Combinatorial Algebraic Topology and applications to Distributed Computing

Dmitry Feichtner-Kozlov

ALTA
Institute for Algebra, Geometry, Topology and their applications

University of Bremen

Institute for Advanced Study, Princeton

2015

A primer in theoretical distributed computing

- **n processes performing a certain task**
- **choice of the model of communication**
  - message passing
  - read/write
  - immediate snapshot
  - read/write with other primitives
- **choice of possible failures**
  - crash failures
  - byzantine failures
- **choice of timing model**
  - asynchronous
  - synchronous
- **impossibility results**

One standard model of communication

- processes communicate via shared memory using two operations:

  (1) **`write`** - process writes its state to the devoted register;
  (2) **`snapshot read`** - process reads the entire shared memory in an atomic step;

- allow crash failures;

- all protocols are asynchronous and **`wait-free`**, meaning there is no upper bound on how long it takes to execute a step;

- all executions are **`immediate snapshot`**, meaning that at each step a group of processes becomes active, all processes in that group first write and then read together;

- all executions are **`layered`**, meaning that the execution is divided into rounds; in each round each non-faulty process reads and writes exactly once.

An example: assigning communication channels

- The processes use shared memory to coordinate among themselves the assignment of unique communication channels.

- One can show that two available communication channels are not enough for two processes.

- Two processes can pick unique communication channels among three available channels c_1, c_2, and c_3 using the following protocol:

  (1) the process writes its id into shared memory, then it reads the entire memory;

  (2) if the other process did not write its id yet - pick c_1;

  (3) if the other process wrote its id - compare the ids and pick c_2 or c_3.

Examples of tasks

- **Binary consensus:** Each process starts with a value 0 or 1 and eventually decides on 0 or 1 as its output value, subject to the following conditions:
  - the output value of all the processes should be the same;
  - if all the processes have the same initial value then they all must decide on that value.

- ***k*-set agreement:** Each process starts with an input value and eventually decides on an output value, subject to the following conditions:
  - at most k different values appear as output values of the processes;
  - only values which appear as input values are allowed to be picked as an output value.

- **Weak Symmetry Breaking:** Processes have no inputs *(an inputless task)* and need to pick outputs from $\{0, 1\}$, such that not all processes pick the same output; the processes are only allowed to compare their id's with each other.

Reference

M. Herlihy, D. Kozlov, S. Rajsbaum:
Distributed Computing Through Combinatorial Topology,
319 pp., Elsevier / Morgan Kaufmann, 2014.

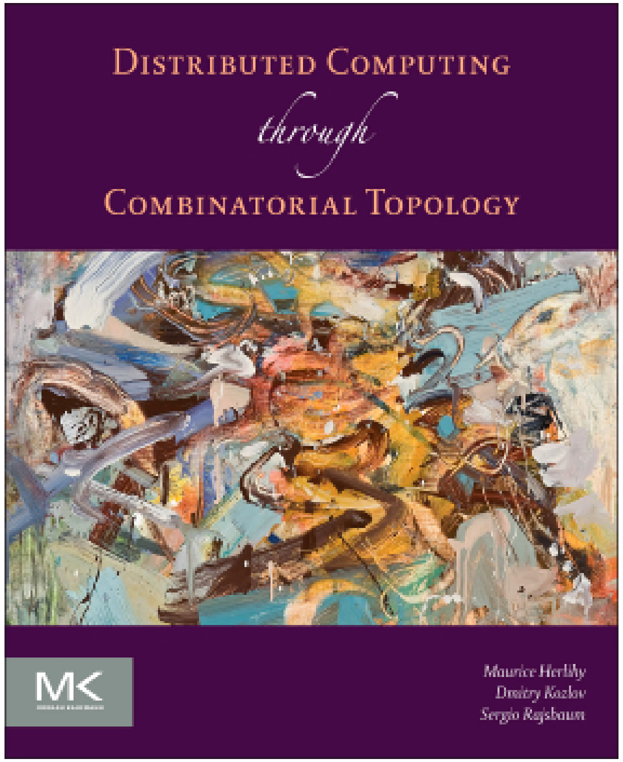

Introducing simplicial structure

- A simplicial complex of initial configurations: *input complex*
- A simplicial complex of final configurations: *output complex*
- A simplicial complex of all executions: *protocol complex*
- A carrier map from the input complex to the output complex: *task specification map*
- A carrier map from the input complex to the protocol complex: *execution map*
- A simplicial map from the protocol complex to the output complex: *decision map*
- A condition connecting task specification, decision, and execution maps

Binary consensus

Simplicial formulation for the binary consensus:

- The input complex $\mathcal{I}$ has vertices labeled (P, v), where $P \in [n]$, $v \in \{0, 1\}$. The vertices $(P_0, v_0), \ldots, (P_l, v_l)$ form a simplex iff $P_i \neq P_j$.
- The output complex $\mathcal{O}$ consists of two disjoint n-simplices; the first simplex has vertices labeled $(P, 0)$, and the second one has vertices labeled $(P, 1)$, for $P \in [n]$.
- The carrier map Δ is defined as follows. Take $\sigma = \{(P_0, v_0), \ldots, (P_l, v_l)\} \in \mathcal{I}$.
  - If $v_0 = \cdots = v_l = 0$, then $\Delta(\sigma)$ is the simplex $\{(P_0, 0), \ldots, (P_l, 0)\}$;
  - If $v_0 = \cdots = v_l = 1$, then $\Delta(\sigma)$ is the simplex $\{(P_0, 1), \ldots, (P_l, 1)\}$;
  - else $\Delta(\sigma)$ is the union of the simplices $\{(P_0, 0), \ldots, (P_l, 0)\}$ and $\{(P_0, 1), \ldots, (P_l, 1)\}$.

Simplicial formalization of other standard tasks

- **Consensus.**
- k**-set agreement.**
- **Weak Symmetry Breaking.**

Simplicial formalization: notion of a task

Definition.

A **task** is a triple $(\mathcal{I}, \mathcal{O}, \Delta)$, where

- $\mathcal{I}$ is a rigid chromatic *input complex* colored by $[n]$ and labeled by V^{in}, such that each vertex is uniquely defined by its color together with its label;
- $\mathcal{O}$ is a rigid chromatic *output complex* colored by $[n]$ and labeled by V^{out}, such that each vertex is uniquely defined by its color together with its label;
- Δ is a chromatic carrier map from $\mathcal{I}$ to $\mathcal{O}$.

A chromatic abstract simplicial complex (K, χ) is called **rigid chromatic** if K is pure of dimension n, and χ is an $(n+1)$-coloring.

Assume K and L are chromatic simplicial complexes, and $\mathcal{M} : K \to 2^L$ is a carrier map. We call $\mathcal{M}$ **chromatic** if for all $\sigma \in K$ the simplicial complex $\mathcal{M}(\sigma)$ is pure of dimension $\dim \sigma$, and we have $\chi_K(\sigma) = \chi_L(\mathcal{M}(\sigma))$;
where $\chi_L(\mathcal{M}(\sigma)) := \{\chi_L(v) \,|\, v \in V(\mathcal{M}(\sigma))\}$.

Simplicial formalization: protocol complexes

Definition.

A **protocol** for $n+1$ processes is a triple $(\mathcal{I}, \mathcal{P}, \mathcal{E})$ where

- $\mathcal{I}$ is a rigid chromatic simplicial complex colored with elements of $[n]$ and labeled with V^{in}, such that each vertex is uniquely defined by (color,label);
- $\mathcal{P}$ is a rigid chromatic simplicial complex colored with elements of $[n]$ and labeled with Views_{n+1}, such that each vertex is uniquely defined by (color,label);
- $\mathcal{E} : \mathcal{I} \to 2^{\mathcal{P}}$ is a chromatic intersection-preserving carrier map, such that $\mathcal{P} = \cup_{\sigma \in \mathcal{I}} \mathcal{E}(\sigma)$, where intersection-preserving means that for all $\sigma, \tau \in \mathcal{I}$ we have

$$\mathcal{E}(\sigma \cap \tau) = \mathcal{E}(\sigma) \cap \mathcal{E}(\tau).$$

Definition.

Assume we are given a task $(\mathcal{I}, \mathcal{O}, \Delta)$ for $n+1$ processes, and a protocol $(\mathcal{I}, \mathcal{P}, \mathcal{E})$. We say that this protocol **solves** this task if there exists a chromatic simplicial map $\delta : \mathcal{P} \to \mathcal{O}$, called **decision map**, satisfying

$$\delta(\mathcal{E}(\sigma)) \subseteq \Delta(\sigma),$$

for all $\sigma \in \mathcal{I}$.

Topological characterization of solvability of colored tasks

Theorem. **Herlihy et al.**

A task $(\mathcal{I}, \mathcal{O}, \Delta)$ has a wait-free layered immediate snapshot protocol if and only if there exists t and a (color-preserving) simplicial map

$$\mu : \mathrm{Ch}^{t}(\mathcal{I}) \to \mathcal{O}$$

carried by Δ.

Sperner's Lemma.

Assume we are given

- a simplicial subdivision of Δ^n, called $\mathrm{Div}\,\Delta^n$;
- an $[n]$-coloring of the vertices of $\mathrm{Div}\,\Delta^n$, such that for every vertex v only the colors of the carrier of v may be used.

Then $\mathrm{Div}\,\Delta^n$ has a properly colored n-simplex.

Sperner's Lemma $\Rightarrow$ k-set agreement is not solvable.

Chromatic subdivision of a simplex

Example.

The protocol complex in the layered immediate snapshot model, where each of the $n+1$ processes acts exactly once, is a chromatic subdivision of an n-simplex, which we call the **standard chromatic subdivision.**
Notation: $\mathrm{Ch}\,\Delta^n$.

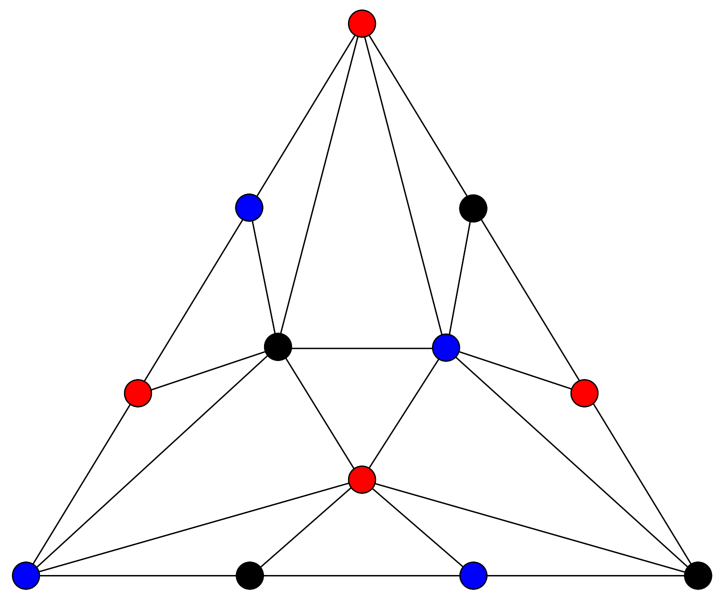

The protocol complex of the standard 1-round wait-free layered immediate snapshot protocol for 3 processes.

Topology of the layered immediate snapshot protocol complexes

Theorem.

Protocol complexes for the standard 1-round wait-free layered immediate snapshot protocols for $n+1$ processes are equivariantly collapsible chromatic subdivisions of an n-simplex.

Notation: $\mathrm{Ch}\,\Delta^n$.

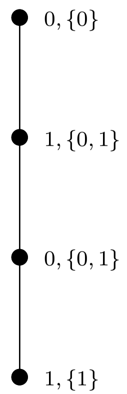

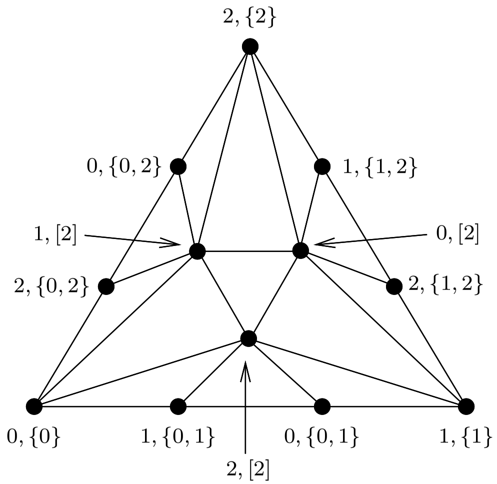

The protocol complexes of the standard 1-round wait-free layered immediate snapshot protocols for 2 and 3 processes.

Topology of the immediate snapshot protocol complexes

Theorem.

Protocol complexes for wait-free IS protocols for $n+1$ processes for an arbitrary round counter are homeomorphic to an n-simplex.

Conjecture.

They are subdivisions of an n-simplex.

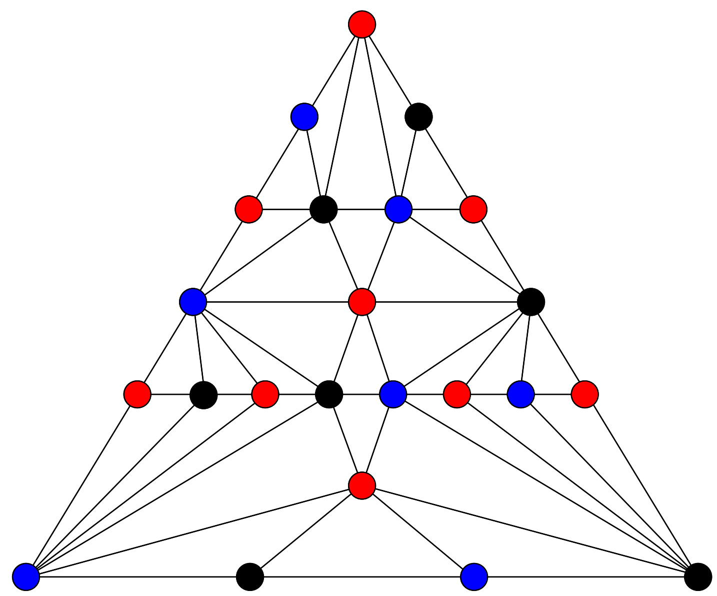

The protocol complex of a wait-free IS protocol for 3 processes.

Topology of the view complex

Theorem.

Protocol complexes for 1-round wait-free read-write protocols of $n+1$ processes are equivariantly collapsible to the iterated standard chromatic subdivision of an n-simplex.

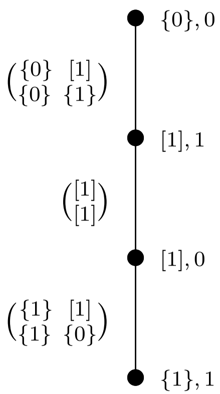

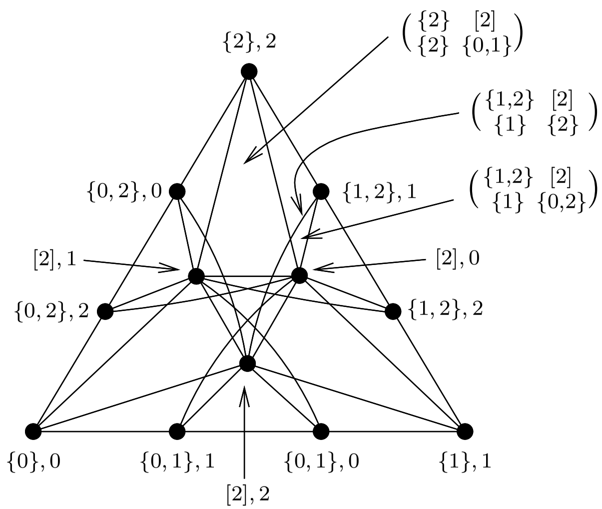

The view complexes for 2 and 3 processes.

Rank-Symmetric Protocols for Weak Symmetry Breaking

Weak Symmetry Breaking is a standard task specified by the following:

- The processes have no input values.
- Their output values are 0 and 1.
- The task is solvable if there exists an IS wait-free protocol such that in every execution in which all processes participate, not all processes decide on the same value.

In addition, this protocol is required to be *rank-symmetric*.

Definition.

Let $I_1, I_2 \subseteq \{0, \ldots, n\}$, such that $|I_1| = |I_2|$, and let $r : I_1 \to I_2$ be the order-preserving bijection. Assume E_1 is an execution of the protocol, in which the set of participating processes is I_1, and E_2 is an execution of the protocol, in which the set of participating processes is I_2. Then, the protocol is called **rank-symmetric** if for every $i \in I_1$, the process with ID i must decide on the same value as the process with ID $r(i)$. The protocol will certainly be rank-symmetric, if each process only compares its ID to the ID's of the other participating processes, and makes the final decision based on the relative rank of its ID.

Index Lemma I

Let M be an orientable n-dimensional pseudomanifold with boundary ∂M, and let c be an $[n]$-coloring of the vertices of M. We say that a simplex is *properly colored* if all colors of its vertices are different.
We want to *count simplices by orientation*, to which end we set

$$C(\sigma, c) := \begin{cases} 1, & \text{if } \sigma \text{ is properly colored and its color orientation is same as induced by} \\ & \text{the manifold orientation;} \\ -1, & \text{if } \sigma \text{ is properly colored and its color orientation is the opposite of} \\ & \text{the one induced by the manifold orientation;} \\ 0, & \text{if } \sigma \text{ is not properly colored.} \end{cases}$$

Definition.

The **content of** c in M is the number of simplices properly colored by c, counted by orientation:

$$C(M, c) := \sum_{\sigma \in M,\ \dim \sigma = n} C(\sigma, c).$$

Index Lemma II

Definition.

For each $i \in [n]$, the ***i*th index** of M is the number of $(n-1)$-simplices of ∂M, properly colored with values from $\mathrm{Face}_i(\Delta^n)$, counted by orientation:

$$I_i(M, c) := \sum_{\tau \in \partial M,\ \dim \tau = n-1} C(\tau, c).$$

Index Lemma.

If M is an oriented pseudomanifold colored by c, then for each $i \in [n]$ we have the identity

$$C(M, c) = (-1)^i I_i(M, c).$$

WSB when the number of processes is a prime power

Herlihy-Shavit Theorem.

WSB is not solvable when the number of processes is a prime power.

Indeed, if WSB is solvable, then there exists a compliant binary coloring, hence there is an $(n+1)$-coloring of $\mathrm{Ch}^r\Delta^n$ with no monochromatic simplices.

On the other hand, one can show that there exist integers $k_0, k_1, \ldots, k_{n-1}$, such that the number of monochromatic simplices in $\mathrm{Ch}^r\Delta^n$, counted by orientation, is

$$(-1)^n + \sum_{q=0}^{n-1} k_q \binom{n+1}{q+1}.$$

The contradiction follows from the fact that the binomial coefficients $\binom{n+1}{1}, \ldots, \binom{n+1}{n}$ are all divisible by a prime p, if $n+1$ is a power of p.

WSB when the number of processes is not a prime power

Castaneda - Rajsbaum Theorem.

WSB is solvable when the number of processes is not a prime power.

The idea is to start with a coloring with 0 content. This means that the number of monochromatic n-simplices is zero, counted with orientation. Then one links simplices of opposite orientation by simplex paths and uses a quite sophisticated *path subdivision algorithm* to eliminate the monochromatic pair.

Protocols solving WSB in a few rounds:

Theorem.

There exists a comparison-based layered immediate snapshot distributed protocol solving the Weak Symmetry Breaking task for $6t$ processes in 3 rounds.

The proof is based on in-depth analysis of the combinatorial structure of the iterated chromatic subdivisions of an $(n-1)$-simplex.