Report from Dagstuhl Seminar 23272

# Epistemic and Topological Reasoning in Distributed Systems

Armando Castañeda<sup>\*1</sup>, Hans van Ditmarsch<sup>\*2</sup>, Roman Kuznets<sup>†3</sup>, Yoram Moses<sup>\*4</sup>, and Ulrich Schmid<sup>\*5</sup>

- 1 National Autonomous University of Mexico, MX. armando.castaneda@im.unam.mx
- 2 CNRS Toulouse, FR & Université de Toulouse, FR. hans.van-ditmarsch@irit.fr
- 3 TU Wien, AT. roman@logic.at
- 4 Technion Haifa, IL. yoram.moses@gmail.com
- 5 TU Wien, AT. s@ecs.tuwien.ac.at

#### — Abstract -

This report documents the program and the outcomes of Dagstuhl Seminar 23272 "Epistemic and Topological Reasoning in Distributed Systems." The seminar brought together experts in combinatorial topology and epistemic logic interested in distributed systems, with the aim of exploring the directions that the recent interaction between those approaches can take, identifying challenges and opportunities.

Seminar July 2-7, 2023 – https://www.dagstuhl.de/23272

- **2012 ACM Subject Classification** Computing methodologies  $\rightarrow$  Distributed algorithms; Mathematics of computing  $\rightarrow$  Graph theory; Mathematics of computing  $\rightarrow$  Topology; Software and its engineering  $\rightarrow$  Distributed systems organizing principles; Software and its engineering  $\rightarrow$  Formal methods; Theory of computation  $\rightarrow$  Distributed computing models; Theory of computation  $\rightarrow$  Data structures design and analysis; Theory of computation  $\rightarrow$  Logic
- Keywords and phrases combinatorial topology, distributed systems, epistemic logic, multi-agent systems, interpreted systems, dynamic epistemic logic, simplicial semantics, knowledge-based approach, distributed computing

Digital Object Identifier 10.4230/DagRep.13.7.34

#### 1 Executive Summary

Armando Castañeda (National Autonomous University of Mexico, MX) Hans van Ditmarsch (CNRS – Toulouse, FR & Université de Toulouse, FR) Yoram Moses (Technion – Haifa, IL) Ulrich Schmid (TU Wien, AT)

Distributed services cover a wide range of our everyday activities. Examples of such services include cloud storage, cryptocurrencies, and collaborative editing, as well as concurrent software that governs modern multicore computers. Reasoning about distributed systems that provide these services is notoriously difficult, however, due to the many sources of uncertainty that can occur: varying execution speeds, unpredictable transmission delays, and

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

Epistemic and Topological Reasoning in Distributed Systems, *Dagstuhl Reports*, Vol. 13, Issue 7, pp. 34–65 Editors: Armando Castañeda, Hans van Ditmarsch, Roman Kuznets, Yoram Moses, and Ulrich Schmid Dagstuhl Reports

<sup>\*</sup> Editor / Organizer

<sup>&</sup>lt;sup>†</sup> Editorial Assistant / Collector

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

partial failures. The design and analysis of protocols and algorithms for distributed systems is, hence, a difficult and error-prone task.

Two approaches have proved to be successful in raising the level of abstraction in the modeling, design, and analysis of distributed systems: combinatorial topology and the epistemic (or knowledge-based) approach for multi-agent systems. It is also worth noting that several different flavors of topology have been successfully applied in general modal logic and related fields as well.

In the context of distributed systems, the epistemic and the topological approach have evolved fairly independently for over more than three decades. In the last few years, however, researchers have started to combine the two approaches in productive ways. In particular, this is true for both of the two main variants of epistemic reasoning in such systems, the more traditional *interpreted systems* (IS) modeling and *dynamic epistemic logics* (DEL), where approaches and/or methods from combinatorial topology have been used successfully already. This is primarily due to a duality between the Kripke models, which underlie epistemic reasoning, and simplicial complexes, which are central to the analysis of distributed protocols using combinatorial topology. And indeed, meanwhile, there are encouraging relations and options for cross-fertilization with the communities of dynamic epistemic logic, knowledge-based analysis, and topological modal logics. This development makes a new level of abstraction possible, based on mutual incorporation of the extensive results established independently by all these approaches.

#### Seminar Goals and Structure

The main purpose of this Dagstuhl Seminar was to further stimulate the process of making this integration happening: Whereas there are researchers with expertise in more than one of the above fields, a majority of them is firmly grounded in only one of those. Consequently, we brought together 28 experts from the respective scientific communities, in an attempt to:

- (A) Introduce the core topics, methods and accomplishments obtained in some field to the others.
- (B) Stimulate discussions among members from different fields, to identify possible crossfertilization opportunities.
- (C) Further explore particularly interesting/promising/challenging cross-fertilization opportunities and shape collaborations on these.

The cornerstones of the seminar program were:

- (i) 5 tutorials (90 min) by experts in the following specific areas:
  - *Alexandru Baltag* 
    - Epistemics and Topology
  - Hans van Ditmarsch

Dynamic Epistemic Logic

- Guy Goren
  - Knowledge and Action
- Jérémy Ledent

Simplicial Models: From Global States to Local States, and What Lies In-between *Sergio Rajsbaum* 

- An Overview of Combinatorial Topology and Its Distributed Computing Perspective
- (ii) 17 presentations (20 min) devoted to specific research topics, from all areas.

- (iii) 3 short presentations (5 min) in a rump session, devoted to research topics that were identified in the course of the seminar.
- (iv) Public sessions devoted to identifying particularly promising cross-fertilization opportunities.
- (v) 3 focused discussion groups, devoted to the following cross-fertilization opportunities:
  - Extending Topology to Handle Data Structures
  - Interpreted Systems and Dynamic Epistemic Logic
  - Representing Epistemic Attitudes via Simplicial Complexes
  - Multiple discussion sessions of these groups were run in parallel, with two public reporting sessions. Seminar participants could join/swap the discussion groups freely.

# 2 Table of Contents

<b>Executive Summary</b> Armando Castañeda, Hans van Ditmarsch, Yoram Moses, and Ulrich Schmid	34
Overview of Talks	
Coalgebraic Modal Logic and Fibrations Henning Basold	39
Asynchronous Speedup Theorems and FLP-Style Proofs <i>Pierre Fraigniaud</i>	40
Epistemic Analysis of the FRR ProblemKrisztina FruzsaKrisztina Fruzsa	40
Semitopology and (Paraconsistent) Logic Murdoch James Gabbay	41
Optimal Eventual Byzantine Agreement Protocols with Omission Failures Joseph Y. Halpern	42
Hypergraphs for Knowledge Roman Kniazev	42
Domain-theoretical Constraint Systems for Knowledge and Belief in Multi-agent Systems Sophia Knight	43
Impure Simplicial Complexes: Local View         Roman Kuznets	43
The Impossibility of Approximate Agreement on a Larger Class of Graphs Shihao (Jason) Liu	44
Product Updates for Partial Epistemic Models and Logical Obstruction to Task Solvability Susumu Nishimura	44
Mixing Combinatorial and Point-set Topology Thomas Nowak	45
A Speedup Theorem for Asynchronous Computation Ami Paz	45
On Simplicial Semantics Rojo Randrianomentsoa	46
Do Genes Argue? Abstract Argumentation and Boolean Networks David A. Rosenblueth	46
Logics for Data Exchange Sonja Smets	46
Synergistic Knowledge Thomas Studer	47
Pattern Models and Consensus Diego A. Velázquez	47

### Working Groups

Extending Topology to Handle Data Structures	
Henning Basold, Armando Castañeda, Faith Ellen, Guy Goren, Shihao (Jason) Liu,	
Susumu Nishimura, Thomas Nowak, Sergio Rajsbaum, Ulrich Schmid	48
Interpreted Systems and Dynamic Epistemic Logic	
Alexandru Baltag, Hans van Ditmarsch, Krisztina Fruzsa, Joseph Y. Halpern,	
Yoram Moses, Hugo Rincón Galeana, Ulrich Schmid, Sonja Smets, Diego A. Velázquez	50
Representing Epistemic Attitudes via Simplicial Complexes	
Alexandru Baltag, Henning Basold, Hans van Ditmarsch, Roman Kniazev, Sophia	
Knight, Roman Kuznets, Jérémy Ledent, David Lehnherr, Rojo Randrianomentsoa,	
Sonja Smets, Thomas Studer	53
articipants	65
	Extending Topology to Handle Data Structures Henning Basold, Armando Castañeda, Faith Ellen, Guy Goren, Shihao (Jason) Liu, Susumu Nishimura, Thomas Nowak, Sergio Rajsbaum, Ulrich Schmid Interpreted Systems and Dynamic Epistemic Logic Alexandru Baltag, Hans van Ditmarsch, Krisztina Fruzsa, Joseph Y. Halpern, Yoram Moses, Hugo Rincón Galeana, Ulrich Schmid, Sonja Smets, Diego A. Velázquez Representing Epistemic Attitudes via Simplicial Complexes Alexandru Baltag, Henning Basold, Hans van Ditmarsch, Roman Kniazev, Sophia Knight, Roman Kuznets, Jérémy Ledent, David Lehnherr, Rojo Randrianomentsoa, Sonja Smets, Thomas Studer

### **3** Overview of Talks

#### 3.1 Coalgebraic Modal Logic and Fibrations

Henning Basold (Leiden University, NL)

In this talk, I provide an overview over coalgebraic modal logic and how it can be abstractly viewed through fibrations. The idea of coalgebras is that they provide an abstraction of state-based systems and enable a category-theoretical study of behaviour, just like (universal) algebra provides tools to reason about different algebraic structures. While in algebra we are interested in structures with operations on them, such as monoids, groups, etc., in coalgebra we consider observations as the fundamental building block. In the talk, we discuss labelled transitions systems as a particular example, in which we can observe for every state the propositional atoms that are true at that state and the outgoing labelled transitions [1]. These can represent epistemic models, linear dynamical systems, Kripke frames, neighbourhood frames, descriptive frames [2], and even proofs of coinductive predicates. To reason about the behaviour induced by coalgebras, we discuss predicates that arise via so-called liftings [3] and rely on the theory of fibrations [4]. We will see the theory illustrated on the labelled box modality for labelled transition systems and we will show how a modal logic with conjunction and box can be given semantics over a coalgebra by interpreting the conjunction via the algebraic structure of meet-semilattices and the modality via a so-called distributive law that is induced by the previously defined lifting. I will present the basic principles of coalgebraic modal logic and how the approach of fibrations can be used to give semantics to modal logic in quite general coalgebras, and to prove soundness and completeness results that extend those of Kupke and Rot [5].

#### References

- 1 Lawrence S. Moss. Coalgebraic logic. Annals of Pure and Applied Logic, 96(1-3):277-317, March 1999. doi:10.1016/S0168-0072(98)00042-6.
- 2 Guram Bezhanishvili, Luca Carai, and Patrick J. Morandi. The Vietoris functor and modal operators on rings of continuous functions. *Annals of Pure and Applied Logic*, 173(1):103029, January 2022. doi:10.1016/j.apal.2021.103029.
- 3 Ichiro Hasuo, Kenta Cho, Toshiki Kataoka, and Bart Jacobs. Coinductive predicates and final sequences in a fibration. In Dexter Kozen and Michael Mislove, editors, Proceedings of the Twenty-ninth Conference on the Mathematical Foundations of Programming Semantics, MFPS XXIX, volume 298 of Electronic Notes in Theoretical Computer Science, pages 197-214. Elsevier, 2013. doi:10.1016/j.entcs.2013.09.014.
- 4 Bart Jacobs. *Categorical Logic and Type Theory*. Number 141 in Studies in Logic and the Foundations of Mathematics. Elsevier, 1999.
- 5 Clemens Kupke and Jurriaan Rot. Expressive logics for coinductive predicates. Logical Methods in Computer Science, 17(4):19:1–19:30, 2021. doi:10.46298/lmcs-17(4:19)2021.

#### 3.2 Asynchronous Speedup Theorems and FLP-Style Proofs

Pierre Fraigniaud (CNRS – Paris, FR & Université Paris Cité, FR)

License ☺ Creative Commons BY 4.0 International license © Pierre Fraigniaud Joint work of Hagit Attiya, Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum

The talk compares two generic techniques for deriving lower bounds or impossibility results in distributed computing. First, we formalize the notion of FLP-style proofs, aiming at capturing the essence of the seminal consensus-impossibility proof [1] and using forward induction. Second, we prove a speedup theorem à la [2], but for wait-free colorless algorithms, aiming at capturing the essence of the seminal round-reduction proof establishing a lower bound on the number of rounds for 3-coloring a cycle [3], and going by backward induction. We show that despite their very different natures, these two forms of proof are tightly connected. In particular, we show that for every colorless task  $\Pi$ , if there is a round-reduction proof establishing the impossibility of solving  $\Pi$  using wait-free colorless algorithms, then there is an FLP-style proof establishing the same impossibility. For 1-dimensional colorless tasks (but for an arbitrarily large number  $n \geq 2$  of processes), we prove that the two proof techniques have exactly the same power, and more importantly, both are complete: if a 1-dimensional colorless task is not wait-free solvable by  $n \geq 2$  processes, then the impossibility can be proved by both proof techniques. Moreover, a round-reduction proof can be automatically derived, and an FLP-style proof can be automatically generated from it. Finally, we illustrate the use of these two techniques by establishing the impossibility of solving any colorless covering task of arbitrary dimension by wait-free algorithms.

#### References

- Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, April 1985. doi: 10.1145/3149.214121.
- 2 Sebastian Brandt. An automatic speedup theorem for distributed problems. In PODC'19: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, pages 379–388. Association for Computing Machinery, 2019. doi:10.1145/3293611.3331611.
- 3 Nathan Linial. Locality in distributed graph algorithms. SIAM Journal on Computing, 21(1):193-201, February 1992. doi:10.1137/0221015.

#### 3.3 Epistemic Analysis of the FRR Problem

Krisztina Fruzsa (TU Wien, AT)

 License 

 © Creative Commons BY 4.0 International license
 © Krisztina Fruzsa

 Joint work of Krisztina Fruzsa, Roman Kuznets, Ulrich Schmid
 Main reference Krisztina Fruzsa, Roman Kuznets, Ulrich Schmid: "Fire!", in Proc. of the Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2021, Beijing, China, June 25-27, 2021, EPTCS, Vol. 335, pp. 139–153, Open Publishing Association, 2021.
 URL https://doi.org/10.4204/EPTCS.335.13

We present an epistemic analysis of a fundamental distributed computing problem called the Firing Rebels with Relay (FRR) within the asynchronous byzantine fault-tolerant model of distributed systems. Essentially, the FRR problem requires all correct agents to perform an action (FIRE) in an all-or-nothing fashion, however, not necessarily at the same time. By using our framework for modeling asynchronous byzantine fault-tolerant distributed systems,

we establish the necessary level of knowledge that needs to be acquired by any correct agent in order to perform FIRE in every protocol that meets the requirements of the FRR problem specification. The corresponding level of knowledge involves the so-called common eventual hope modality, which plays a crucial role in reaching (eventual) agreement among agents.

#### 3.4 Semitopology and (Paraconsistent) Logic

Murdoch James Gabbay (Heriot-Watt University – Edinburgh, GB)

```
    License 
        Creative Commons BY 4.0 International license
        Murdoch James Gabbay

    Joint work of Murdoch James Gabbay, Giuliano Losa
    Main reference Murdoch J. Gabbay, Giuliano Losa: "Semitopology: a new topological model of heterogeneous consensus", CoRR abs/2303.09287, 2023. [1]

    URL https://doi.org/10.48550/ARXIV.2303.09287
```

A semitopology is like a topology, but without the constraint that intersections of open sets be open. This models an idea in distributed consensus that participants can progress locally if they reach agreement amongst a local quorum (= open neighbourhood) of trusted peers.

Semitopologies are related to Horn clause theories: just fix some predicate atoms, write any Horn clause theory over those atoms that you like, and consider the set of assignments of truth values to the atoms that satisfy the clauses. These assignments generate the closed sets of a semitopology as their corresponding answer sets (the set of atoms to which the assignment assigns "true"). It is very easy to see that this is a semitopology.

But now we want to assert and reason about semitopological properties, by applying techniques from logic (such as derivability) to the semitopology as represented as a Horn clause theory. In essence, we have converted safety properties of distributed consensus into derivability (or non-derivability) questions about a simple logical theory.

For this, we require a paraconsistent modal logic (over three values: T, B, and F), with a modality K asserting validity in all valuations.

Time permitting, I will sketch the maths above and show the logic in action on some simple examples.

Some parts of the material described above are in a draft journal paper [1] online.

#### References

 Murdoch J. Gabbay and Giuliano Losa. Semitopology: a new topological model of heterogeneous consensus. Eprint 2303.09287, arXiv, 2023. doi:10.48550/arXiv.2303.09287.

#### 3.5 Optimal Eventual Byzantine Agreement Protocols with Omission Failures

Joseph Y. Halpern (Cornell University – Ithaca, NY, US)

License O Creative Commons BY 4.0 International license

© Joseph Y. Halpern

42

Joint work of Kaya Alpturer, Joseph Y. Halpern, Ron van der Meyden

- Main reference Kaya Alpturer, Joseph Y. Halpern, Ron van der Meyden: "Optimal Eventual Byzantine Agreement Protocols with Omission Failures", in Proc. of the 2023 ACM Symposium on Principles of Distributed Computing, PODC 2023, Orlando, FL, USA, June 19-23, 2023, pp. 244–252, ACM, 2023.
  - **URL** https://doi.org//10.1145/3583668.3594573

Work on optimal protocols for Eventual Byzantine Agreement (EBA) – protocols that, in a precise sense, decide as soon as possible in every run and guarantee that all nonfaulty agents decide on the same value – has focused on full-information protocols (FIPs), where agents repeatedly send messages that completely describe their past observations to every other agent. While it can be shown that, without loss of generality, we can take an optimal protocol to be an FIP, full information exchange is impractical to implement for many applications due to the required message size. We separate protocols into two parts, the information-exchange protocol and the action protocol, so as to be able to examine the effects of more limited information exchange. We then define a notion of optimality with respect to an information-exchange protocol. Roughly speaking, an action protocol  $\mathbf{P}$  is optimal with respect to an information-exchange protocol  $\mathcal{E}$ , if with **P**, agents decide as soon as possible among action protocols that exchange information according to  $\mathcal{E}$ . We present a knowledge-based EBA program for omission failures all of whose implementations are guaranteed to be correct and are optimal if the information exchange satisfies a certain safety condition. We then construct concrete programs that implement this knowledge-based program in two settings of interest that are shown to satisfy the safety condition. Finally, we show that a small modification of our program results in an FIP that is both optimal and efficiently implementable, settling an open problem posed by Halpern, Moses, and Waarts [1].

#### References

1 Joseph Y. Halpern, Yoram Moses, and Orli Waarts. A characterization of eventual Byzantine agreement. *SIAM Journal on Computing*, 31(3):838–865, 2001. doi:10.1137/S0097539798340217.

#### 3.6 Hypergraphs for Knowledge

Roman Kniazev (École Polytechnique – Palaiseau, FR)

License  $\textcircled{\mbox{\scriptsize only}}$  Creative Commons BY 4.0 International license

© Roman Kniazev

Main reference Éric Goubault, Roman Kniazev, Jérémy Ledent: "A Many-Sorted Epistemic Logic for Chromatic Hypergraphs", in Proc. of the 32nd EACSL Annual Conference on Computer Science Logic, CSL 2024, February 19-23, 2024, Naples, Italy, LIPIcs, Vol. 288, pp. 30:1–30:18, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
 URL https://doi.org//10.4230/LIPICS.CSL.2024.30

After recalling semantics of epistemic logic in epistemic coverings, we discuss how chromatic hypergraphs generalize them and can resolve some inconveniences appearing in practice.

#### 3.7 Domain-theoretical Constraint Systems for Knowledge and Belief in Multi-agent Systems

Sophia Knight (University of Minnesota – Duluth, MN, US)

License  $\bigcirc$  Creative Commons BY 4.0 International license

Sophia Knight

Joint work of Michell Guzmán, Stefan Haar, Sophia Knight, Catuscia Palamidessi, Prakash Panangaden, Salim Perchy, Santiago Quintero, Sergio Ramírez, Camilo Rueda, Frank Valencia

Main reference Michell Guzmán, Sophia Knight, Santiago Quintero, Sergio Ramírez, Camilo Rueda, Frank Valencia: "Reasoning About Distributed Knowledge of Groups with Infinitely Many Agents", in Proc. of the 30th International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands, LIPIcs, Vol. 140, pp. 29:1–29:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.

 ${\tt URL}\ {\tt https://doi.org//10.4230/LIPICS.CONCUR.2019.29}$ 

We begin by reviewing the connection between topology and domain theory. We discuss the importance and difficulties in representing knowledge, belief, and information in multi-agent distributed systems, especially systems with fallible agents. Then we explain how to generalize concurrent constraint programming (CCP) with an underlying epistemic constraint system, representing agents' knowledge and information. We discuss some of the details of epistemic constraint systems and the more general multi-agent spatial constraint systems. We present a new semantics for CCP extended with epistemic and spatial constraint systems and an agent space operator. We discuss how to extend these systems to infinite groups of agents. In a purely logical setting, representing group, common, and distributed knowledge of infinite groups of agents presents difficulties, but in the domain-theoretic constraint system setting, these knowledge operators can be represented in a natural way. Finally, we discuss epistemic and spatial information systems, and their correspondence to constraint systems.

#### Impure Simplicial Complexes: Local View 3.8

Roman Kuznets (TU Wien, AT)

License 
Creative Commons BY 4.0 International license © Roman Kuznets

Joint work of Hans van Ditmarsch, Roman Kuznets, Rojo Randrianomentsoa

Main reference Rojo Randrianomentsoa, Hans van Ditmarsch, Roman Kuznets: "Impure Simplicial Complexes: Complete Axiomatization", Logical Methods in Computer Science, Vol. 19(4), October 2023. URL https://doi.org//10.46298/LMCS-19(4:3)2023

We discuss the choices that the local view suggests for impure simplicial complexes, i.e., complexes where some of the agents may have crashed. We settle on a three-valued semantics with the third value being "undefined" and provide a sound and complete axiomatization by means of a novel canonical simplicial model construction.

# **3.9** The Impossibility of Approximate Agreement on a Larger Class of Graphs

Shihao (Jason) Liu (University of Toronto, CA)

44

License 

 Creative Commons BY 4.0 International license
 Shihao (Jason) Liu

 Main reference Shihao Liu: "The Impossibility of Approximate Agreement on a Larger Class of Graphs", in Proc. of the 26th International Conference on Principles of Distributed Systems, OPODIS 2022, December 13-15, 2022, Brussels, Belgium, LIPIcs, Vol. 253, pp. 22:1–22:20, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.
 URL https://doi.org//10.4230/LIPICS.OPODIS.2022.22

Approximate agreement is a variant of consensus in which processes receive input values from a domain and must output values in that domain that are sufficiently close to one another. We study the problem when the input domain is the vertex set of a connected graph. In asynchronous systems where processes communicate using shared registers, there are wait-free approximate agreement algorithms when the graph is a path or a tree, but not when the graph is a cycle of length at least 4. For many graphs, it is unknown whether a wait-free solution for approximate agreement exists.

We introduce a set of impossibility conditions and prove that approximate agreement on graphs satisfying these conditions cannot be solved in a wait-free manner. In particular, the graphs of all triangulated *d*-dimensional spheres that are not cliques satisfy these conditions. The vertices and edges of an octahedron are an example of such a graph. We also present a reduction from approximate agreement on one graph to another graph. This enables us to extend known impossibility results to even more graphs. Finally, we show that extensionbased proofs cannot be used to prove the impossibility of wait-free approximate agreement on any connected graph, which demonstrates the necessity of using combinatorial arguments.

#### 3.10 Product Updates for Partial Epistemic Models and Logical Obstruction to Task Solvability

Susumu Nishimura (Kyoto University, JP)

We introduce the notion of partial product update. This refines the original notion of product update to encompass distributed tasks and protocols modeled by impure simplicial complexes. With this refined notion, we obtain a logical method that is generalized to allow the application of logical obstruction to show unsolvability results in a distributed environment where the failure of agents is detectable. We demonstrate the use of the logical method by giving a concrete logical obstruction and showing that the consensus task is unsolvable by single-round synchronous message-passing protocol.

#### 3.11 Mixing Combinatorial and Point-set Topology

Thomas Nowak (ENS - Gif-sur-Yvette, FR)

License 

 © Creative Commons BY 4.0 International license
 © Thomas Nowak

 Joint work of Hagit Attiya, Armando Castañeda, Thomas Nowak
 Main reference Hagit Attiya, Armando Castañeda, Thomas Nowak: "Topological Characterization of Task Solvability in General Models of Computation", CoRR abs/2301.13837, 2023.
 URL https://doi.org//10.48550/ARXIV.2301.13837

The famous asynchronous computability theorem (ACT) relates the existence of an asynchronous wait-free shared-memory protocol for solving a task with the existence of a simplicial map from a subdivision of the simplicial complex representing the inputs to the simplicial complex representing the allowable outputs. The original theorem relies on a correspondence between protocols and simplicial maps in finite models of computation that induce a compact topology. This correspondence, however, is far from obvious for computational models that induce a non-compact topology, and indeed previous attempts to extend the ACT have failed.

We show first that in every non-compact model, protocols solving tasks correspond to simplicial maps that need to be continuous. This correspondence is then used to prove that the approach used in ACT that equates protocols and simplicial complexes actually works for every compact model, and to show a generalized ACT that applies also to non-compact computation models.

Our study combines combinatorial and point-set topological aspects of the executions admitted by the computational model.

#### 3.12 A Speedup Theorem for Asynchronous Computation

Ami Paz (CNRS – Gif-sur-Yvette, FR)

License Creative Commons BY 4.0 International license
 Ami Paz
 Joint work of Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum
 Main reference Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum: "A Speedup Theorem for Asynchronous Computation with Applications to Consensus and Approximate Agreement", in Proc. of the 2022 ACM Symposium on Principles of Distributed Computing, PODC 2022, Salerno, Italy, July 25 – 29, 2022, pp. 460–470, ACM, 2022.

URL https://doi.org//10.1145/3519270.3538422

Speedup theorems have recently gained increasing attention in studying distributed graph algorithms in synchronous systems. Using topological tools, we present a variant of this technique that applies to asynchronous shared-memory systems. At the core of our technique is a round reduction theorem: given a distributed task  $\Pi$ , we define a new task called the closure of  $\Pi$ , such that if  $\Pi$  is solvable in t rounds, then its closure is solvable in t - 1 rounds.

We illustrate the power of our speedup theorem by providing new proof of the wait-free impossibility of consensus using read/write registers. This is done merely by showing that the closure of consensus is consensus itself. The simplicity of our technique allows us to study additional communication objects, namely test&set and binary consensus. We analyze the approximate agreement task in systems augmented with the two objects and show that while these objects are more powerful than read/write registers from the computability perspective, they do not help to reduce the time complexity of solving approximate agreement.

#### 3.13 **On Simplicial Semantics**

Rojo Randrianomentsoa (TU Wien, AT)

License	© Creative Commons BY 4.0 International license
	© Rojo Randrianomentsoa
Joint work of	Hans van Ditmarsch, Roman Kuznets, Rojo Randrianomentsoa
Main reference	Hans van Ditmarsch, Roman Kuznets, Rojo Randrianomentsoa: "On Two- and Three-valued
	Semantics for Impure Simplicial Complexes", in Proc. of the Fourteenth International Symposium on
	Games, Automata, Logics, and Formal Verification, GandALF 2023, Udine, Italy, 18-20th September
	2023, EPTCS, Vol. 390, pp. 50–66, Open Publishing Association, 2023.
URL	https://doi.org//10.4204/EPTCS.390.4

Impure simplicial complexes are used to model distributed systems with crashing agents. We discuss language extension for three-valued semantics, and provide the translations from that semantics to the two-valued facet semantics. In my talk, I present some of our results as well as some of the challenges we face in our current work.

#### 3.14 **Do Genes Argue? Abstract Argumentation and Boolean Networks**

David A. Rosenblueth (National Autonomous University of Mexico, MX)

License 🐵 Creative Commons BY 4.0 International license David A. Rosenblueth Joint work of Eugenio Azpeitia, Stalin Muñoz, David A. Rosenblueth, Octavio Zapata

Stemming from the work by Phan Minh Dung in the 90s, there is a plethora of formalisms of abstract argumentation, from Dung's original work to abstract dialectical frameworks. Stemming from the work by François Jacob and Jacques Monod in the 60s, there are numerous mathematical models of gene regulation, from differential equations to Boolean networks. These two areas resemble each other: synchronous Boolean networks are reminiscent of abstract dialectical frameworks. The question is whether or not there is more than this resemblance. We will explore possibilities of bridging these two areas.

#### 3.15 Logics for Data Exchange

Sonja Smets (University of Amsterdam, NL)

License 🐵 Creative Commons BY 4.0 International license Sonja Smets Joint work of Alexandru Baltag, Sonja Smets

We present a new family of dynamic logics that can model complex acts of data exchange and communication between different systems or agents. In the context of multi-agent systems, these are acts by which individual agents, as well as groups of agents, can publicly or privately access and transmit all the information stored at specific locations. In addition to having the full power of standard dynamic epistemic logics (DEL) to model acts of propositional communication, our logics can handle the type of communication protocols in which the data that is being communicated cannot be explicitly captured by a proposition.

### 3.16 Synergistic Knowledge

Thomas Studer (Universität Bern, CH)

License 
 Creative Commons BY 4.0 International license
 Thomas Studer

 Joint work of Christian Cachin, David Lehnherr, Thomas Studer
 Main reference Christian Cachin, David Lehnherr, Thomas Studer: "Synergistic Knowledge", in Proc. of the 25th International Symposium on Stabilization, Safety, and Security of Distributed Systems, SSS 2023, Jersey City, NJ, USA, October 2-4, 2023, LNCS, Vol. 14310, pp. 552–567, Springer, 2023.
 URL https://doi.org//10.1007/978-3-031-44274-2\_41

In formal epistemology, group knowledge is often modeled as the knowledge that the group would have if the agents shared all their individual knowledge. However, this interpretation does not account for relations between agents. In this work, we propose the notion of synergistic knowledge, which makes it possible to model those relationships. As an example, we investigate the use of consensus objects.

#### 3.17 Pattern Models and Consensus

Diego A. Velázquez (National Autonomous University of Mexico, MX)

License ☺ Creative Commons BY 4.0 International license © Diego A. Velázquez Joint work of Armando Castañeda, Hans van Ditmarsch, David A. Rosenblueth, Diego A. Velázquez

Recently, dynamic epistemic logic has started to be exploited for analyzing distributed systems. Pattern model logic was designed to that end. We present a simple-to-test sufficient condition that ensures consensus unsolvability for a given oblivious dynamic-network model.

#### 4 Working Groups

The presentations and discussions in the unique setting of Dagstuhl, with its relaxed and stimulating atmosphere, fully achieved their purpose: Long discussions took place during the official seminar, and many fruitful cross-community interactions spontaneously occurred during the free times, which even exceeded the amount of available time.

As a conclusion, we are convinced that our seminar will contribute to the further development of epistemic reasoning and topology in distributed computing, and the integration of the respective approaches.

#### 4.1 Extending Topology to Handle Data Structures

Henning Basold, Armando Castañeda, Faith Ellen, Guy Goren, Shihao (Jason) Liu, Susumu Nishimura, Thomas Nowak, Sergio Rajsbaum, Ulrich Schmid

License o Creative Commons BY 4.0 International license

© Henning Basold, Armando Castañeda, Faith Ellen, Guy Goren, Shihao (Jason) Liu, Susumu Nishimura, Thomas Nowak, Sergio Rajsbaum, Ulrich Schmid

#### 4.1.1 Context

The combinatorial topology approach to distributed computing has been shown to be powerful for studying solvability of fundamental distributed problems such a consensus, set agreement and renaming [42]. However, the current approach has a clear limitation as it can be applied only to distributed problems that can be modeled as (distributed) tasks. A *task* models a one-shot distributed problem where each process starts with a private input, and the processes are required to communicate with each other in order to decide private outputs that satisfy the task's specification. Nonetheless, there are fundamental distributed problems that cannot be specified as tasks [17, 18]. An important class of such problems are *concurrent objects* (or data structures), e.g., queues, stacks, trees and dictionaries. One seeks for implementations of them that are *linearizable* [43]. Roughly speaking, in a concurrent object, a process can perform an unbounded number of the operations provided by the object (differently from tasks where there is a single operation) and, in a linearizable implementation, responses to operations must be valid for the object that is implemented, somehow respecting the order operations are performed in a given execution.

#### 4.1.2 Known results

Formally, a task is a triple  $T = \langle I, O, \Delta \rangle$ , where I and O are input and output complexes specifying the input and output domains of the task, and  $\Delta$  is a carrier map specifying the complex  $\Delta(\sigma) \subseteq O$  with valid outputs for each input simplex  $\sigma \in I$ . An execution E is *valid* for the task if  $\tau \in \Delta(\sigma)$ , where  $\sigma$  and  $\tau$  are the simplexes with inputs and outputs in E, respectively.

Our discussion centered around finding a topological specification for concurrent objects. It first centered on previous work [17, 18], which already showed that tasks are too weak to even represent a one-shot queue, where each process can only access once by performing a fixed operation. The main problem is that carrier maps of tasks have limited ability to specify the order in which operations are performed in an execution.

The aforementioned work defined refined tasks (called long-lived task in [17]), where a vertex in the input complex represents a specific invocation to the object by a specific process. When using a simplex to represent an input configuration, multiple invocations of the same operation by the same process are represented by multiple distinct vertices in the simplex. Starting from any simplex in the input complex, each execution E specifies a simplex  $\tau_E$  where each vertex of  $\tau_E$  consists of a process identifier, an invocation to the object by this process, the response to this invocation in E, and the set-view when the response was returned. A set-view is the set of all invocations (by all processes) on the object up to a particular point in the execution E. The output complex consists of simplexes  $\tau_E$  specified by all possible executions (concurrent and sequential) E whose operations appear in simplexes of the input complex. A refined task is a triple  $T = \langle I, O, \Delta \rangle$  where  $\Delta(\sigma) \subseteq O$  contains simplexes  $\tau_E$  specified by all possible sequential E where processes perform the operations in a simplex  $\sigma \in I$ , in some order. A concurrent execution E' with operations in a simplex  $\sigma \in I$ is considered *valid* for the refined task if there exists a simplex  $\tau_E \in \Delta(\sigma)$  such that (1) the outputs to operations in E' and E agree, and (2) the set-view of each vertex in  $\tau_{E'}$  contains the set-view of its corresponding vertex in  $\tau_E$ .

Refined tasks have been shown to be expressive enough to model any linearizable concurrent object. However, this does not seem to resolve the limitation mentioned above because the notion of solvability for refined tasks is based on two components: outputs and set-views. The problem is that set-views are features of the execution, i.e., they are not produced or implied by processes' states. It is not clear how set-views can be captured by simplicial maps in a way that the combinatorial topology approach can be applied to study solvability of refined tasks. Let us recall that in the combinatorial topology approach protocols are modeled as simplicial maps that specify outputs as function on *only* process states, and the main question that needs to be answered is whether there exists a simplicial map from the protocol complex to the output complex, respecting the specification of the task in question. It may be possible that refined tasks can be used to study solvability of linearizable concurrent object S in an indirect manner as suggested in [33], where a task T is constructed such that if there is a linearizable implementation of S then T is solvable.

#### 4.1.3 Proposed approaches

The group discussed two possible alternative approaches to specify concurrent objects that may be more suitable for the combinatorial topology approach. Applications of these approaches remain to be seen.

# 4.1.3.1 Approach I: using a simplicial set to represent the possible outputs instead of an simplicial complex.

Here each vertex in the input complex I is defined by a process and a sequence of operations to be performed by the process. Each vertex in the output simplicial set O is defined by a process, a sequence of operations performed by the process, and the sequence of responses to these operations. Each facet of the simplicial set is labelled by a linearization, which is a total order of the operations in all of its vertices. The linearization labelling each facet needs to respect the order of operations within each of its vertices. A map  $\Delta$  specifies the subset of allowed output simplexes for each input simplex. Starting from any simplex  $\sigma \in I$ , an execution Ewith inputs described by  $\sigma$  is *correct* if there exists a linearization of the operations in E and a simplex  $\tau \in O$  labelled by this linearization such that, for each process, p, there is a vertex in  $\tau$  with process identifier p, the sequence of operations performed by p during E, and the sequence of responses p obtained during E.

#### 4.1.3.2 Approach II: modifying $\Delta$ so that it depends on the partial order of operations.

For simplicity, we first consider only one-shot objects. To represent the order in which operations are performed on a one-shot object, we can use a partial order on the set of processes. Specifically, if the operation by process p finishes before the operation by process q begins, then  $p \prec q$ . If the operation by process p is concurrent with the operation by process q, then p and q are incomparable.

Here a one-shot object is represented by a tuple  $(I, \Pi, O, \Delta)$ . *I* is the input complex. Each vertex in *I* is of the form (id, op), where *id* is a process identifier and *op* is an operation to be performed by this process.  $\Pi$  is the set of all possible partial orders on the set of processes. *O* is the output complex. Each vertex in *O* is of the form (id, op, resp), where

(id, op) is a vertex in I and resp is a possible response to the operation op performed by the process with identifier id.  $\Delta : I \times \Pi \to 2^O$  is a chromatic map which maps a simplex  $\sigma \in I$  and a partial order  $\pi \in \Pi$  to a subcomplex of O. If the object has a sequential specification, then each simplex in this subcomplex describes an output configuration resulting from a sequential execution whose order of operations respects the given partial order. An algorithm correctly implements an object if, for every execution E, there exists a simplex of  $\Delta(\sigma, \pi)$  that describes the response to each operation performed during E, where  $\sigma$  is the simplex of I describes the operation each process performs during the execution and  $\pi \in \Pi$  describes the partial order of these operations.

Now, consider any execution in which each process may access the object multiple times. For each process p, let  $n_p$  denote the number of operations performed by p during this execution. We can use a partial order on the set  $\{(p, i) \mid p \text{ is a process and } 1 \leq i \leq n_p\}$  to describe the relevant information about the operations performed during the execution. Specifically, if the *i*th operation by process p finishes before the *j*th operation by process q begins, then  $(p, i) \prec (q, j)$ . If the *i*th operation by process p is concurrent with the *j*th operation by process q, then (p, i) and (q, j) are incomparable.

Vertices in the input complex I are now of the form  $(p, seq_p)$ , where  $seq_p$  is a sequence of operations to be performed by the process p. Likewise, vertices in the output complex O are now of the form  $(p, seq_p, resp)$ , where resp is a sequence of possible responses to the operations in  $seq_p$ .  $\Delta$  now maps a simplex in I and a partial order on the set  $\{(p, i) \mid p \text{ is a process and } 1 \leq i \leq |seq_p|\}$  to a subcomplex of O.

Finally, it was also briefly discussed that the order of operations in the two proposed approaches could be captured through *higher-dimensional automata* (HDA) [53], which have an inherent geometric nature that may facilitate the use of combinatorial topology techniques.

#### 4.2 Interpreted Systems and Dynamic Epistemic Logic

Alexandru Baltag, Hans van Ditmarsch, Krisztina Fruzsa, Joseph Y. Halpern, Yoram Moses, Hugo Rincón Galeana, Ulrich Schmid, Sonja Smets, Diego A. Velázquez

License C Creative Commons BY 4.0 International license

© Alexandru Baltag, Hans van Ditmarsch, Krisztina Fruzsa, Joseph Y. Halpern, Yoram Moses, Hugo Rincón Galeana, Ulrich Schmid, Sonja Smets, Diego A. Velázquez

The Interpreted Systems (IS) approach ([31]) has, over the last four decades, been used to study distributed systems and has served as a tool in the design and analysis of distributed and multi-agent protocols and systems. The discussion began in this with the question of whether and how Dynamic Epistemic Logic (DEL) can contribute in a similar fashion.

Broadly speaking, we thus considered the following specific questions:

- What are the main differences between IS and DEL modeling?
- What could DEL add to IS modeling?
- What could IS add to DEL modeling?

The (partial) answers developed in this discussion group will be summarized in dedicated subsections below.

#### 4.2.1 What are the main differences between the IS and DEL modeling?

While both DEL and IS provide means to track the manner in which knowledge and belief change over time in a multi-agent system, the two approaches differ dramatically in the way that they model the effect that actions have. In DEL, actions operate directly on the

epistemic state of the system (this epistemic state is represented by a Kripke structure). For every action, the framework must provide an *action model*, which maps a Kripke structure and an action to a new Kripke structure, thereby describing the manner in which the epistemic state changes when the action is performed. In IS actions are taken to modify the physical configuration (or global state) of the system. As in automata theory, this is modeled by way of a transition function, mapping a configuration and an action to the configuration that results from performing the action. The epistemic state of the system in IS is also modeled using a Kripke structure, but this structure is induced by the system configurations, where each agent's indistinguishability relation is determined by its local state.

In the IS way of modeling of distributed systems, based on the runs-and-systems framework [31], agent protocols and the environment are outside of the formal logic used for epistemic reasoning. In fact, they are represented using *transition systems* (TS) which are the primary objects of study. Since the executions of a distributed protocol can typically be described in terms of a transition system (see [47]), the IS approach can readily be applied to a broad range of distributed systems problems.

In DEL [26], agent protocols (as well as the environment) can be expressed in the formal logic used for epistemic reasoning. In particular, action models enable one to model many forms of communication between the agents while focusing on capturing the resulting knowledge evolution of the involved agents. Moreover, precondition formulas can be used to tie communication actions to particular epistemic states of the agents (that are necessary to hold when communicating). Works such as [37], which even connect DEL and combinatorial topology, show that the idea of using DEL to help with topological reasoning about distributed systems is promising.

DEL and IS have been used for the analysis of communication protocols. Specific features of the protocols themselves indicate which framework fits best. The features we are looking at are a) temporal aspects (synchronous versus asynchronous communications) and b) the private versus public nature of the information contained in the events or actions that happen as part of the protocol. Here, we note that in DEL, time is not explicitly represented. Embeddings from DEL into IS (or rather into temporal epistemic logics) have been proposed. This includes both synchronous embeddings [10], wherein the execution of an action corresponds to a clock tick, and asynchronous embeddings [21], wherein histories of actions of variable length correspond to time in some way. While some believe that DEL and IS can be extended with ingredients so they can model the same kind of protocols, the starting base of these frameworks does differ. The built-in features of the starting base (such as the temporal aspects in IS and the action models in DEL) will make the application of one or the other framework a more straightforward choice for a given type of protocol. For instance when we focus on cryptographic protocols with complex cryptographic primitives as part of the actions or events that happen in the protocol, we see an added advantage of starting with DEL which offers us the expressive power of modeling the security aspects in the actions separately and allows us to automatically compute the state-transitions via its product update mechanism. DEL is used in the study of cryptographic protocols in works like [44, 20, 23, 30].

An advantageous feature of DEL is the fact that its built-in model update mechanism makes it immediately well-suited for automated model checking tasks. The implementation was done in different programming languages, but the most elaborate tool uses Haskell in DEMO [29]. The power of DEMO for model checking tasks has been illustrated on a small number of well-known sample protocols including comparison to temporal epistemic model checkers like MCK [27], which is based on IS modeling.

A disadvantage of DEL in general is the potential blowup of the model (or the precondition formulas). There are various reasons for such a blowup. As the models used in DEL are small (updates, and their temporal interpretation, are computable from an initial model and the sequence of actions), complexities with respect to model size are high [3]. Another reason is the iteration of action models, which leads to exponential growth of the size with each model update. This can be addressed with symmetry reductions [34], restriction to point-generated submodels [12], dedicated tools for specific frame classes (such as models with equivalence relations) instead of arbitrary binary relations [29], and bisimulation contraction [12].

A feature of DEL often forgotten is that such logics and frameworks are essentially logics of *observation*, not of action. One can investigate the consequence of actions if they were to happen, but one cannot state "let this action happen now". This is of course very different from IS. Once we add factual change to the DEL framework, this drawback is somewhat overcome. Now there are actions such as turning on and off the light, setting values of keys, etcetera. But agency remains lacking.

#### 4.2.2 What could DEL add to IS modeling?

Operationalizing the implicit Kripke model construction resulting from executing the system primarily requires the construction of action models for typical communication protocols considered in epistemic analysis via IS. The challenge here is to develop action models that avoid state explosion. One promising starting point could be communication patterns, updates that compare in interesting ways to action models [55, 15, 14]. The availability of such action models would open up some interesting avenues such as:

Could correctness proofs of protocols formulated via suitable precondition formulas of such action models be automated, via analyzing the Kripke model obtained by the product updates w.r.t. satisfying the desired properties? This would replace the translation of knowledge-based protocols obtained via IS to TS, and the obligatory proof that the resulting protocol guarantees that the agents attain their respective necessary level of knowledge.

Good target practice will be to model the consensus protocol for processes with sendomission faults from [1] in DEL, by (i) defining an action model for send-omission faults, (ii) modeling the protocol implementations via suitable precondition formulas, and (iii) generating and analyzing the Kripke model after suitably many rounds. We were optimistically thinking of including these results already in the report, however that was a bridge too far. Still, we are committed to this and other benchmarks to compare DEL and IS.

An area in DEL that is highly targeted towards protocols and automation is that of epistemic planning [13, 6, 48], and in particular concerning decidability of epistemic planning problems [2, 46]. Could epistemic planning techniques be used for synthesizing protocols in a distributed setting?

#### 4.2.3 What could IS add to DEL modeling?

The IS approach has been used to provide insights into the interaction between knowledge and communication in distributed systems ([39, 19, 7]), has allowed to generalize existing protocols ([40]), and to design optimal and state-of-the-art protocols solving well-known problems ([28, 50, 35, 16, 1]). These works provide insight into the role of synchrony and partial synchrony, refine our analysis of achievable types of optimality in different models of distributed computation, and into various aspects of fault-tolerance. The Knowledge of

Preconditions principle [49], which is naturally formulated in IS, provides model-independent theorem capturing the connection between knowledge and action in multi-agent systems. It has been shown to facilitate the design and analysis of efficient distributed systems protocols. Although protocols are not explicitly represented in the IS framework, it is possible to consider high-level "knowledge-based programs" in this setting, in which actions depend on tests for knowledge. In general, however, such programs are specifications and do not always uniquely determine the agents' behavior. Finding ways to implement them may be challenging. The more developed literature applying IS for the study of distributed computing can serve as an example and inspiration for directions in which DEL may be expanded. Fundamentally, DEL can be viewed as a subclass of IS. Its added declarative nature described above comes at a cost, however. Modeling and studying aspects of timing, including the role of synchrony, asynchrony and partial synchrony have been carried out successfully using the IS framework, while handling asynchrony, and partial synchrony appear to pose significant challenges for DEL. Similarly, incorporating failures, which have been studied in many IS settings, seems likely to give rise to cumbersome DEL models.

In order to be able to apply the IS framework, one only needs a description of the set of possible histories of a system, and a clear definition of the local states of agents in every configuration. As mentioned above, the fact that systems are often described by way of a transition system obviates the need to translate a given system into the formalism. Treating, say, an existing internet protocol using DEL would typically require a significant reformulation effort. Making DEL more flexible in this sense seems to be a considerable albeit worthwhile challenge. Finally, while the IS framework has given rise to new and improved protocols as well as tight bounds and impossibility results in different models of distributed computing, this has not been a focus of DEL works. It will be very interesting to see how well DEL can match such tasks, in light of the results obtained using IS.

#### 4.3 Representing Epistemic Attitudes via Simplicial Complexes

Alexandru Baltag, Henning Basold, Hans van Ditmarsch, Roman Kniazev, Sophia Knight, Roman Kuznets, Jérémy Ledent, David Lehnherr, Rojo Randrianomentsoa, Sonja Smets, Thomas Studer

License C Creative Commons BY 4.0 International license

© Alexandru Baltag, Henning Basold, Hans van Ditmarsch, Roman Kniazev, Sophia Knight, Roman Kuznets, Jérémy Ledent, David Lehnherr, Rojo Randrianomentsoa, Sonja Smets, Thomas Studer

The general header 'Epistemic Attitudes via Simplicial Complexes' permits a wealth of topics, also involving group epistemics such as distributed knowledge and common knowledge, and generalizations of simplicial complexes such as simplicial sets. However, the overall focus of the discussions turned towards a single issue: how to model *belief* on simplicial complexes, and, to a somewhat lesser extent, how to model dynamic epistemic attitudes such as deceit and Byzantine behaviour in general. In this report we focus on belief, the structures needed to encode it, their further generalization to polychromatic simplicial complexes (where simplexes can contain multiple vertices with the same colour), and the use of those for neighbourhood (epistemic) semantics on complexes.

The discussion was very much example driven. We therefore give a minimum of formal precision concerning complexes and their epistemic semantics. A simplicial complex C is a set of subsets of a set V of vertices that is downward closed, and a simplicial model is a triple  $(C, \chi, \ell)$  where C is a simplicial complex,  $\chi$  is a chromatic map assigning an agent

(or a colour) to each vertex and such that different vertices in a face are assigned different colours, and  $\ell$  is a *valuation* of propositional variables to vertices (or, alternatively, to faces). A maximal face is a facet. Intuitively a vertex represents the local state of an agent and a facet represents a global state. In a *simplicial set* there can be multiple occurrences of a simplex (thus it is a multi-set of subsets and not a set).

The semantics of knowledge on simplicial complexes replaces the usual Kripke model semantics with equivalence classes for agents a, given binary (equivalence) relations, with adjacency of facets (or faces) to a vertex with colour a. The basic definition is as follows, where we refer to works like [45, 37, 24, 36] for details. Let X be a facet in C ( $X \in \mathcal{F}(C)$ ) of simplicial model  $\mathcal{C} = (C, \chi, \ell)$ . We consider formulas in a standard epistemic logical language. Then:

$$\mathcal{C}, X \models K_a \varphi$$
 iff  $\mathcal{C}, Y \models \varphi$  for all  $Y \in \mathcal{F}(C)$  such that  $a \in \chi(X \cap Y)$ .

In fact we can interpret formulas in arbitrary faces of a complex<sup>1</sup> where the case of a vertex, a singleton face, captures the knowledge semantics as follows, where  $\chi(v) = a$ :

 $\mathcal{C}, v \models K_a \varphi$  iff  $\mathcal{C}, Y \models \varphi$  for all  $Y \in C$  with  $v \in Y$ 

We will use a running example to demonstrate all semantics for knowledge, belief, and other epistemic attitudes. And we start by explaining the knowledge semantics. Hopefully, this will suffice for the reader to prepare the ground for the subsequent different belief semantics.

Consider a simplicial model C as above for three agents (colours) a, b, c, where for convenience the vertices are named xy for x = a, b, c and y = 0, 1, 2 and where these numbers simultaneously represent local values, formalized as propositional variables  $y_x$  only true there. (So in facet X variables  $0_a$ ,  $0_c$ , and  $1_b$  are true and  $1_a, 2_a, 1_c, 0_b, 2_b$  are false.) We also assume that variables are local for vertices but that does not seem to be of prime importance (one could think of assigning variables to facets).



 $\mathcal{C}^i$ 

We now have that in facet X agent a knows that the value of c is 0, because in the facets W and X intersecting with X in vertex a0 the value of c is 0. However, agent b is uncertain there whether the value of c is 0 or 1 (where we use  $\hat{K}_b$  to abbreviate  $\neg K_b \neg$ ):

$$\mathcal{C}, X \models K_a 0_c \qquad \qquad \mathcal{C}, X \models \widehat{K}_b 0_c \wedge \widehat{K}_b 1_c$$

#### 4.3.1 Irrevocable Belief

Let us now model belief  $(B_a)$  instead of knowledge  $(K_a)$ . Unlike knowledge, belief may be incorrect. A common way to model belief in modal logic, interpreted on Kripke models, is by way of relations that are serial, transitive, and Euclidean, for which the logic is KD45. That is, it satisfies:

<sup>&</sup>lt;sup>1</sup> To interpret formulas in arbitrary faces we use the multi-pointed semantics of [24], section 'Local semantics for simplicial complexes', wherein  $\mathcal{C}, \mathbf{X} \models \varphi$  for a set of faces  $\mathbf{X}$  iff  $\mathcal{C}, \mathbf{X} \models \varphi$  for all  $X \in \mathbf{X}$ . In particular,  $\mathcal{C}, v \models \varphi$  then means that  $\mathcal{C}, \mathsf{star}(v) \models \varphi$ , where  $\mathsf{star}(v) = \{Y \in C \mid v \in Y\}$ .

- $\blacksquare B_a \varphi \to \widehat{B}_a \varphi \text{ (consistency)}$
- $= B_a \to B_a B_a \varphi$  (positive introspection)

 $\blacksquare B_a \varphi \to B_a B_a \varphi \text{ (negative introspection)}$ 

This is known as consistent belief [25] and also as irrevocable belief (or as conviction) [54]. Recalling the KD45 "balloons" (clusters of indistinguishable worlds) with "ropes" (asymmetric pairs in the accessibility relation) from isolated worlds pointing to these clusters, a way to model belief on complexes is with a *belief function*  $f : \mathcal{V}(C) \to \mathcal{V}(C)$  that is an idempotent and color-preserving vertex map. It is not a chromatic map as it may not preserve other simplices. Let  $X_a$  be the vertex of color a that belongs to a facet X. Belief is now defined as

 $\mathcal{C}, X \models B_a \varphi$  iff  $\mathcal{C}, Y \models \varphi$  for all  $Y \in \mathcal{F}(C)$  with  $f(X_a) \in Y$ 

If, for a world v, f(v) = v, belief becomes knowledge. The map needs to be idempotent to ensure that the properties of belief hold, we can see the map f as "pointing" vertex v to a cluster of f(v) adjoining facets, in words referring to a corresponding Kripke model, to a set of b-indistinguishable worlds.

As an example, we reconsider the model above but now enriched with a belief function that maps vertex a0 to vertex a2 and otherwise maps all vertices to themselves, as below.



This represents that in local state a0 agent a believes that her local state is a2 and that the global state is Z. Observe that this also implies that a incorrectly believes that the value of her local state is 2 and not 0. However, a reasonable assumption may be that an agent correctly knows its local state, and may only be mistaken about what it believes to be the local states of other agents. (This assumption underlies some of our later epistemic explorations, wherein different vertices of the same colour jointly represent an agent's beliefs.) We now have that, for example,

$\mathcal{C}, X \models 1_b \land \neg B_a 1_b \land B_a 2_b$	agent $a$ incorrectly believes that the local state of $b$ is 2
$\mathcal{C}, Z \models 2_b \wedge B_a 2_b$	but $a \ knows$ that the local state of $b$ is 2 in global state $Z$
$\mathcal{C}, X \models B_b 1_b \wedge B_b B_a 2_b$	agent $b$ knows that agent $a$ has incorrect beliefs about him

and so forth ... Note that we only have one modality in the logical language here, the informal "agent a knows" and "agent b knows" above are only written because in the facets (global states) where these formulas are interpreted the agents' beliefs are correct.

As a side issue, let us now consider the belief function g such that g(a0) = a2, g(c0) = c1, and otherwise the identity. In the model  $C^i$ , the distributed knowledge of agents a and cin facet W is represented by the edge  $\{c0, a0\}$ . Whereas in the model  $C^{ii}$  but with belief function g, agents a and c have distributed belief embodied in the edge  $\{c1, a2\}$ . Group epistemic notions other than for knowledge were not discussed, but might worth to be investigated later.

The belief function is an enrichment of the simplicial model. It does not, on first sight, seem very topological in nature, and therefore not very simplicial. It rather looks like an accessibility relation between the worlds of a Kripke model. In fact, taken as a relation, a

belief function is serial, transitive, and Euclidean, just as required for belief. Let us see how we can stretch this analogy a bit further before we return to more topological interpretations, in Subsection 4.3.3.

#### 4.3.2 Revocable Belief, Safe Belief, and Knowledge

Now assume any function that is a transitive relation. Consider the belief function f that maps vertex a0 to a1 and vertex a1 to a2, and therefore, by transitivity, a0 to a2. As follows (transitivity is implicit):

$$b0 \xrightarrow{a0} a0 \xrightarrow{a1} a2$$

$$W \xrightarrow{X} \xrightarrow{Y} \xrightarrow{Z} c1 \xrightarrow{b1} b2$$

$$C^{iii}$$

This induces a (strict) total order on the *a* vertices such that a0 > a1 > a2, with an associated weak total order  $\geq$ . We can now interpret this total order as sorting comparable *a*-vertices into more and less plausible, where in this case *a*2 is most plausible, *a*1 somewhat less, and *a*0 least, by analogy of the plausibility models of [4] (and not dissimilar to [8, 22], see also [11]). Let us additionally write  $a0 \geq a2$  to indicate that *a*2 is the most plausible local state for *a* in this total order (or, dually,  $\leq$  and  $\leq$ ),  $\equiv$  for equally plausible, and ~ for comparable.

We now can assume a multi-modal language and define notions of belief (revocable/defeasible belief), safe belief (strong belief), and knowledge. Depending on one's preferred dialect this can be done in different ways. In the semantics for *a*-adjoining facets we would get, where  $X \in \mathcal{F}(C)$ :

$$\begin{array}{lll} \mathcal{C}, X \models B_a \varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ and } v \leq X_a \text{ with } v \in X \\ \mathcal{C}, X \models \Box_a \varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ and } v \leq X_a \text{ with } v \in X \\ \mathcal{C}, X \models K_a \varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ and } v \sim X_a \text{ with } v \in X \end{array}$$

The relation  $\geq$  between *a*-vertices induces a relation  $\geq_a$  between facets, namely such that  $X \geq_a Y$  if  $X_a \geq Y_a$ , and analogously we can induce relations  $\geq_a$ ,  $\equiv_a$  and  $\sim_a$ . Note that as a consequence all facets containing a given local state for agent *a* are equally plausible for her. Using this relation between facets, for the semantics of defeasible belief  $B_a$  we would now get (where those for  $\Box_a$  and  $K_a$  are analogous) the following equivalent semantics, where for good measure we have added two more equivalent formulations. Which one looks most simplicial?

$$\begin{array}{ll} \mathcal{C}, X \models B_a \varphi & \text{iff} \quad \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ with } Y \leq a X \\ \mathcal{C}, x \models B_a \varphi & \text{iff} \quad \mathcal{C}, Y \models \varphi \text{ for all } Y \in C \text{ with } v \leq x \text{ and } v \in Y \\ \mathcal{C}, x \models B_a \varphi & \text{iff} \quad \mathcal{C}, v \models \varphi \text{ for all } v \in \mathcal{V}(C) \text{ with } v \leq x \end{array}$$

As an example of the semantics above we can easily determine in the simplicial model  $\mathcal{C}^{iii}$  that

$\mathcal{C}^{iii}, X \models B_a 2_b$	Note $a2 \leq a0 = X_a$ , and Z is the only facet containing $a2$
$\mathcal{C}^{iii}, Y \models \Box_a \widehat{B}_c \mathbb{1}_b$	Belief in $\widehat{B}_c 1_b$ is safe for $a$ in $Y$
	as $Y, Z \leq Y$ and both satisfy $\widehat{B}_c 1_b$ .
$\mathcal{C}^{iii}, Y \not\models \Box_a 2_b$	But belief in $2_b$ is not safe for $a$ as $Y \not\models 2_b$ .
$\mathcal{C}^{iii}, X \not\models K_a(0_b \lor 1_b)$	This is because $Z \leq_a X$ and $Z \not\models 0_b \lor 1_b$ .

This notion of knowledge is different from the standard notion of simplicial knowledge applied in model  $\mathcal{C}^i$ . We recall that  $\mathcal{C}^i, X \models K_a(0_b \vee 1_b)$ , as knowledge in X means truth in W and X. Whereas  $\mathcal{C}^{iii}, X \not\models K_a(0_b \vee 1_b)$ , as this 'novel' knowledge in X means truth in W, X, Y, Z, and  $0_b \vee 1_b$  is false in Z.

This can be adjusted by **tentatively** proposing a belief quotient of a simplicial model with respect to a belief function inducing a total order  $\leq$  as above. As we will see this may not always be a simplicial model, but it seems it always will be a simplicial set. The belief quotient of a complex  $C = (C, \chi, \ell)$  would be  $C_{\leq} = (C_{\leq}, \chi_{\leq}, \ell_{\leq})$  as follows. Let for any  $v \in \mathcal{V}(C)$  vertex  $v_{\leq}$  be defined as  $\{w \in \mathcal{V}(C) \mid w \sim v\}$ , then  $\mathcal{V}(C_{\leq}) := \{v_{\leq} \mid v \in \mathcal{V}(C)\}$ . Note that the choice of vertex determines its colour, so  $v_{\leq}$  is the equivalence class of all vertices w that are more or less plausible then v for agent  $\chi(v)$ , so there is no need for an 'a' as in  $\leq_a$ . The faces of  $C_{\leq}$  now are Y for which there is a  $Z \in C$  with  $Y = \{v_{\leq} \mid v \in Z\}$ . As comparable vertices may have different labelings of propositional variables, proposing  $\ell_{\leq}(v_{\leq}) := \bigcup \{\ell(v) \mid w \sim_a v\}$  may not be very useful. However, we recall that requiring labelings to be the same for comparable vertices was already mentioned as a reasonable assumption for belief functions. Let us therefore assume that all  $\leq$ -comparable vertices have the same labeling. We now can produce some nice examples.

In  $C^i$ , the variables  $0_a, 1_a, 2_a$  were only true in a0, a1, a2, respectively. Let us now, instead, assume that they are true in all three a0, a1, a2 so that  $0_a \leftrightarrow 1_a$  and  $1_a \leftrightarrow 2_a$  on the model. We now simply take  $a0 \leq$  to be a0 itself. Given that, and some visual trickery, the belief quotient of  $C^{ii}$  is  $C^{iv}$  below.



Similarly, the belief quotient of  $C^{iii}$  is  $C^v$  below, where we note that this is a simplicial set, as the edges  $\{a0, b1\}$  and  $\{a0, c1\}$  occur twice in the complex (denoted by the double-lined edge).



This representation does not make any difference for the semantics, but we are now back to a possible more pleasing, and familiar, notation for more and less plausible facets adjoining a vertex. In  $\mathcal{C}^{iv}$  we have that vertex a0 intersects with W, X, Z where W and Xare equally plausible and Z is more (and most) plausible. In  $\mathcal{C}^v$  we have that vertex a0intersects with W, X, Y, Z where W and X are equally plausible, both are more plausible than Y, which is more plausible than Z.

#### 4.3.3 Polychromatic Simplicial Complexes

A simple belief function for irrevocable belief, or a more complex belief function to model defeasible or strong belief, or even knowledge in a different way, is not a very topological tool. Nor is ordering the different facets intersecting in a vertex in the belief quotient version. It was widely discussed how to make such ideas into something more topological. Reconsidering the simplicial model  $C^{ii}$  once more, instead of considering pairs (a0, a1), (a1, a2) (and (a0, a2)) in a relation, we can also consider such a pair as an edge, which is topology.

A structure where both vertices of an edge have the same colour would no longer be a chromatic simplicial complex, where in any simplex all vertices must have different colours. We tentatively call complexes where the  $\chi$  map may assign the same colour to more than one vertex in a simplex a *polychromatic simplicial complex*.

Starting out with edges  $\{a0, a1\}$  and  $\{a0, a1\}$  in the model  $\mathcal{C}^{iii}$  we would get the model  $\mathcal{C}^{vi}$  as below.



This model  $\mathcal{C}^{vi}$  allows for additional epistemic notions, using the novel topological information. For example, reminiscent of the plausibility order (but also essentially different, as we will see) consider the shortest *a*-path linking vertices. Given *a*0, *a*1 is one step away, and *a*2 is two steps away. Based on that, and somewhat dual to the prior  $\mathcal{C}^{iii}$ , we can interpret this as agent *a* considering *a*0 its most plausible local state, *a*1 less, and *a*2 least (the dual where *a*2 is most and *a*0 least plausible from the perspective of *a*0 seems less intuitive in this interpretation). Similarly we can order facets that way, from the perspective of a given agent *a*, where the distance between facets is the length of the shortest *a*-path between them. We now interpret  $B_a\varphi$  as before, as truth in the most plausible facet or vertex.

Note that there is a difference with the previous semantics for total  $\leq$  (and  $\leq_a$ ) orders: the truth of  $B_a\varphi$  depends on the actual facet. This is unlike the notions of belief and knowledge employing plausibility models in the previous subsection. These are subjective notions that do not depend on the actual point of evaluation (however, unlike strong belief, that is an 'objective' notion as well, that is, a notion depending on the point of evaluation given an equivalence class for an agent).

In model  $\mathcal{C}^{vi}$  we would now have that  $\mathcal{C}^{vi}, X \models B_a 0_c$ , whereas  $\mathcal{C}^{vi}, Y \models B_a 1_b$ , and  $\mathcal{C}^{vi}, Z \models B_a 2_b$  (where all these are distinguishing formulas in the model, that is, only true in those facets and in none other).

If we let pair (a0, a2) in (relation induced by) the belief function count as well, we get this model  $C^{vii}$  and there seems nothing against considering them as the three edges of a triangle, as in  $C^{viii}$ :



 $\mathcal{C}^{vii}$ 



In this more topological interpretation involving simplexes (edges or triangles, here) of vertices with the same colour, we can maybe somehow interpret this as different mental states of the same agent, a structured form of the beliefs the agent considers possible. The representation clearly allows agent a to reason about what her beliefs are in vertex a0 as different from what her beliefs are in vertex a1, or a2. Relations and generalizations to (semi-simplicial sets and) simplicial sets are again conceivable. So far, only semi-simplicial sets have been studied.

It would also be interesting to investigate what happens if we allow *higher-dimensional* paths, such as the chaining in order to obtain *common distributed knowledge* [5, 24].

Finally, how can we interpret formulas on simplexes where some but not all vertices are of the same colour, such as the one below?



Another surely fertile direction of research would be to consider models like  $C^{vii}$  and  $C^{viii}$ in a neighbourhood semantics [41, 51], a common framework in modal logical semantics although infrequently applied to epistemic notions [9]. Exploring neighbourhood semantics might lead to new insights on *coalition logic* which found recent application to the specification and verification of smart contracts [52].

In neighbourhood semantics, going into simplicial mode, we would have that  $\mathcal{C}, X \models K_a \varphi$ if there is a *neighbourhood* (a set of facets) of X for agent a such that  $\varphi$  is true in all the facets in that neighbourhood. Or, possibly more intuitively, given a vertex v with colour a,  $K_a \varphi$  is true in v if there is an a-neighbourhood of v where  $\varphi$  is true.

In  $C^{vii}$  (or  $C^{viii}$ ) facet W has three neighbourhoods for a:  $\{W, X\}$ ,  $\{Z\}$ , and  $\{Y\}$ . Agents b and c only have a singleton neighbourhood in these structures, we revert to the standard knowledge semantics here. Therefore true in W are:  $K_a 0_c$  (in  $\{W, X\}$ ),  $K_a 1_b$ , and  $K_a 2_b$ . Observe that these three neighbourhoods do not intersect (there are not any two such that their intersection contains a facet), which, in Kripke semantics, is common fare for neighbourhoods. (In this sort of interpretation, it seems even impossible.) Although all this is crying for an intuitive interpretation, we will leave this as such in this report and instead relay it to further research.

#### 4.3.4 Simplicial Structures for Byzantine Actions

It was agreed upon that before modelling Byzantine agents (that is, unreliable information or malicious intentions), we needed to have a fitting notion of belief. And indeed that occupied the main part of this report. It is clear that, as usual in action model logic, similar structures as presented above for the static notion of belief can be employed to encode dynamic notions involving deceit, error, and Byzantine conduct in general. Indeed, in simplicial semantics the "Kripke-style action models" have natural counterparts as *simplicial action models* [45, 24].

 $C^{viii}$ 

However, asymmetric versions of that needed for Byzantine communication have not yet been investigated. Additionally, novel update mechanisms such as communication patterns and related [5, 14] could also be contemplated to model Byzantine conduct. The update of simplicial models with communication patterns has been described in [14], including summary suggestions how to represent them as "simplicial communication patterns". All these horizons seem currently pretty far away, although mouth-watering in their presumed modelling features.

Let us close this far too short subsection on dynamics with a simplicial action model  $C^x$  representing that agent c dies (that process c crashes), although agents a and b only partially observe that, they remain uncertain about this. The preconditions of all three vertices are  $\top$  (the always true proposition). As before, there are two a-b edges, the lower edge is a maximal simplex (the complex is an impure simplicial set). If we update a simplicial model with this simplicial action model, any simplex is duplicated into one where c is alive and a for a and b indistinguishable one where c is dead. Even more elementary, simplicial action model  $C^{xi}$  consists of the a-b edge part of  $C^x$  and models that c is dying, pure and simple. And then, finally, to have at least some Byzantine in the picture, the action  $C^{xii}$  where a incorrectly believes c to be dying, to the despair of b and, obviously, c who both observe that.



#### References

- 1 Kaya Alpturer, Joseph Y. Halpern, and Ron van der Meyden. Optimal eventual Byzantine agreement protocols with omission failures. In PODC'23: Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing, pages 244–252. Association for Computing Machinery, 2023. doi:10.1145/3583668.3594573.
- 2 Guillaume Aucher and Thomas Bolander. Undecidability in epistemic planning. In Francesca Rossi, editor, IJCAI-13, Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence: Beijing, China, 3-9 August 2013, pages 27-33. AAAI Press, 2013. Available from: https://www.ijcai.org/Abstract/13/015.
- 3 Guillaume Aucher and François Schwarzentruber. On the complexity of dynamic epistemic logic. In Burkhard C. Schipper, editor, TARK 2013: Theoretical Aspects of Rationality and Knowledge, Proceedings of the 14th Conference Chennai, India, pages 19–28. University of California, Davis, 2013. doi:10.48550/arXiv.1310.6406.
- 4 Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, editors, Logic and the Foundations of Game and Decision Theory (LOFT 7), volume 3 of Texts in Logic and Games, pages 11–58. Amsterdam University Press, 2008. Available from: https://www.jstor.org/stable/j.ctt46mz4h.4.
- 5 Alexandru Baltag and Sonja Smets. Learning what others know. In Elvira Albert and Laura Kovács, editors, LPAR23. LPAR-23: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, volume 73 of EPiC Series in Computing, pages 90-119. EasyChair, 2020. doi:10.29007/plm4.
- 6 Vaishak Belle, Thomas Bolander, Andreas Herzig, and Bernhard Nebel. Epistemic planning: Perspectives on the special issue. Artificial Intelligence, 316:103842, March 2023. doi: 10.1016/j.artint.2022.103842.

- 7 Ido Ben-Zvi and Yoram Moses. Beyond Lamport's *Happened-before*: On time bounds and the ordering of events in distributed systems. *Journal of the ACM*, 61(2):13:1–13:26, April 2014. doi:10.1145/2542181.
- 8 Johan van Benthem. Dynamic logic for belief revision. Journal of Applied Non-Classical Logics, 17(2):129–155, 2007. doi:10.3166/jancl.17.129-155.
- 9 Johan van Benthem, David Fernández-Duque, and Eric Pacuit. Evidence and plausibility in neighborhood structures. Annals of Pure and Applied Logic, 165(1):106–133, January 2014. doi:10.1016/j.apal.2013.07.007.
- 10 Johan van Benthem, Jelle Gerbrandy, Tomohiro Hoshi, and Eric Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, 38(5):491–526, October 2009. doi:10.1007/ s10992-008-9099-x.
- 11 Johan van Benthem and Sonja Smets. Dynamic logics of belief change. In Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi, editors, *Handbook of Epistemic Logic*, pages 313–393. College Publications, 2015.
- 12 Patrick Blackburn, Maarten de Rijke, and Yde Venema. Modal Logic, volume 53 of Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001. doi:10.1017/CB09781107050884.
- 13 Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single- and multi-agent systems. Journal of Applied Non-classical Logics, 21(1):9–34, 2011. doi: 10.3166/jancl.21.9-34.
- 14 Armando Castañeda, Hans van Ditmarsch, David A. Rosenblueth, and Diego A. Velázquez. Communication pattern logic: Epistemic and topological views. *Journal of Philosophical Logic*, 2023. doi:10.1007/s10992-023-09713-8.
- 15 Armando Castañeda, Hans van Ditmarsch, David A. Rosenblueth, and Diego A. Velázquez. Comparing the update expressivity of communication patterns and action models. In Rineke Verbrugge, editor, Proceedings Nineteenth conference on Theoretical Aspects of Rationality and Knowledge, Oxford, United Kingdom, 28–30th June 2023, volume 379 of Electronic Proceedings in Theoretical Computer Science, pages 157–172. Open Publishing Association, 2023. doi:10.4204/EPTCS.379.14.
- 16 Armando Castañeda, Yannai A. Gonczarowski, and Yoram Moses. Unbeatable consensus. Distributed Computing, 35(2):123–143, April 2022. doi:10.1007/s00446-021-00417-3.
- 17 Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. Long-lived tasks. In Amr El Abbadi and Benoît Garbinato, editors, Networked Systems: 5th International Conference, NETYS 2017, Marrakech, Morocco, May 17–19, 2017, Proceedings, volume 10299 of Lecture Notes in Computer Science, pages 439–454. Springer, 2017. doi: 10.1007/978-3-319-59647-1\_32.
- 18 Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. Unifying concurrent objects and distributed tasks: Interval-linearizability. *Journal of the ACM*, 65(6):45:1–45:42, December 2018. doi:10.1145/3266457.
- K. Mani Chandy and Jayadev Misra. How processes learn. Distributed Computing, 1(1):40–52, March 1986. doi:10.1007/BF01843569.
- 20 Francien Dechesne and Yanjing Wang. To know or not to know: epistemic approaches to security protocol verification. Synthese, 177(Supplement 1):51-76, 2010. doi:10.1007/s11229-010-9765-8.
- 21 Cédric Dégremont, Benedikt Löwe, and Andreas Witzel. The synchronicity of dynamic epistemic logic. In Krzysztof R. Apt, editor, *TARK XIII, Theoretical Aspects of Rationality and Knowledge: Proceedings of the Thirteenth Conference (TARK 2011)*, pages 145–152. Association for Computing Machinery, 2011. doi:10.1145/2000378.2000395.
- 22 Hans van Ditmarsch. Prolegomena to dynamic logic for belief revision. Synthese, 147(2):229–275, November 2005. doi:10.1007/s11229-005-1349-7.

- 23 Hans van Ditmarsch, Jan van Eijck, Ignacio Hernández-Antón, Floor Sietsma, Sunil Simon, and Fernando Soler-Toscano. Modelling cryptographic keys in dynamic epistemic logic with DEMO. In Javier Bajo Pérez et al., editor, *Highlights on Practical Applications of Agents and Multi-Agent Systems: 10th International Conference on Practical Applications of Agents and Multi-Agent Systems*, volume 156 of Advances in Intelligent and Soft Computing, pages 155–162. Springer, 2012. doi:10.1007/978-3-642-28762-6\_19.
- 24 Hans van Ditmarsch, Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. Knowledge and simplicial complexes. In Björn Lundgren and Nancy Abigail Nuñez Hernández, editors, *Philosophy of Computing: Themes from IACAP 2019*, volume 143 of *Philosophical Studies Series*, pages 1–50. Springer, 2022. doi:10.1007/978-3-030-75267-5\_1.
- 25 Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi. An introduction to logics of knowledge and belief. In Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi, editors, *Handbook of Epistemic Logic*, pages 1–51. College Publications, 2015.
- 26 Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. Dynamic Epistemic Logic, volume 337 of Synthese Library. Springer, 2007. doi:10.1007/978-1-4020-5839-4.
- 27 Hans van Ditmarsch, Wiebe van der Hoek, Ron van der Meyden, and Ji Ruan. Model checking Russian Cards. In Proceedings of the Third Workshop on Model Checking and Artificial Intelligence (MoChArt 2005), volume 149(2) of Electronic Notes in Theoretical Computer Science, pages 105–123. Elsevier, 2006. doi:10.1016/j.entcs.2005.07.029.
- 28 Cynthia Dwork and Yoram Moses. Knowledge and common knowledge in a Byzantine environment: Crash failures. *Information and Computation*, 88(2):156–186, October 1990. doi:10.1016/0890-5401(90)90014-9.
- 29 Jan van Eijck. DEMO A demo of epistemic modelling. In Johan van Benthem, Dov Gabbay, and Benedikt Löwe, editors, Interactive Logic: Selected Papers from the 7th Augustus de Morgan Workshop, London, volume 1 of Texts in Logic and Games, pages 303–362. Amsterdam University Press, 2007. Available from: https://www.jstor.org/stable/j.ctt45kdbf.15.
- 30 Jan van Eijck and Malvin Gattinger. Elements of epistemic crypto logic (extended abstract). In AAMAS'15: Proceedings of the 2015 International Conference on Autonomous Agents & Multiagent Systems, pages 1795–1796. Association for Computing Machinery, 2015. Available from: https://www.ifaamas.org/Proceedings/aamas2015/aamas/p1795.pdf.
- 31 Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. Reasoning About Knowledge. MIT Press, 1995. doi:10.7551/mitpress/5803.001.0001.
- 32 Krisztina Fruzsa, Roman Kuznets, and Ulrich Schmid. Fire! In Joseph Halpern and Andrés Perea, editors, Proceedings Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, Beijing, China, June 25–27, 2021, volume 335 of Electronic Proceedings in Theoretical Computer Science, pages 139–153. Open Publishing Association, 2021. doi: 10.4204/EPTCS.335.13.
- 33 Eli Gafni. Snapshot for time: The one-shot case. Eprint 1408.3432, arXiv, 2014. doi: 10.48550/arXiv.1408.3432.
- 34 Nina Gierasimczuk and Jakub Szymanik. A note on a generalization of the Muddy Children puzzle. In Krzysztof R. Apt, editor, TARK XIII, Theoretical Aspects of Rationality and Knowledge: Proceedings of the Thirteenth Conference (TARK 2011), pages 257–264. Association for Computing Machinery, 2011. doi:10.1145/2000378.2000409.
- 35 Guy Goren and Yoram Moses. Silence. Journal of the ACM, 67(1):3:1–3:26, February 2020. doi:10.1145/3377883.
- 36 Éric Goubault, Roman Kniazev, Jérémy Ledent, and Sergio Rajsbaum. Semi-simplicial set models for distributed knowledge. In 2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). IEEE, 2023. doi:10.1109/LICS56636.2023.10175737.

- 37 Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. A simplicial complex model for dynamic epistemic logic to study distributed task computability. *Information and Computation*, 278:104597, June 2021. doi:10.1016/j.ic.2020.104597.
- 38 Michell Guzmán, Sophia Knight, Santiago Quintero, Sergio Ramírez, Camilo Rueda, and Frank Valencia. Reasoning about distributed knowledge of groups with infinitely many agents. In Wan Fokkink and Rob van Glabbeek, editors, 30th International Conference on Concurrency Theory: CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands, volume 140 of Leibniz International Proceedings in Informatics (LIPIcs), pages 29:1–29:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.CONCUR. 2019.29.
- **39** Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, July 1990. doi:10.1145/79147.79161.
- 40 Joseph Y. Halpern and Lenore D. Zuck. A little knowledge goes a long way: Knowledgebased derivations and correctness proofs for a family of protocols. *Journal of the ACM*, 39(3):449–478, July 1992. doi:10.1145/146637.146638.
- 41 Helle Hvid Hansen, Clemens Kupke, and Eric Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 5(2):2:1–2:38, April 2009. doi:10.2168/LMCS-5(2:2)2009.
- 42 Maurice P. Herlihy, Dmitry Kozlov, and Sergio Rajsbaum. Distributed Computing through Combinatorial Topology. Morgan Kaufmann, 2014. doi:10.1016/C2011-0-07032-1.
- 43 Maurice P. Herlihy and Jeannette M. Wing. Linearizability: A correctness condition for concurrent objects. ACM Transactions on Programming Languages and Systems, 12(3):463– 492, July 1990. doi:10.1145/78969.78972.
- 44 Arjen Hommersom, John-Jules Meyer, and Erik de Vink. Update semantics of security protocols. *Synthese*, 142(2):229–267, September 2004. doi:10.1007/s11229-004-2247-0.
- 45 Jérémy Ledent. *Geometric semantics for asynchronous computability*. PhD thesis, Paris-Saclay University, Palaiseau, France, 2019. Prepared at École polytechnique. Available from: https://theses.hal.science/tel-02445180.
- 46 Qiang Liu and Yongmei Liu. Multi-agent epistemic planning with common knowledge. In Jérôme Lang, editor, Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, Stockholm, 13–19 July 2018, pages 1912–1920. International Joint Conferences on Artificial Intelligence, 2018. doi:10.24963/ijcai.2018/264.
- 47 Nancy A. Lynch. Distributed Algorithms. Morgan Kaufmann, 1996.
- 48 Bastien Maubert, Sophie Pinchinat, François Schwarzentruber, and Silvia Stranieri. Concurrent games in Dynamic Epistemic Logic. In Christian Bessiere, editor, Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, Yokohama, January 2021, pages 1877–1883. International Joint Conferences on Artificial Intelligence, 2020. doi:10.24963/ijcai.2020/260.
- 49 Yoram Moses. Relating knowledge and coordinated action: The knowledge of preconditions principle. In R. Ramanujam, editor, Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, Carnegie Mellon University, Pittsburgh, USA, June 4–6, 2015, volume 215 of Electronic Proceedings in Theoretical Computer Science, pages 231–245. Open Publishing Association, 2016. doi:10.4204/EPTCS.215.17.
- 50 Yoram Moses and Mark R. Tuttle. Programming simultaneous actions using common knowledge. Algorithmica, 3(1–4):121–169, November 1988. doi:10.1007/BF01762112.
- 51 Eric Pacuit. Neighborhood Semantics for Modal Logic. Short Textbooks in Logic. Springer, 2017. doi:10.1007/978-3-319-67149-9.
- 52 Marc Pauly and Rohit Parikh. Game Logic An overview. *Studia Logica*, 75(2):165–182, November 2003. doi:10.1023/A:1027354826364.

- 53 Vaughan Pratt. Modeling concurrency with geometry. In Conference Record of the Eighteenth Annual ACM Symposium on Principles of Programming Languages, pages 311–322. Association for Computing Machinery, 1991. doi:10.1145/99583.99625.
- 54 Krister Segerberg. Irrevocable belief revision in dynamic doxastic logic. *Notre Dame Journal* of Formal Logic, 39(3):287–306, summer 1998. doi:10.1305/ndjfl/1039182247.
- 55 Diego A. Velázquez, Armando Castañeda, and David A. Rosenblueth. Communication pattern models: An extension of action models for dynamic-network distributed systems. In Joseph Halpern and Andrés Perea, editors, Proceedings Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, Beijing, China, June 25–27, 2021, volume 335 of Electronic Proceedings in Theoretical Computer Science, pages 307–321. Open Publishing Association, 2021. doi:10.4204/EPTCS.335.29.

#### Participants

Alexandru Baltag
 University of Amsterdam, NL

Henning Basold
 Leiden University, NL

 Armando Castañeda
 National Autonomous University of Mexico, MX

Hans van Ditmarsch
 CNRS – Toulouse, FR &
 Université de Toulouse, FR

■ Faith Ellen University of Toronto, CA

Pierre Fraigniaud
 CNRS – Paris, FR &
 Université Paris Cité, FR

Krisztina Fruzsa TU Wien, AT

■ Murdoch James Gabbay Heriot-Watt University – Edinburgh, GB

Guy Goren
 Protocol Labs – Kibbutz
 Nahsholim, IL

Joseph Y. Halpern Cornell University – Ithaca, NY, US

Roman Kniazev
 École Polytechnique –
 Palaiseau, FR

Sophia Knight
 University of Minnesota –
 Duluth, MN, US

Roman Kuznets TU Wien, AT

Jérémy Ledent University of Strathclyde – Glasgow, GB

David Lehnherr Universität Bern, CH

Shihao (Jason) Liu
 University of Toronto, CA

■ Yoram Moses Technion – Haifa, IL

Susumu Nishimura Kyoto University, JP Thomas Nowak
ENS – Gif-sur-Yvette, FR
Ami Paz
CNRS – Gif-sur-Yvette, FR
Sergio Rajsbaum
National Autonomous University of Mexico, MX
Rojo Randrianomentsoa
TU Wien, AT
Hugo Rincón Galeana
TU Wien, AT

 David A. Rosenblueth
 National Autonomous University of Mexico, MX

Ulrich Schmid TU Wien, AT

Sonja Smets
 University of Amsterdam, NL

Thomas Studer Universität Bern, CH

Diego A. Velázquez National Autonomous University of Mexico, MX

