known that the flip-chain converges to the uniform distribution over connected d -regular graphs, and it was conjectured that an expander graph is obtained after $O(nd \log n)$ steps, w.h.p., where n is the number of vertices. In this talk, we describe a new analysis of a natural flip-chain instantiation, which shows that starting from any connected d -regular graph, for $d = \Omega(\log n)$, an expander is obtained after $O(nd \log^2 n)$ steps, w.h.p.

3.8 Dynamic Implicit Coloring Algorithms For Low-Arboricity Graphs

Christoph Grunau (ETH Zürich, CH)

License © Creative Commons BY 4.0 International license
© Christoph Grunau
Joint work of Mohsen Ghaffari, Christoph Grunau
Main reference Mohsen Ghaffari, Christoph Grunau: “Dynamic $O(\text{arboricity})$ coloring in polylogarithmic worst-case time”, CoRR, Vol. abs/2410.19536, 2024.
URL <https://doi.org/10.48550/ARXIV.2410.19536>

A recent work by Christiansen, Nowicki, and Rotenberg [STOC’23] provides dynamic algorithms for coloring sparse graphs, concretely as a function of the graph’s arboricity α . They give two randomized algorithms: $O(\alpha \log \alpha)$ implicit coloring in $\text{poly}(\log n)$ worst-case update and query times, and $O(\min\{\alpha \log \alpha, \alpha \log \log \log n\})$ implicit coloring in $\text{poly}(\log n)$ amortized update and query times (against an oblivious adversary). We improve these results

in terms of the number of colors and the time guarantee: First, we present an extremely simple algorithm that computes an $O(\alpha)$ -implicit coloring with $\text{poly}(\log n)$ amortized update and query times. Second, and as the main technical contribution of our work, we show that the time complexity guarantee can be strengthened from amortized to worst-case. That is, we give a dynamic algorithm for implicit $O(\alpha)$ -coloring with $\text{poly}(\log n)$ worst-case update and query times (against an oblivious adversary).

3.9 Towards Scalable and Practical Batch-Dynamic Connectivity

Jakub Lacki (Google – New York, US), Adam Karczmarz (University of Warsaw, PL & IDEAS NCBR – Warsaw, PL)

License © Creative Commons BY 4.0 International license

© Jakub Lacki and Adam Karczmarz

Joint work of Jakub Lacki, Julian Shun, Adam Karczmarz, Laxman Dhulipala, Quinten De Man, Zhongqi Wang

Main reference Quinten De Man, Laxman Dhulipala, Adam Karczmarz, Jakub Lacki, Julian Shun, Zhongqi Wang: “Towards Scalable and Practical Batch-Dynamic Connectivity”, CoRR, Vol. abs/2411.11781, 2024.

URL <https://doi.org/10.48550/ARXIV.2411.11781>

We study the problem of dynamically maintaining the connected components of an undirected graph subject to edge insertions and deletions. We give the first parallel algorithm for the problem which supports updates in polylogarithmic time and uses linear space. On the empirical side, we provide the first implementation of the (sequential) cluster forest algorithm. We evaluate it empirically and find that it uses up to $19.7\times$ less space and is up to $6.2\times$ faster than the best existing baseline.

3.10 One Round Distributed Clique Listing

Quanquan C. Liu (Yale University – New Haven, US)

License © Creative Commons BY 4.0 International license

© Quanquan C. Liu

Main reference Quanquan C. Liu: “A note on improved results for one round distributed clique listing”, Inf. Process. Lett., Vol. 181, p. 106355, 2023.

URL <https://doi.org/10.1016/J.IPL.2022.106355>

In this talk, we investigate listing cliques of arbitrary sizes in bandwidth-limited, dynamic networks. The problem of detecting and listing triangles and cliques was originally studied in great detail by Bonne and Censor-Hillel [ICALP 2019]. We extend this study to dynamic graphs where more than one update may occur as well as resolve an open question posed by Bonne and Censor-Hillel [2019]. Our algorithms and results are based on some simple observations about listing triangles under various settings and we show that we can list larger cliques using such facts. Specifically, we show that our techniques can be used to solve an open problem posed in the original paper: we show that detecting and listing cliques (of any size) can be done using $O(1)$ -bandwidth after one round of communication under node insertions and node/edge deletions. We conclude with an extension of our techniques to obtain a small bandwidth 1-round algorithm for listing cliques when more than one node insertion/deletion and/or edge deletion update occurs at any time.

3.11 Smoothed Analysis of Dynamic Networks

Ami Paz (CNRS – Gif-sur-Yvette, FR)

License © Creative Commons BY 4.0 International license
© Ami Paz

Joint work of Ami Paz, Seth Gilbert, Uri Meir, Gregory Schwartzman

Smoothed analysis is a framework suggested for mediating gaps between worst-case and average-case complexities. In the case of dynamic networks, this technique aims to explain the gaps between real-world networks that function well despite being dynamic and the strong theoretical lower bounds for arbitrary networks. In the seminar, I presented our recent applications of smoothed analysis to dynamic networks. Our work studies different models of adversarial behavior and concerns problems such as information propagation and load balancing. Finally, I mentioned our preliminary results regarding applying smoothed analysis to centralized dynamic graph algorithms.

3.12 Distributed Replacement Paths and Distance Sensitivity Oracles

Vijaya Ramachandran (University of Texas – Austin, US)

License © Creative Commons BY 4.0 International license
© Vijaya Ramachandran

Joint work of Vignesh Manoharan, Vijaya Ramachandran

Main reference Vignesh Manoharan, Vijaya Ramachandran: “Computing Replacement Paths in the CONGEST Model”, in Proc. of the Structural Information and Communication Complexity – 31st International Colloquium, SIROCCO 2024, Vietri sul Mare, Italy, May 27-29, 2024, Proceedings, Lecture Notes in Computer Science, Vol. 14662, pp. 420–437, Springer, 2024.

URL https://doi.org/10.1007/978-3-031-60603-8_23

Main reference Vignesh Manoharan, Vijaya Ramachandran: “Distributed Distance Sensitivity Oracles”, CoRR, Vol. abs/2411.13728, 2024.

URL <https://doi.org/10.48550/ARXIV.2411.13728>

We consider the problem of computing shortest paths in the presence of an edge failure in a graph in a distributed setting. In the sequential context, two versions of this problem have been studied widely: Replacement Paths (RP), where a specific source-destination pair is given, and we need to compute a shortest path in the graph when any one edge fails, and Distance Sensitivity Oracles (DSO), where we need to preprocess the graph in order to efficiently return shortest path information for any source-destination pair, and any failed edge. We present several upper and lower bounds for both of these problems in the distributed CONGEST model, and we pose several open problems for further investigation, especially for DSO.

Details of this work are presented in the following two write-ups:

1. Vignesh Manoharan, Vijaya Ramachandran, “Computing replacement paths in the CONGEST model,” SIROCCO 2024, LNCS 14662, pp. 420–437, 2024. https://doi.org/10.1007/978-3-031-60603-8_23
2. Vignesh Manoharan, Vijaya Ramachandran, “Distributed distance sensitivity oracles,” arXiv:2411.13728, November 2024, SIROCCO 2025, to appear.

3.13 Vizing's Theorem in Near-Linear Time

Shay Solomon (Tel Aviv University, IL)

License © Creative Commons BY 4.0 International license
© Shay Solomon

Joint work of Sepehr Assadi, Soheil Behnezhad, Sayan Bhattacharya, Martín Costa, Shay Solomon, Tianyi Zhang
Main reference Sepehr Assadi, Soheil Behnezhad, Sayan Bhattacharya, Martín Costa, Shay Solomon, Tianyi Zhang: “Vizing’s Theorem in Near-Linear Time”, CoRR, Vol. abs/2410.05240, 2024.
URL <https://doi.org/10.48550/ARXIV.2410.05240>

Vizing’s theorem states that any n -vertex m -edge graph of maximum degree Δ can be edge colored using at most $\Delta + 1$ different colors [Vizing, 1964]. Vizing’s original proof is algorithmic and shows that such an edge coloring can be found in $O(mn)$ time. This was subsequently improved to $\tilde{O}(m\sqrt{n})$ time, independently by [Arjomandi, 1982] and by [Gabow et al., 1985].

Very recently, independently and concurrently, using randomization, this runtime bound was further improved to $\tilde{O}(n^2)$ by [Assadi, 2024] and $\tilde{O}(mn^{1/3})$ by [Bhattacharya, Carmon, Costa, Solomon and Zhang, 2024] (and subsequently to $\tilde{O}(mn^{1/4})$ time by [Bhattacharya, Costa, Solomon and Zhang, 2024]).

In this talk, we present a randomized algorithm that computes a $(\Delta + 1)$ -edge coloring in near-linear time – in fact, only $O(m \log \Delta)$ time – with high probability, *giving a near-optimal algorithm for this fundamental problem.*

3.14 Locality in distributed, dynamic, online, and quantum settings

Jukka Suomela (Aalto University, FI)

License © Creative Commons BY 4.0 International license
© Jukka Suomela

I will talk about new, surprising connections between dynamic, online, distributed, and quantum worlds. In ICALP 2023 (<https://arxiv.org/abs/2109.06593>), we introduced a new model of computing, “online-LOCAL”, which gives a unified perspective for discussing locality in online, dynamic, and sequential graph algorithms. Recently we have discovered that in the context of local graph problems, the randomized version of online-LOCAL is strong enough to capture so-called non-signaling distributions, and as a corollary it is also strong enough to capture distributed quantum algorithms (<https://arxiv.org/abs/2403.01903>). In our very recent work (<https://arxiv.org/abs/2409.13795>) we demonstrate the power of this new perspective: if we show that a family of problems has the same locality in classical deterministic LOCAL and randomized online-LOCAL, as a corollary we also get a full characterization of locality in e.g. dynamic and quantum settings.

3.15 Fully Dynamic Exact SSSP

Jan van den Brand (Georgia Institute of Technology – Atlanta, US)

License © Creative Commons BY 4.0 International license
© Jan van den Brand

Joint work of Adam Karcmarz, Jan van den Brand

Main reference Jan van den Brand, Adam Karcmarz: “Deterministic Fully Dynamic SSSP and More”, in Proc. of the 64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023, pp. 2312–2321, IEEE, 2023.

URL <https://doi.org/10.1109/FOCS57990.2023.00142>

In this talk, I outline recent progress on dynamic single source shortest paths. The result is based on algebraic methods for maintaining distance information for short paths.

We present the first non-trivial fully dynamic algorithm maintaining exact single-source distances in unweighted graphs. This resolves an open problem stated by Sankowski [CO-COON 2005] and van den Brand and Nanongkai [FOCS 2019]. Previous fully dynamic single-source distances data structures were all approximate, but so far, non-trivial dynamic algorithms for the exact setting could only be ruled out for polynomially weighted graphs (Abboud and Vassilevska Williams, [FOCS 2014]). The exact unweighted case remained the main case for which neither a subquadratic dynamic algorithm nor a quadratic lower bound was known.

3.16 Dynamic Rounding: Success for Matching; Open Questions Beyond

David Wajc (Technion – Haifa, IL)

License © Creative Commons BY 4.0 International license
© David Wajc

The relax-and-round framework is one of the most ubiquitous in approximation algorithms. In the active area of dynamic (approximation) algorithms, it has mostly played a key role in advances in dynamic matching. This is explained by a gap in our understanding of the rounding step of this framework in dynamic settings: While numerous general approaches have been used to design dynamic fractional algorithms (solving the relaxation), including the primal-dual method, the multiplicative weights method, congestion balancing and continuous optimization methods, the scope of general dynamic rounding approaches is still limited mostly to (unweighted) matching problems.

In this talk, I will discuss some techniques and ideas relating to dynamic rounding of fractional matchings, based on [ACCSW ICALP18, W STOC20, BK ICALP21, K ITCS22, BKSWSODA23, BKSWS TOC24, Dudeja ICALP24, Dudeja arXiv:2402.03068, CST SODA25, BCDLST SODA25]. Time permitting, I will highlight some techniques and tools which should prove useful in designing a more general dynamic rounding toolkit.

3.17 Faster $(\Delta + 1)$ -Edge Coloring: Breaking the $m\sqrt{n}$ Time Barrier

Tianyi Zhang (ETH Zürich, CH)

License © Creative Commons BY 4.0 International license
© Tianyi Zhang

Joint work of Sayan Bhattacharya, Din Carmon, Martín Costa, Shay Solomon, Tianyi Zhang
Main reference Sayan Bhattacharya, Din Carmon, Martín Costa, Shay Solomon, Tianyi Zhang: “Faster $(\Delta+1)$ -Edge Coloring: Breaking the $m\sqrt{n}$ Time Barrier”, in Proc. of the 65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024, pp. 2186–2201, IEEE, 2024.

URL <https://doi.org/10.1109/FOCS61266.2024.00128>

Edge coloring is a fundamental graph algorithm problem that has been studied extensively in both distributed and dynamic settings. Vizing’s theorem states that any n -vertex m -edge graph of maximum degree Δ can be *edge colored* using at most $\Delta + 1$ different colors [Diskret. Analiz, ’64]. Vizing’s original proof is algorithmic and shows that such an edge coloring can be found in $\tilde{O}(mn)$ time. This was subsequently improved to $\tilde{O}(m\sqrt{n})$, independently by Arjomandi [1982] and by Gabow et al. [1985]. In this talk we present a new algorithm that computes such an edge coloring in $\tilde{O}(mn^{1/3})$ time, giving the first polynomial improvement for this fundamental problem in over 40 years.

4 Open problems

4.1 An Open Problem in Dynamic Edge Coloring

Martín Costa (University of Warwick – Coventry, GB)

License © Creative Commons BY 4.0 International license
© Martín Costa

URL https://martin-costa.github.io/martincosta.com/files/Dynamic_Edge_Coloring_Open_Problem.pdf

Let $G = (V, E)$ be a simple, undirected graph with n vertices, m edges and maximum degree Δ . It was shown by Vizing in 1964 that the graph G can always be edge colored with at most $\Delta + 1$ colors.


In the dynamic setting, we have a fixed set of n vertices V , and the graph G evolves over time via a sequence of edge insertions and deletions. We assume that we have some fixed value Δ and that the maximum degree of G is at most Δ at all times. Our objective is to design an algorithm that explicitly maintains an edge coloring χ of G at all times and has a small update time, which is the time it takes to update the edge coloring χ after an edge is either inserted or deleted from the graph G .

Dual et al. [SODA’19] showed how to maintain a $(1 + \epsilon)\Delta$ -edge coloring in $\tilde{O}(\text{poly}(1/\epsilon))$ update time.

A natural question is whether or not we can improve the dependence on ϵ in the update time. In particular, can we maintain a $(\Delta + \tilde{O}(\Delta^{1-\gamma}))$ -edge coloring of a dynamic graph G with $\tilde{O}(1)$ update time, for some $\gamma > 0$?

4.2 Open Problem: Low-Degree Spanning Tree

Thatchaphol Saranurak (University of Michigan – Ann Arbor, US)

License  Creative Commons BY 4.0 International license
© Thatchaphol Saranurak

Given a graph with m edges and n vertices, Furer and Raghavachari [1] showed a simple $O(mn)$ -time algorithm for computing a spanning tree, whose maximum degree is minimized within $+1$ additive factor from the optimal bound. Is there a near-linear-time algorithm for computing a spanning tree, whose maximum degree is within $O(1)$ multiplicative factor from the optimal bound.

If a $O(\text{polylog}(n))$ multiplicative factor is allowed, there are many possible approaches to obtain near-linear-time algorithms.

References

- 1 M. Furer and B. Raghavachari, “Approximating the minimum-degree Steiner tree to within one of optimal,” *Journal of Algorithms*, vol. 17, no. 3, pp. 409–423, 1994.

4.3 Open Problem: Dynamic Brooks’ Theorem

Tianyi Zhang (ETH Zürich, CH)

License  Creative Commons BY 4.0 International license
© Tianyi Zhang

Vertex coloring represents a foundational problem in graph theory. For any graph $G = (V, E)$ with maximum degree Δ , a simple greedy algorithm guarantees a proper vertex coloring using at most $\Delta + 1$ colors. A deeper characterization is provided by Brooks’ theorem [1], which states that a connected graph G admits a Δ -coloring if and only if G is neither a complete graph nor an odd cycle. For disconnected graphs, this condition applies independently to each connected component.

Efficient Δ -coloring algorithms have been extensively studied across computational models. In sequential settings, linear-time algorithms exist [2]. Parallel implementations have been developed for PRAM models [3, 4], while distributed solutions span LOCAL and CONGEST frameworks [5, 6, 7, 8]. Recent work has also addressed streaming models [9].

Notably, the dynamic setting remains an open frontier for Δ -coloring, where only $(\Delta + 1)$ -coloring has seen progress [10, 11, 12]. This gap presents an intriguing challenge: Can techniques from distributed and streaming paradigms – such as limited communication strategies or space-efficient data structures – enable efficient dynamic Δ -coloring algorithms? Bridging these domains could yield novel insights for dynamic graph algorithms.

References

- 1 Brooks, R. On colouring the nodes of a network. *Mathematical Proceedings Of The Cambridge Philosophical Society*. **37**, 194-197 (1941)
- 2 Skulrattanakulchai, S. Δ -list vertex coloring in linear time. *Information Processing Letters*. **98**, 101-106 (2006)
- 3 Karloff, H. An NC algorithm for Brooks’ theorem. *Theoretical Computer Science*. **68**, 89-103 (1989)
- 4 Hajnal, P. & Szemerédi, E. Brooks coloring in parallel. *SIAM Journal On Discrete Mathematics*. **3**, 74-80 (1990)

- 5 Panconesi, A. & Srinivasan, A. The local nature of Δ -coloring and its algorithmic applications. *Combinatorica*. **15** pp. 255-280 (1995)
- 6 Grable, D. & Panconesi, A. Fast distributed algorithms for brooks-vizing colorings. *Journal Of Algorithms*. **37**, 85-120 (2000)
- 7 Ghaffari, M., Hirvonen, J., Kuhn, F. & Maus, Y. Improved distributed delta-coloring. *Proceedings Of The 2018 ACM Symposium On Principles Of Distributed Computing*. pp. 427-436 (2018)
- 8 Fischer, M., Halldórsson, M. & Maus, Y. Fast distributed Brooks' theorem. *Proceedings Of The 2023 Annual ACM-SIAM Symposium On Discrete Algorithms (SODA)*. pp. 2567-2588 (2023)
- 9 Assadi, S., Kumar, P. & Mittal, P. Brooks' theorem in graph streams: a single-pass semi-streaming algorithm for Δ -coloring. *Proceedings Of The 54th Annual ACM SIGACT Symposium On Theory Of Computing*. pp. 234-247 (2022)
- 10 Bhattacharya, S., Chakrabarty, D., Henzinger, M. & Nanongkai, D. Dynamic algorithms for graph coloring. *Proceedings Of The Twenty-Ninth Annual ACM-SIAM Symposium On Discrete Algorithms*. pp. 1-20 (2018)
- 11 Bhattacharya, S., Grandoni, F., Kulkarni, J., Liu, Q. & Solomon, S. Fully dynamic $(\Delta + 1)$ -coloring in $O(1)$ update time. *ACM Transactions On Algorithms (TALG)*. **18**, 1-25 (2022)
- 12 Behnezhad, S., Rajaraman, R. & Wasim, O. Fully Dynamic $(\Delta + 1)$ -Coloring Against Adaptive Adversaries. *Proceedings Of The 2025 Annual ACM-SIAM Symposium On Discrete Algorithms (SODA)*. pp. 4983-5026 (2025)

Participants

- Sepehr Assadi
University of Waterloo, CA
- Alkida Balliu
Gran Sasso Science Institute –
L'Aquila, IT
- Aaron Bernstein
New York University, US
- Sayan Bhattacharya
University of Warwick –
Coventry, GB
- Joakim Blikstad
KTH Royal Institute of
Technology – Stockholm, SE
- Sebastian Brandt
CISPA – Saarbrücken, DE
- Karl Bringmann
Universität des Saarlandes –
Saarbrücken, DE
- Keren Censor-Hillel
Technion – Haifa, IL
- Yi-Jun Chang
National University of
Singapore, SG
- Shiri Chechik
Tel Aviv University, IL
- Keerti Choudhary
Indian Institute of Technology –
New Delhi, IN
- Martín Costa
University of Warwick –
Coventry, GB
- Tijn de Vos
Paris Lodron Universität
Salzburg, AT
- Michal Dory
University of Haifa, IL
- Aditi Dudeja
Paris Lodron Universität
Salzburg, AT
- Faith Ellen
University of Toronto, CA
- George Giakkoupis
INRIA – Rennes, FR
- Seth Gilbert
National University of
Singapore, SG
- Christoph Grunau
ETH Zürich, CH
- Magnús M. Halldórsson
Reykjavik University, IS
- Adam Karczmarz
University of Warsaw, PL &
IDEAS NCBR – Warsaw, PL
- Fabian Daniel Kuhn
Universität Freiburg, DE
- Jakub Lacki
Google – New York, US
- Quanquan C. Liu
Yale University – New Haven, US
- Yannic Maus
TU Graz, AT
- Danupon Nanongkai
MPI für Informatik –
Saarbrücken, DE
- Yasamin Nazari
VU Amsterdam, NL
- Alexandre Nolin
CISPA – Saarbrücken, DE
- Krzysztof Nowicki
Wrocław, PL
- Dennis Olivetti
Gran Sasso Science Institute –
L'Aquila, IT
- Ami Paz
CNRS – Gif-sur-Yvette, FR
- Richard Peng
Carnegie Mellon University –
Pittsburgh, US
- Vijaya Ramachandran
University of Texas – Austin, US
- Eva Rotenberg
Technical University of Denmark
– Lyngby, DK
- Thatchaphol Saranurak
University of Michigan –
Ann Arbor, US
- Shay Solomon
Tel Aviv University, IL
- Jukka Suomela
Aalto University, FI
- Jan van den Brand
Georgia Institute of Technology –
Atlanta, US
- Virginia Vassilevska Williams
MIT – Cambridge, US
- David Wajc
Technion – Haifa, IL
- Tianyi Zhang
ETH Zürich, CH
- Anna Zych-Pawlewicz
University of Warsaw, PL

