ISSN 2186-7437

NII Shonan Meeting Report

No.162

Distributed Graph Algorithms

Guy Even
Gregory Schwartzman

November 3–7, 2019

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Overview of the Meeting

The field of distributed graph algorithms is a rapidly growing field in Theoretical Computer Science. The main goals in this field are to develop algorithms and methodologies that will support the correct and efficient operation of decentralized systems. The correctness, functionality, and efficiency of many decentralized systems rely on distributed algorithms over graphs. The classical example is the Internet. Newer applications include blockchains (e.g. Bitcoin), communication networks between data-centers, autonomous vehicles, and modern cellular networks (e.g. 5G).

A distributed system is modeled as a communication graph, where the nodes of the graph are the processors connected to each other via communication links, and communication occurs in synchronous rounds. The goal is to develop fast algorithms for fundamental graph problems, where the input graph is the topology of the network. Efficiency is measured in terms of the number of communication rounds of the algorithm. This abstraction formalizes the empirical observation that, in many modern distributed systems, the main bottleneck is communication between nodes rather than local computation.

The field of distributed graph algorithms is rather active, keeps growing in popularity year by year (there are two main annual conferences and a few smaller conferences). This field builds on the classical field of graph algorithms while taking into account the constraints imposed in a distributed environment. Many new algorithms and algorithmic techniques have been developed recently. These recent results are promising candidates for becoming the foundational basis of this field.

The meeting brought together researchers in this dynamic field so that they could share their recent results and insights. New ideas and collaborations have resulted from the meeting. The participants of the meeting came from various countries and institutions, specifically from Israel, Japan, Europe and the USA. Following the Dagstuhl tradition of welcoming researchers in early stages of their career, we have invited a few PhD students to the meeting. We believe the meeting strengthened the academic connection between this international research communities, especially between Japanese and non-Japanese researchers.

The participants came from various subfields within the field of distributed graph algorithms. The research interests of our invitees range from solving fundamental graph problems in the distributed setting, to modeling swarms of autonomous robots in 3D space. We that our meeting, with this diverse mix of researchers, had a very positive impact on the field, as ideas and techniques from one subfield can be applied to related subfield.

The main research topics that were discussed in the meeting are:

Bridging the gap between lower and upper bounds. Designing algorithms with an optimal running time has both theoretical and practical importance. Such a rare feat demonstrates a deep understanding of the computational character of a problem. Recently, there has been an increasing number of results in the field of distributed graph algorithms achieving optimal running times that meet the lower bounds. Nevertheless, some fundamental problems still exhibit a gap between their lower and upper bounds. The most famous such open problem is that of computing a Maximal Independent Set (MIS) in graphs. The MIS problem is a fundamental tool for symmetry breaking in distributed computing, and any improvement in this problem would immediately translate into an improved running time for a range of important distributed algorithms. MIS is only an example to one of many fundamental algorithmic problems (Set Cover, Spanning Tree, Dominating Set, Shortest Path computation, etc...) we wish to consider.

The power of randomization in distributed computing. Randomization is a fundamental tool in algorithm design. In fact, randomized algorithms for some problems are faster than the best known deterministic algorithms. This raises the open question of whether randomization is absolutely needed in order to achieve faster running times. Thus, the possibility of closing the gap in distributed computing between deterministic and randomized upper bounds for many problems is a long standing open problem. For example, the MIS problem admits a logarithmic time randomized algorithm, but no sublinear time deterministic algorithm is known. Understanding the power of randomization in the design of distributed algorithms is of fundamental importance.

Harnessing algorithmic techniques for distributed algorithms. Many techniques have been developed in the field of Algorithms. These techniques are useful in many models such as approximation algorithms of NP-hard problems, online algorithms that make decisions without knowledge of the future, streaming algorithms that read the input once and have limited memory, dynamic algorithms in which the input changes over time, and distributed algorithms that are characterized by a limited local view of the "world". Techniques such as Linear Programming, Primal-Dual algorithms, Randomization, Dynamic Programming, etc. have been proven to be useful in the development of algorithms for all of these settings.

Fundamental algorithmic building blocks for applications. Each application has its particular characteristics that influence the applicability of distributed algorithmic solutions. Some examples of these special characteristics include: Coordination and synchronization of data centers is characterized by having huge traffic volumes over modest sized networks of tens or hundreds of nodes that are placed far away from each other. Vehicular communication networks, on the other hand, consist of vehicles that are close to each other and require relatively small amounts of communication. However, delay is a crucial factor in vehicular communication networks to obtain safety. Another example is Blockchain networks that consist of thousands of nodes with a dense communication pattern. Indeed, theoretical problems such as reaching consensus in a distributed fashion in the presence of faults turn out to be fundamental problems in the design of Blockchain systems.

Meeting Schedule

Check-in day (November 3rd)

- Welcome banquet

Day 1 (November 4th)

- Invited talk by Seth Pettie
- Lecture session 1
- Group photo shooting
- Lecture session 2
- Lecture session 3
- Free discussion

Day 2 (November 5th)

- Invited talk by Václav Rozhoň
- Lecture session 1
- Lecture session 2
- Quantum computing tutorial
- Free discussion

Day 3 (November 6th)

- Invited talk by Alkida Balliu
- Lecture session
- Excursion
- Main banquet

Day 4 (November 7th)

- Invited talk by Fabian Kuhn
- Lecture session
- Wrap up

Overview of Talks

Survey of the Distributed Lovasz Local Lemma

Seth Pettie, University of Michigan

The Lovasz Local Lemma is a well known tool to prove the existence of combinatorial objects (such as graph colorings or packet-routing schedules) and algorithmic versions of the LLL can be applied to efficiently find those objects. The LLL and its variants are fairly well understood in the usual model of sequential computation.

In this talk I will define the Distributed LLL problem and survey its role in distributed computing. In short, the Distributed LLL problem is a useful prism to view many recent developments in the LOCAL model, such as graph shattering and the role of randomness, modern derandomization techniques, and the existence of complexity gaps in the LOCAL time hierarchy.

Topological Perspective on Local Computing in Networks

Pierre Fraigniaud, CNRS and Universite de Paris

More than two decades ago, combinatorial topology was shown to be useful for analyzing distributed fault-tolerant algorithms in shared memory systems and in message passing systems. In this work, we show that combinatorial topology can also be useful for analyzing distributed algorithms in networks of arbitrary structure. To illustrate this, we analyze consensus, set- agreement, and approximate agreement in networks, and derive lower bounds for these problems under classical computational settings, such as the LOCAL model and dynamic networks.

Distance-2 Coloring in CONGEST

Magnus Halldorsson, Reykjavik University

We discuss some very recent results on randomized and distributed CONGEST algorithm for distance-2 coloring

Optimal Distributed Covering Algorithms

Guy Even, Tel Aviv University

We present a time-optimal deterministic distributed algorithm for approximating a minimum weight vertex cover in hypergraphs of rank f. This problem is equivalent to the Minimum Weight Set Cover problem in which the frequency of every element is bounded by f. The approximation factor of our algorithm is $(f + \epsilon)$. Let Δ denote the maximum degree in the hypergraph. Our algorithm runs in the CONGEST model and requires $O(\log \Delta / \log \log \Delta)$ rounds, for constants $\epsilon \in (0, 1]$ and $f \in N^+$. This is the first distributed algorithm for this problem whose running time does not depend on the vertex weights nor the number of vertices. Thus adding another member to the exclusive family of *provably optimal* distributed algorithms.

Joint work with Ran Ben-Basat, Ken-ichi Kawarabayashi, Gregory Schwartzman

Parameterized Distributed Algorithms

Gregory Schwartzman, NII

In this work, we initiate a thorough study of graph optimization problems parameterized by the *output size* in the distributed setting. In such a problem, an algorithm decides whether a solution of size bounded by k exists and if so, it finds one. We study fundamental problems, including Minimum Vertex Cover , Maximum Independent Set , Maximum Matching , and many others, in both the LOCAL and CONGEST distributed computation models.

We present lower bounds for the round complexity of solving parameterized problems in both models, together with optimal and near-optimal upper bounds.

Joint work with Ran Ben-Basat and Ken-ichi Kawarabayashi

Quantum Distributed Graph Algorithms

Francois Le Gall, Kyoto University

The subject of this talk will be quantum distributed computing, i.e., distributed computing when the processors of the network can exchange quantum information. After describing the basics of distributed computing, I will explain a result obtained with Frédéric Magniez (arXiv:1804.02917) on quantum algorithms computing the diameter of the network in the CONGEST model. I will then present others results (arXiv:1810.10838 and arXiv:1908.11488) that show separations between the computational powers of quantum and classical distributed algorithms in the LOCAL model and the CONGEST-CLIQUE model. I will conclude my talk by mentioning interesting and important open questions in quantum distributed computing.

Deterministic Distributed Coloring (Through Recursive List Coloring)

Fabian Kuhn, University of Freiburg

We describe deterministic distributed coloring algorithms to compute a $(\Delta+1)$-vertex coloring in time $2^{O(\sqrt{log\Delta})} \cdot \log n$ and a $(2\Delta-1)$-edge coloring in time $2^{O(\sqrt{log\Delta})} + O(\log^* n)$. The main technical contribution leading to these algorithms is a recursive distributed algorithm to solve list coloring problems with lists of size $\Delta^{(1+o(1))}$. Given some list coloring problem and an orientation of the edges, we show how to recursively divide the global color space into smaller subspaces, assign one of the subspaces to each node of the graph, and compute a new edge orientation such that for each node, the list size to degree ratio degrades at most by a constant factor on each recursion level.

The landscape of distributed time complexity

Dennis Olivetti, Aalto University

Locally checkable labelings (LCLs) are a class of graph problems in which a globally correct solution can be verified locally, by checking all constant-radius neighbourhoods. Many important graph problems fall into this class, such as vertex coloring, maximal independent set, maximal matching, etc.

How much do we need to know about the graph in order to solve an LCL problem? How much does randomness help? What are the possible time complexity classes of LCLs? These are some of the most fundamental questions in theory of distributed computing.

In this context, there are several relevant graph classes one can study; we focus on bounded-degree graphs, and the model we consider is the LOCAL model of distributed computing. This presentation will go through the state of the art of the landscape of all possible distributed time complexities of LCLs, both for deterministic and randomised algorithms.

Distributed Algorithms for LP-Type Problems

Christian Scheideler, Paderborn University

In this talk I will present distributed algorithms for LP-type problems in the well-known gossip model. LP-type problems include many important classes of problems such as (integer) linear programming, geometric problems like smallest enclosing ball and polytope distance, and set problems like hitting set and set cover. In the gossip model, a node can only push information to or pull information from nodes chosen uniformly at random. Protocols for the gossip model are usually very practical due to their fast convergence, their simplicity, and their stability under stress and disruptions. Our algorithms are very efficient (logarithmic rounds or better with just polylogarithmic communication work per node per round) whenever the combinatorial dimension of the given LP-type problem is constant, even if the size of the given LP-type problem is polynomially large in the number of nodes.

Distributed Computation of MIS in Hypergraphs

Chaodong Zheng, Nanjing University

Given a graph, a maximal independent set (MIS) is a maximal subset of pairwise non-adjacent vertices. Finding an MIS is a fundamental problem in distributed computing. Although the problem is extensively studied and well understood in simple graphs, our knowledge is still quite limited when solving it in hypergraphs. In this talk, I will present some known previous results on efficient distributed computation of MIS in hypergraphs. I will also present a randomized algorithm for computing an MIS in a restricted class of hypergraphs, namely, linear hypergraphs. This algorithm has poly-logarithmic runtime and works in the CONGEST model.

Trade-offs in Distributed Interactive Proofs

Ami Paz, University of Vienna

The study of interactive proofs in the context of distributed network computing is a novel topic, recently introduced by Kol, Oshman, and Saxena [PODC 2018]. In the spirit of sequential interactive proofs theory, we study the power of distributed interactive proofs. This is achieved via a series of results establishing trade-offs between various parameters impacting the power of interactive proofs, including the number of interactions, the certificate size, the communication complexity, and the form of randomness used. Our results also connect distributed interactive proofs with the established field of distributed verification. In general, our results contribute to providing structure to the landscape of distributed interactive proofs.

Lower bounds for maximal matchings and maximal independent sets

Alkida Balliu, Aalto University

There are distributed graph algorithms for finding maximal matchings and maximal independent sets in $O(\Delta + \log^* n)$ communication rounds; here n is the number of nodes and Δ is the maximum degree. The lower bound by Linial (1987, 1992) shows that the dependency on n is optimal: these problems cannot be solved in $o(\log^* n)$ rounds even if $\Delta = 2$.

However, the dependency on Δ is a long-standing open question, and there is currently an exponential gap between the upper and lower bounds.

We prove that the upper bounds are tight. We show that maximal matchings and maximal independent sets cannot be found in $o(\Delta + \log\log n/\log\log\log n)$ rounds with any randomized algorithm in the LOCAL model of distributed computing.

As a corollary, it follows that there is no deterministic algorithm for maximal matchings or maximal independent sets that runs in $o(\Delta + \log n/\log\log n)$ rounds; this is an improvement over prior lower bounds also as a function of n.

Distributed derandomization via network decomposition

Václav Rozhoň, ETH Zurich

We present a simple polylogarithmic-time deterministic distributed algorithm for network decomposition. This improves on a celebrated $2^{O(\sqrt{\log n})}$-time algorithm of Panconesi and Srinivasan [STOC'93] and settles one of the long-standing and central questions in distributed graph algorithms. It also leads to the first polylogarithmic-time deterministic distributed algorithms for numerous other graph problems, hence resolving several open problems, including Linial's well-known question about the deterministic complexity of maximal independent set [FOCS'87]. Put together with the results of Ghaffari, Kuhn, and Maus [STOC'17] and Ghaffari, Harris, and Kuhn [FOCS'18], we get a general distributed derandomization result that implies P-RLOCAL = P-LOCAL. That is, for any distributed problem whose solution can be checked in polylogarithmic-time, any polylogarithmic-time randomized algorithm can be derandomized to a polylogarithmic-time deterministic algorithm. By known connections, our result leads also to substantially faster randomized algorithms for a number of fundamental problems including $(\Delta + 1)$-coloring, MIS, and Lovász Local Lemma.

Distributed Uniformity Testing

Uri Meir, Tel Aviv University

In uniformity testing, we are given an oracle access to a distribution over the set [n], and parameter , and we are asked to distinguish with high probability between the case that the distribution is uniform, and the case that it is ϵ-far from uniform (measured in statistical distance). Our goal in this setting is to minimize the number of samples to know the answer with high enough probability.

One can define a parallelized version of the problem, where k players are given samples from the same distribution, each sends a single-bit answer to a referee, and the referee decides by aggregating all their individual "decisions" together. Our complexity measure is now the number of samples per player.

We show that if the aggregation function is AND (we only accept if everyone sent 'yes'), one can devise such a tester, but the number of samples each player would need is almost enough to solve the problem on its own. (a later work shows that you cannot do much better than that) On a more optimistic note, for a general aggregation rule (a threshold rule would suffice), there exist a testing algorithm where each players only takes a fraction of $1/\sqrt{k}$ from the sample complexity of the centralized version. (later work shows that this is optimal, up to constants)

Recent Developments on Distributed Min-Cut Problem

Mohit Daga, KTH-Royal Institute of Technology

In this talk, I would give the highlights of recent developments on distributed min-cut problem in the CONGEST model. I will also discuss on the possible directions of improvents.

On the swap game with local information

Yukiko Yamauchi, Kyushu University

In the swap game, each player is associated to a vertex of a graph and can perform an edge swap, i.e., removing an incident edge and creating a new edge. We consider players that can access strategies of other players within a specified distance and show the effect of the locality on the Price of Anarchy and the dynamics.

Leader Election in Population Protocols

Yuichi Sudo, Osaka University

This talk Introduces an state of art algorithm that solves the leader election problem in the population protocol model.

Distributed Algorithms for MCMC Sampling

Yitong Yin, Nanjing University

In this talk, I will report the recent progresses on distributed algorithms for sampling a random LCL solution. In probability theory, such distributions over constraint satisfying solutions defined by local constraints are represented as Gibbs distributions. A primary tool for sampling from Gibbs distribution is the Markov chain Monte Carlo (MCMC) method. We will talk about distributed variants of the MCMC method, and will focus on a recent result on distributed algorithm that perfectly simulates the Metropolis-Hastings algorithm (a classic MCMC sampling procedure) in the CONGEST model with ideal parallel speedup.

Distributed expander decomposition and its applications

Yi-Jun Chang, ETH Zurich

In [Chang, Pettie, Zhang, SODA 2019] we introduced a framework of designing distributed CONGEST algorithms using expander decomposition. This approach led to a near-optimal triangle listing algorithm [Chang, Saranurak, PODC 2019]. In this talk I will discuss this line of research.

Distributed Algorithms below the Sequential Greedy Regime

Yannic Maus, Technion

We present newest developments and techniques of distributed algorithms that can solve classical problems such as vertex cover, dominating set or graph coloring variants. The focus lies on algorithmic techniques that obtain better results than simple centralized greedy algorithms, e.g., we focus on $(1+\epsilon)$ approximation for vertex cover instead of 2-approximation or Delta-vertex coloring instead of $(\Delta+1)$-vertex coloring.

An Automatic Speedup Technique for the LOCAL Model

Sebastian Brandt, ETH Zurich

Recent years have seen the emergence of a new technique, called "round elimination", that has become one of the main tools in proving round complexity lower bounds for locally checkable problems in the LOCAL model of distributed computing. The fundamental idea behind this technique is that the existence of a t-round algorithm for some given problem P can be used to infer the existence of a (t-1)-round algorithm for either the same or a different problem

P_1. By applying this idea iteratively, obtaining problems $P_1, P_2, ...$, the round complexity of P can then be related to the question of which of the P_i is the first problem in the sequence that can be solved in 0 rounds. As we will see, the above sequence of problems can be obtained in an automated way, opening new and exciting possibilities for obtaining improved complexity bounds for locally checkable problems. In this talk, we will take an in-depth look at the round elimination technique, its inner workings, potential and limitations.

Dynamic Algorithms for the Massively Parallel Computation Model

Nikos Parotsidis, University of Copenhagen

The Massive Parallel Computing (MPC) model gained popularity during the last decade and it is now seen as the standard model for processing large scale data. One significant shortcoming of the model is that it assumes to work on static datasets while, in practice, real world datasets evolve continuously. To overcome this issue, in this paper we initiate the study of dynamic algorithms in the MPC model. We first discuss the main requirements for a dynamic parallel model and we show how to adapt the classic MPC model to capture them. Then we analyze the connection between classic dynamic algorithms and dynamic algorithms in the MPC model. Finally, we provide new efficient dynamic MPC algorithms for a variety of fundamental graph problems, including connectivity, minimum spanning tree and matching.

List of Participants

- Alkida Balliu, Aalto University
- Ami Paz, University of Vienna
- Bingkai Lin, Nanjing University
- Boaz Patt-Shamir, Tel Aviv University
- Chaodong Zheng, Nanjing University
- Christian Scheideler, Paderborn University
- Dennis Olivetti, Aalto University
- Fabian Kuhn, University of Freiburg
- Francois Bonnet, Tokyo Institute of Technology
- Francois Le Gall, Kyoto University
- Gregory Schwartzman, NII
- Guy Even, Tel Aviv University
- Ken-ichi Kawarabayashi, NII

- Magnus Halldorsson, Reykjavik University
- Nikos Parotsidis, University of Copenhagen
- Penghui Yao, Nanjing University
- Pierre Fraigniaud, CNRS and Universite de Paris
- Sagnik Mukhopadhyay, KTH-Royal Institute of Technology
- Sebastian Brandt, ETH Zurich
- Seth Pettie, University of Michigan
- Taisuke Izumi, Nagoya Institute of Technology
- Toshimitsu Masuzawa, Osaka University
- Yannic Maus, Technion
- Yi-Jun Chang, ETH Zurich
- Yitong Yin, Nanjing University
- Yuichi Sudo, Osaka University
- Yukiko Yamauchi, Kyushu University
- Manuela Fischer, ETH Zurich
- Mohit Daga, KTH-Royal Institute of Technology
- Uri Meir, Tel Aviv University
- Václav Rozhoň, ETH Zurich