Report from Dagstuhl Seminar 16282

Topological Methods in Distributed Computing

Edited by

Dmitry Feichtner-Kozlov[1] **and Damien Imbs**[2]

1 **Universität Bremen, DE,** `dfk@math.uni-bremen.de`
2 **Universität Bremen, DE,** `imbs@math.uni-bremen.de`

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 16282 "Topological Methods in Distributed Computing", which was attended by 22 international researchers, both junior and senior. In the last 10–15 years, there has been an increasing body of work on applications of topology in theoretical distributed computing. This seminar brought together computer scientists working in theoretical distributed computing and mathematicians working in combinatorial topology, leading to interesting new collaborations. This report gathers abstracts of the talks given by the participants and of the group research sessions that happened during the seminar.

Seminar July 10–15, 2016 – http://www.dagstuhl.de/16282
1998 ACM Subject Classification C.2.4 Distributed Systems, H.1.1 Systems and Information Theory
Keywords and phrases combinatorial topology, concurrency, distributed protocols, shared-memory communication, solvability
Digital Object Identifier 10.4230/DagRep.6.7.31

1 Executive Summary

Dmitry Feichtner-Kozlov
Damien Imbs

This seminar brought together 22 researchers in combinatorial topology and in theoretical distributed computing. Participants came from Germany, France, Israel, Switzerland, United States, Canada, Japan and Austria. The seminar featured a combination of 1-hour talks, group sessions and an open problems session.

Scientific background and topics of the seminar

In the classical sequential computational model, computability is viewed through the Church-Turing thesis, where computations are reduced to those done by Turing machines, and complexity issues are of central importance. In the distributed setting, the situation is quite different. Since the threads of executions may intertwine in various ways (depending on the model), one of the central issues becomes dealing with execution ambiguity, and deciding whether certain standard tasks (Consensus, Renaming, etc.) are computable at all.

In this sense, the distributed setting is harder to analyze rigorously than the sequential one, or at least the difficulties are of quite different type. At the same time very many

Topological Methods in Distributed Computing, *Dagstuhl Reports*, Vol. 6, Issue 7, pp. 31–41
Editors: Dmitry Feichtner-Kozlov and Damien Imbs
DAGSTUHL REPORTS Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

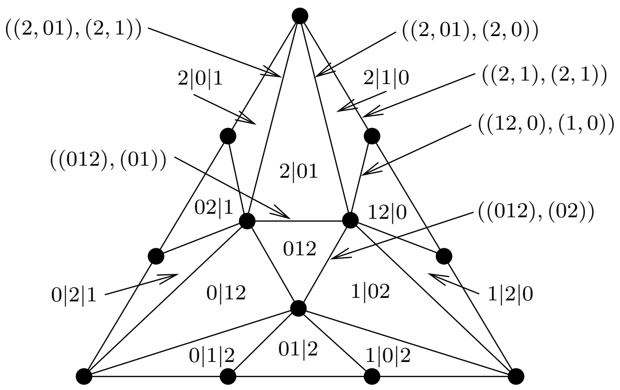

Figure 1 The protocol complex for the one-round execution of the standard immediate snapshot protocol for 3 processors.

real-life situations need to be modeled by the distributed setting. These include networks of banking machines, or networks of flight controllers and airplanes, who need to reach a common decision in a decentralized setting. Another example is the parallel chip design, where we need to understand what type of elementary operations – so-called computational primitives, have to be implemented on the hardware level, so that the resulting computational system is powerful enough for our needs.

In the 80s it was realized (due in particular to the work of Fischer, Lynch and Paterson) that certain standard tasks (Consensus) cannot be solved in standard computational models (such as Message-Passing) in the presence of even simple processor crash failures. As spectacular as it is, it has become one of the steps in the development of a sophisticated and beautiful subject of theoretical distributed computing; we refer here to the classical books of Lynch and Attiya & Welch.

In the late 90s and in the early years of our millenium, it was realized by at least 3 independent groups of researchers that topological methods are applicable in proving impossibility results in theoretical distributed computing. There followed a process of further penetration of topological methods, which by now have gained a definite foothold in distributed computing. Additionally, there has also been some work on mathematical foundations, though much remains to be done when it comes to precise definitions and rigorous proofs. Independently, we feel that it is of great interest to develop the mathematics which is inspired by these methods.

The state-of-the-art of the subject was recently summarized in a book by Herlihy, Kozlov and Rajsbaum. One of the paradigms introduced there is to replace the computational task specification by the triple: input complex, output complex, and task specification map, there the input and the output complexes are simplicial complexes with additional structure, and the task specification map is what we call a carrier map, whose definition reflects our desire to restrict ourselves to the wait-free protocols. All the wait-free tasks can be encoded this way, and as a result one obtains both well-known as well as new structures from combinatorial topology.

Furthermore, one can consider the simplicial model for the totality of all executions of a given protocol – the so-called protocol complex. In the full formal setting one actually considers a triple of two simplicial complexes and a carrier map, each one equipped with an additional structure. The intuition here is that the second simplicial complex, as well as the carrier map depend heavily on the model of computation that we choose. One standard example is to take the so-called Immediate Snapshot model. On Figure 1 we show the protocol complex for the one-round execution of the standard immediate snapshot protocol for 3 processors. As already this example shows, frequently there is a purely combinatorial description of the arising simplicial structure. The question of wait-free computability of a given task in a given computational model reduces then to the question of existence of a simplicial map from the protocol complex to the output complex, the so-called decision map, which satisfies certain conditions, which in essence mean that the outputs obtained by the protocol are valid under the task specification. Furthermore, we also have mathematical models for anonymous tasks, and anonymous protocols, as well as for colorless tasks.

As one can see, the mathematics needed for the current model is essentially that of simplicial complexes and carrier maps between them. With subsequent deepening of the theory and diversification of the considered questions, many further mathematical fields are coming in: for example, one needs to consider group actions and equivariant maps, as well as simplicial and carrier maps which satisfy other, less standard conditions. Many of the questions which arise in this setup are somewhat different from the questions classically studied in the simplicial context.

2 Table of Contents

3 Overview of Talks

3.1 The Entropy of Shared Memory

Peva Blanchard (EPFL – Lausanne, CH) and Julien Stainer (EPFL – Lausanne, CH)

We ask to what extent two processes communicating through shared memory can extract randomness from a stochastic scheduler, e.g., to generate random numbers for cryptographic applications. We introduce the quantitative notions of entropy rate and information capacity of a distributed algorithm. Whilst the entropy rate measures the Shannon information that may pass from a given scheduler to the processes executing the algorithm, the information capacity measures the optimal entropy rate over all possible schedulers.

We present a general method for computing these quantities by classifying distributed algorithms according to their pattern of shared memory accesses. We then address the issue of effectively extracting, online, this information into a meaningful format at every process. We present an algorithm solving this problem with an optimal memory consumption.

3.2 Anonymous Graph Exploration with Binoculars

Jeremie Chalopin (Aix-Marseille University, FR)

Joint work of Jeremie Chalopin, Emmanuel Godard, Antoine Naudin
Main reference J. Chalopin, E. Godard, A. Naudin, "Anonymous Graph Exploration with Binoculars", in Proc. of the 29th Int'l Symp. on Distributed Computing (DISC'15), LNCS, Vol. 9363, pp. 107–122, Springer, 2015.
URL http://dx.doi.org/10.1007/978-3-662-48653-5_8

We investigate the exploration of networks by a mobile agent. It is long known that, without global information about the graph, it is not possible to make the agent halts after the exploration except if the graph is a tree. We therefore endow the agent with binoculars, a sensing device that enables the agent to see the graph induced by its neighbors.

We show that, with binoculars, it is possible to explore and halt in a large class of non-tree networks. We give a complete characterization of the class of networks that can be explored using binoculars using standard notions of discrete topology. This class is much larger than the class of trees: it contains in particular chordal graphs, plane triangulations and triangulations of the projective plane. Our characterization is constructive: we present an Exploration algorithm that is universal; this algorithm explores any network explorable with binoculars, and never halts in non-explorable networks.

3.3 Introduction to matroids (with a view on computations in topology)

Emanuele Delucchi (University of Fribourg, CH)

Matroid theory is a structurally rich branch of combinatorics with pervasive connections with geometry, topology, and algebra.

In this expository talk I will gently introduce the notion of matroid and illustrate some of its applications and main open problems. Then – with an view towards topological applications – I will go on to more refined combinatorial structures such as oriented matroids. I will conclude with a recent and already thriving new character among matroidal theories: Baker's matroids over hyperfields.

3.4 Evasiveness

Etienne Fieux (Université Paul Sabatier – Toulouse, FR)

This is an introductory talk on evasiveness. It will not give a presentation of recent results on this question, the aim being to give an illustration of the use of tools from algebraic topology in order to solve questions in combinatorics. More precisely, we will emphasize the role of collapsibility for making a bridge between the original version of the evasiveness conjecture (about query complexity of boolean functions) and the "topological" version in terms of simplicial complexes. We also will present important results obtained by this way.

3.5 What is Distributed Computability? (Personal Perspective)

Eli Gafni (UCLA, US)

Early research in Distributed-Computing (DC) à la Dijkstra and Lamport centered on the Mutual-Exclusion (MX) problem. In this problem a processor with distinguished privilege, has to signal to other processors that it relinquishes its privilege. Hence, the "Wait-until Condition" construct naturally had to be used. They did not stop and ask whether it is essential in a one-shot case of just securing the privilege once.

In a landmark paper that started distributed-computability, Fischer, Lynch and Patterson in 1983 observed that the one-shot model of MX is of runs such that any processor that took a single step in a run will continue to take steps infinitely-often. They then showed that any meaningful relaxation of this set of runs in the form of allowing a single processor out of any number, to appear only finite number of times greater than 0, is an obstruction to the ability to secure the privilege. Thus, the power of coordination depends on the set of runs under consideration: Coordinating between processors that have limited agreement about what run they might be in. Thus, DC is the study of the coordination power of a given set of runs. Theoreticians found it convenient to consider models which are not "linear" like runs. Nevertheless, such models should be used only insofar as they correspond to some set of runs, i.e. the coordination-power of the non-linear model under study should be exactly the same as the coordination power of a set of runs.

3.6 Topological methods for distributed computability in communication networks

Emmanuel Godard (Aix-Marseille University, FR)

Coverings and fibrations are topological morphisms that capture very well the locality of distributed computations in communication networks. First introduced to generalize the seminal impossibility result of Angluin in anonymous networks, they proved to be a very versatile tool to get complete distributed computability characterizations in many models. We present here a general framework to investigate distributed computability in communication networks (that could be anonymous, homonymous or with ids) for any scheduler, both for deterministic and probabilistic algorithms. This framework comprehends all known characterisations and provides a direct and effective methodology to investigate new models.

3.7 Homotopy theory and concurrency

John Frederick Jardine (University of Western Ontario – London, CA)

Higher dimensional automata are defined as finite cubical complexes. The path category P(K) of such an object K has objects all vertices (states) of K, and its morphisms are the execution paths between states. An algorithm for computing P(K) is derived from an explicit 2-category resolution. Complexity reduction methods are discussed. A first parallel technique for computing P(K) is displayed.

3.8 Concurrency as an Iterated Affine Task

Petr Kuznetsov (Telecom Paris Tech, FR)

We consider models of computations expressed via sets of runs bounding the concurrency level: the number of processes that can be concurrently active. The model of k-concurrency is equivalent to the read-write shared-memory systems equipped with k-set-agreement objects. We show that every such model can be characterized by an affine task: a simple combinatorial structure (a simplicial complex), defined as a subset of simplices in the second degree of the standard chromatic subdivision. Our result implies the first combinatorial representation of models equipped with abstractions other than read-write registers.

3.9 Tight Bounds for Connectivity and Set Agreement in Byzantine Synchronous Systems

Hammurabi Mendes (University of Rochester, US)

Joint work of Hammurabi Mendes, Maurice Herlihy

In this talk, we discuss how the protocol complex of a Byzantine synchronous system can remain $(k-1)$-connected for up to $\rceil t/k \lceil$ rounds, where t is the maximum number of Byzantine processes, with $t \geq k \geq 1$. In crash-failure systems, the connectivity upper bound is $\rfloor t/k \lfloor$ rounds, so Byzantine ambiguity can potentially persist for one extra round, delaying the solution to k-set agreement and other related problems. We see how our connectivity bound is tight, at least when the number of processes is suitably large compared to the number of failures, by solving an appropriate formulation of k-set agreement in $\rceil t/k \lceil + 1$ rounds.

3.10 Carrier Complex: A Poset Topology for Finding Distributed Protocols

Susumu Nishimura (Kyoto University, JP)

We study the problem of finding a wait-free protocol for read-write shared memory distributed systems that meets a given topological specification of colored distributed task. Due to the asynchronous computability theorem by Herlihy and Shavit, this amounts to finding a pair of a subdivision on the input complex and a chromatic simplicial map onto the output complex, which is an undecidable problem.

Though there can be no perfect solution to the problem, we claim that there is a partial solution. Given a carrier map as a task specification, we first set up a variant of a well-known poset topology, called the carrier complex, which is a simplicial complex obtained by a combinatorial construction on the carrier map. The solution is then reached by finding an appropriate sequence of simplicial collapses that deforms the input complex embedded as a subcomplex within the carrier complex.

To put the above method to work for concrete instances of tasks in practice, we have to find a reasonably efficient algorithm. The complexity bound of the algorithm is polynomial when the number of processes is fixed, but the constant factor rapidly increases for a larger number of processes. Based on a preliminary experiment on a prototype implementation, we report that the algorithm is expected to work in practice for up to 4 processes.

3.11 Algebraic Methods in Distributed Graph Algorithms

Ami Paz (Technion – Haifa, IL)

Joint work of Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, Jukka Suomela
Main reference K. Censor-Hillel, P. Kaski, J. H. Korhonen, C. Lenzen, A. Paz, J. Suomela, "Algebraic Methods in the Congested Clique", in Proc. of the 2015 ACM Symp. on Principles of Distributed Computing (PODC'15), pp. 143–152, ACM, 2015.
URL http://dx.doi.org/10.1145/2767386.2767414

We will survey algorithm and lower bound techniques in three distributed message passing models – local, congest and congested clique. Then, we will discuss our recent results [PODC'15] regarding these problems: implementing matrix multiplication in the congested clique model, and using it to solve various distance computation and subgraph detection problems in this model.

If time suffices , we will discuss some restricted lower bounds for these problems, and the failure to achieve general lower bounds in the congested clique model.

3.12 Topological characterisation of 2-set agreement in communication networks with omission faults

Eloi Perdereau (Aix-Marseille University, FR)

We consider an arbitrary communication network G where at most f messages can be lost at each round. We give a topological characterisation when $G = K_3$ and $G = P_3$. We then show it is possible to reduce any arbitrary graph G to one of this case.

3.13 Easy Impossibility Proofs for k-Set Agreement

Ulrich Schmid (TU Wien, AT)

Despite of being quite similar agreement problems, distributed consensus (1-set agreement) and general k-set agreement require surprisingly different techniques for proving impossibilities. In particular, the relatively simple bivalence arguments used in the impossibility proof for consensus in the presence of a single crash failure are superseded by a reasoning based on algebraic topology resp. variants of Sperner's Lemma in the case of k-set agreement for $f \geq k > 1$ crash failures. In this talk, we provide an overview of a generic theorem for proving the impossibility of k-set agreement in various message passing settings, which has been published at OPODIS'11 [1]. Our BRS Theorem is based on a reduction to the consensus impossibility in a certain subsystem resulting from a partitioning argument, which facilitates easy impossibility proofs for k-set agreement. Its broad applicability is demonstrated in several message passing models.

The relevance of the BRS Theorem for this Dagstuhl Seminar results from the fact that both the core idea (reduction to consensus) and the main precondition (partitioning) of the BRS Theorem have counterparts in the algebraic topology setting. Exploring the exact nature of these is not only interesting in itself, but may also lead to a better understanding of the actual relation of partitioning-vs. Sperner-Lemma-based impossibility results.

References

1 Martin Biely, Peter Robinson and Ulrich Schmid, "Easy Impossibility Proofs for k-Set Agreement in Message Passing Systems", in Proc. of the 15th Int'l Conf. on Principles of Distributed Systems (OPODIS'11), LNCS, Vol. 7109, pp. 299–312, Springer, 2011. DOI: 10.1007/978-3-642-25873-2_21

4 Working groups and the open problems session

Four group sessions, of two and a half hours each, were organized in the afternoons. These sessions served as a complement to the lectures and provided a less formal setting in which the researchers could exchange ideas and locate areas in which progress could be made in the near future. The sessions also served as an excellent opportunity to jump-start joint projects. Due to the extensive interest and the limited amount of time available to the organizers, the group sessions were scheduled in parallel and both tracks were well attended.

One of the highlights was an extensive session on higher-dimensional expansion (by Roy Meshulam), a modern topic bringing a new approach to the classical graph expanders. Due to the technical nature of the subject and its novelty to most participants, the working session provided a very fitting setting for the discussions. Another topic which was addressed was the possible connections of higher homotopy and the study of concurrency to the distributed setting (by Rick Jardine). This later theme has already brought some fruit, and there are upcoming workshops which will continue in this new research direction.

The open problems session was also attended by everyone and resulted in lively discussions of several central problems. Just to mention one: the round-complexity of the Weak Symmetry Breaking, and the intricate mathematics connected to it, was the source of interest and fascination. It was clear to the participants, that in this, and in many other instances, much remains to be understood by the community.

We feel that the sessions were a success and made use of the unique setting provided at the Dagstuhl Research Center, versus a usual conference venue.

Participants

- Peva Blanchard
EPFL – Lausanne, CH
- Jeremie Chalopin
Aix-Marseille University, FR
- Emanuele Delucchi
University of Fribourg, CH
- Dmitry Feichtner-Kozlov
Universität Bremen, DE
- Etienne Fieux
Université Paul Sabatier – Toulouse, FR
- Eli Gafni
UCLA – Los Angeles, US
- Emmanuel Godard
Aix-Marseille University, FR
- Hans Hagen
TU Kaiserslautern, DE
- Damien Imbs
Universität Bremen, DE
- John Frederick Jardine
University of Western Ontario – London, CA
- Petr Kuznetsov
Telecom Paris Tech, FR
- Jan-Philipp Litza
Universität Bremen, DE
- Hammurabi Mendes
University of Rochester, US
- Roy Meshulam
Technion – Haifa, IL
- Susumu Nishimura
Kyoto University, JP
- Ami Paz
Technion – Haifa, IL
- Eloi Perdereau
Aix-Marseille University, FR
- Thibault Rieutord
Telecom Paris Tech, FR
- Matthieu Roy
LAAS – Toulouse, FR
- Ulrich Schmid
TU Wien, AT
- Jan Felix Senge
Universität Bremen, DE
- Julien Stainer
EPFL – Lausanne, CH