Simplicial models in Distributed Computing

Dmitry Kozlov

CALTOP

ALTA
Institute for Algebra, Geometry, Topology and their applications

University of Bremen

Paris, May 2022

What is . . .

Combinatorial Algebraic Topology ?

Combinatorial Algebraic Topology is about computing invariants in Algebraic Topology

- for **combinatorial cell complexes**
- by **combinatorial means.**

Combinatorial cell complexes

locally = simple cell attachments

for example - simplicial or prodsimplicial complexes;

globally = cells are indexed by combinatorial objects

such as graphs, partitions, permutations, various combinations and enrichments of these;

taking boundary = combinatorial rule for the indexing objects

for example - removing vertices in graphs, merging blocks in partitions, relabeling, etc.

Combinatorial means

using matchings, orderings, labelings, et cetera, to simplify or to completely eliminate algebraic computations and topological deformations.

Book Reference

D. Kozlov: *Combinatorial Algebraic Topology*,
405 pp., Springer-Verlag, 2008.

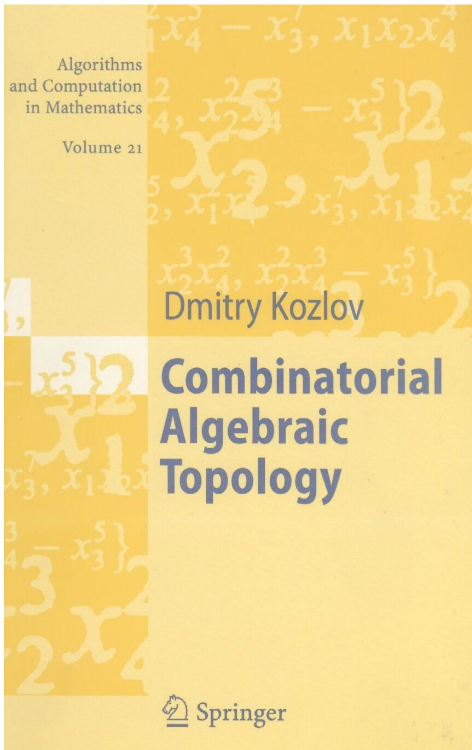

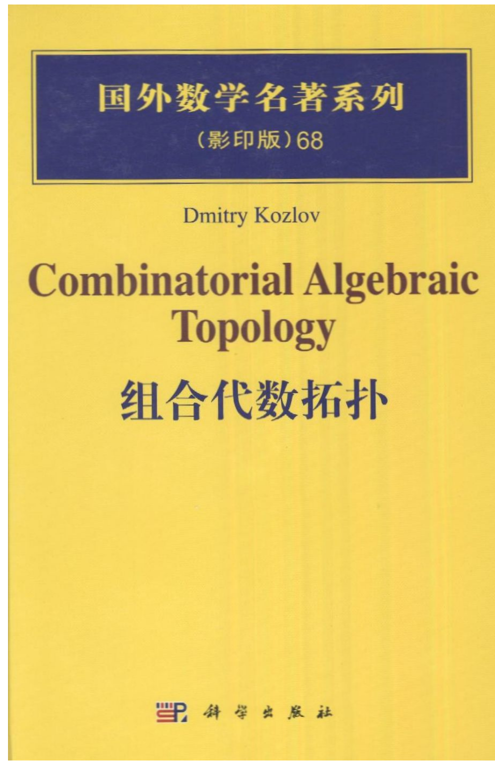

On to Distributed Computing !

The setting

- **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.

Further 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.

Think simplicially !

Introducing simplicial structure in distributed computing

- 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 $\{(P_0, 1), \ldots, (P_l, 1)\}$.

Standard Chromatic Subdivision

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

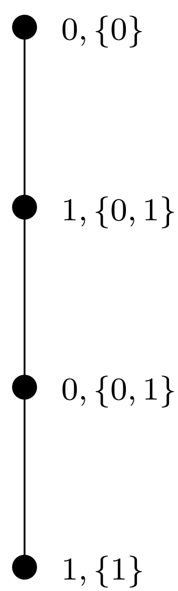

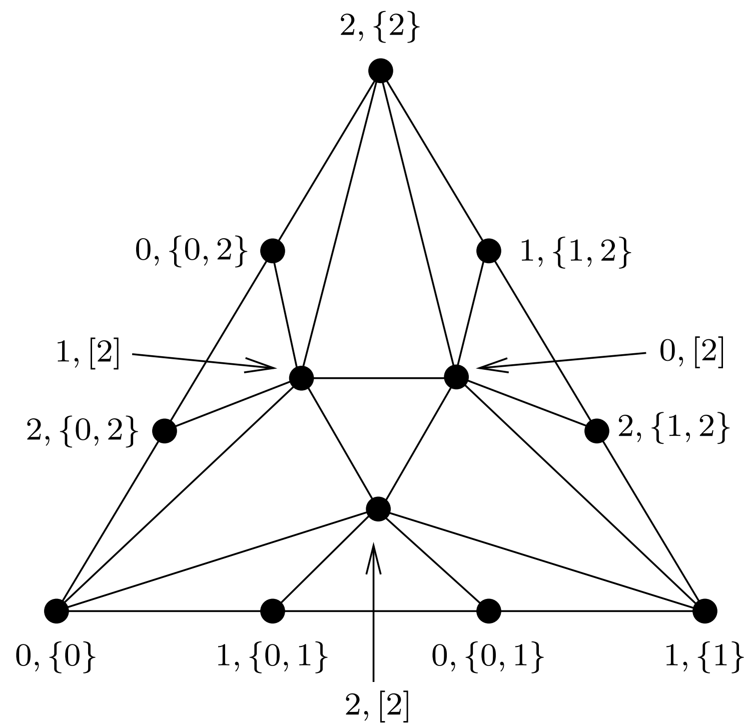

The protocol complexes of the standard 1-round wait-free layered immediate snapshot protocols for 2 and 3 processes.

Simplicial formalization: notion of a task

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 $\chi_L(\mathcal{M}(\sigma)) := \{\chi_L(v) \,|\, v \in V(\mathcal{M}(\sigma))\}$.

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 $\mathcal{O}$.

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

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

Simplicial formalization: protocol solves a task

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

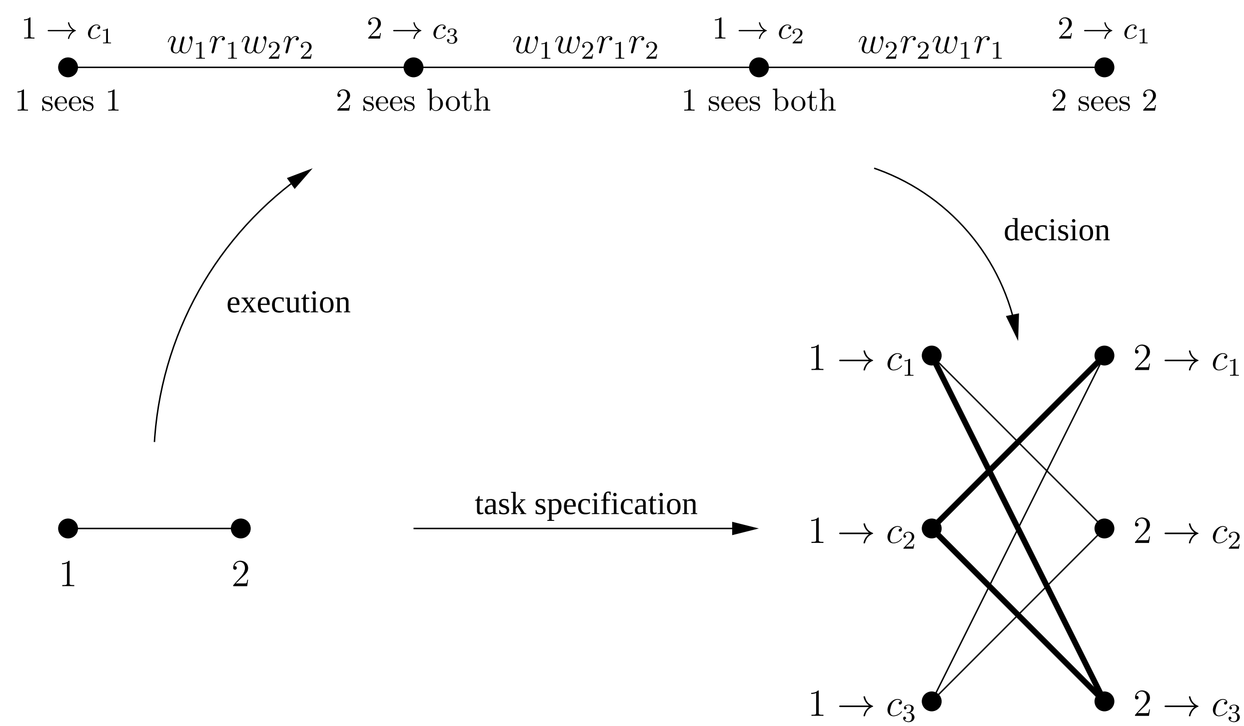

The channel assignment for 2 processes viewed simplicially.

Book Reference

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

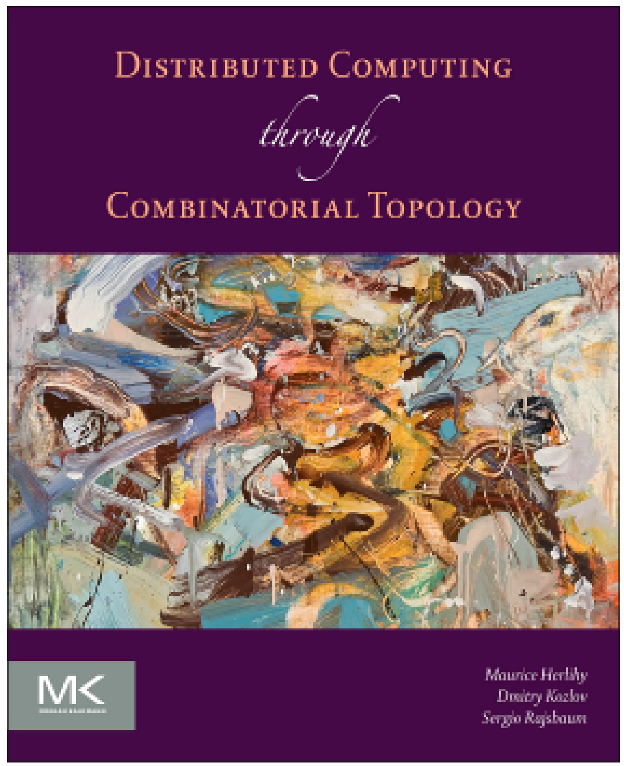

SUMMARY

- Simplicial methods are useful in theoretical distributed computing.
- Recent work provides rigorous foundation for this interdisciplinary area.
- Protocol complexes are configuration spaces encoding all possible executions of protocols.
- Combinatorial topology of protocol complexes is sophisticated, contains answers to distributed computing questions, depends on the communication model.
- Study of combinatorial topology of protocol complexes leads to new and interesting mathematics

Mathematics of Weak Symmetry Breaking

Standard chromatic subdivision I

For all $n \geq 0$, we can define a simplicial subdivision of an n-simplex, which we call the **standard chromatic subdivision.**

Notation: $\chi(\Delta^n)$.

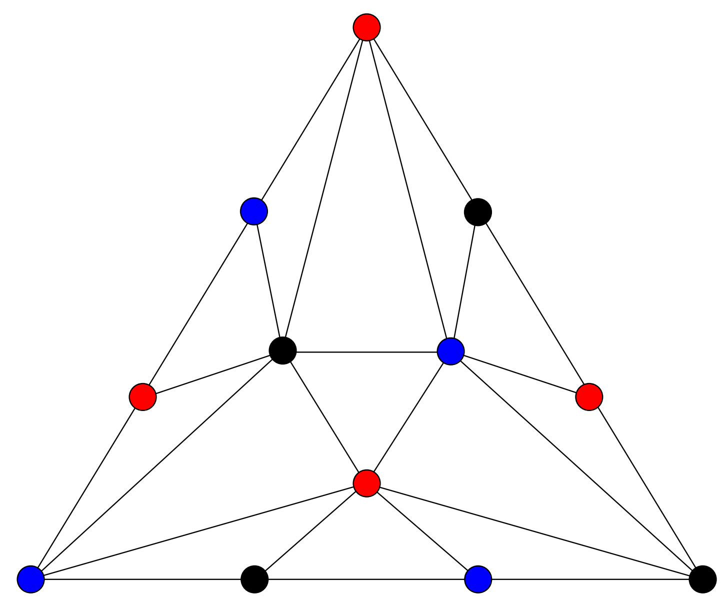

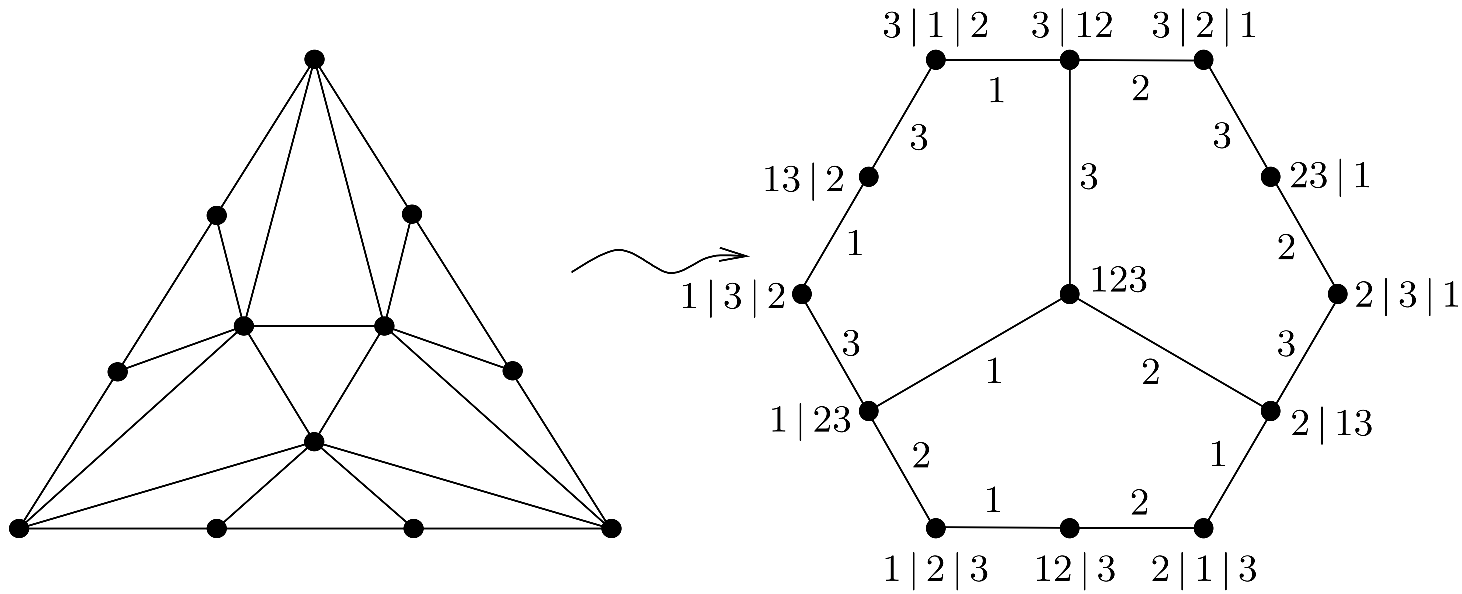

The standard chromatic subdivision and the flip graph.

Standard chromatic subdivision II

The construction χ has the following properties:

- χ can be applied to simplicial complexes;
- χ can be iterated;
- for all $I, J \subseteq [n]$, such that $|I| = |J|$, the linear isomorphism $\varphi_{I,J}$ between Δ^I and Δ^J is also an isomorphism of the subdivision restrictions $\chi(\Delta^I)$ and $\chi(\Delta^J)$.

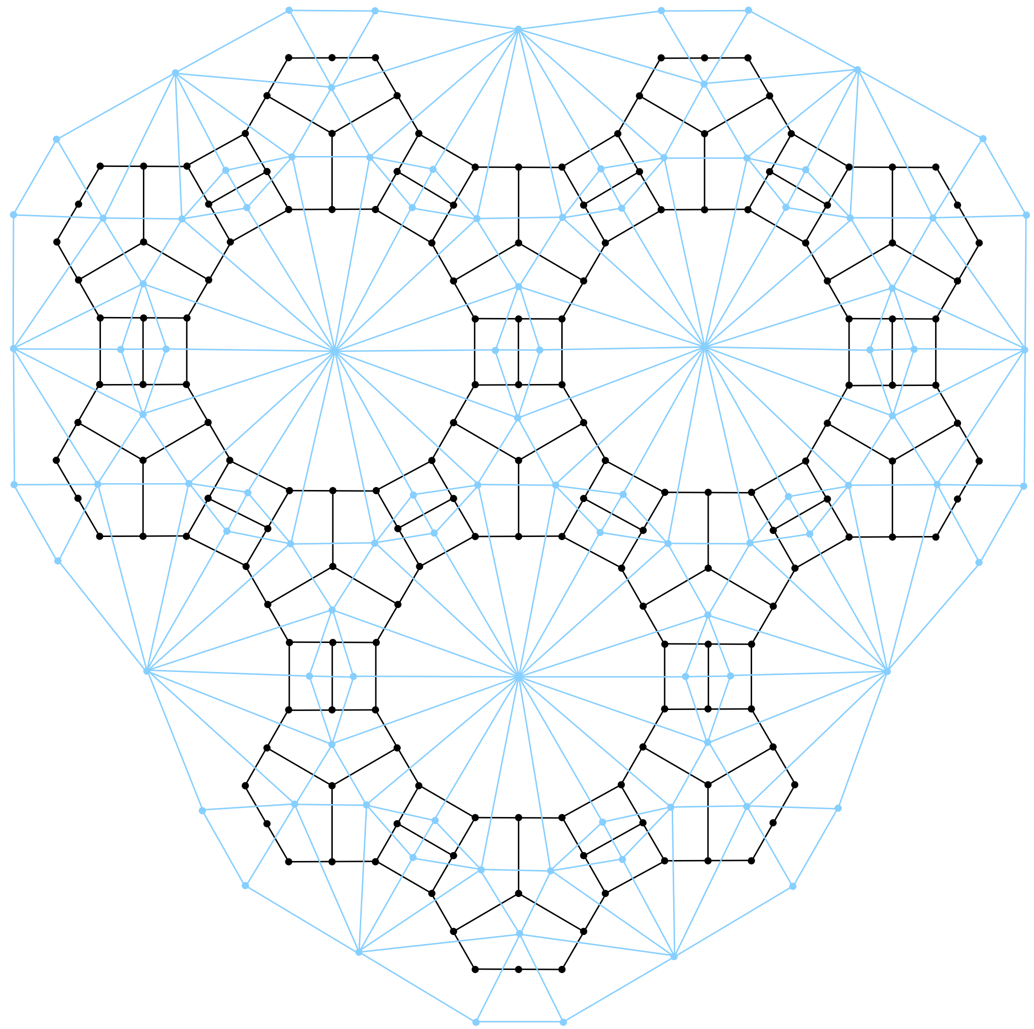

The flip graph Γ_3^2.
In the background: the second standard chromatic subdivision of a triangle.

Standard chromatic subdivision III

A labeling $\lambda : V(\chi^t(\Delta^n)) \to L$, where L is an arbitrary set of labels, is called **compliant** if for all $I, J \subseteq [n]$ such that $|I| = |J|$, and all vertices $v \in V(\chi^t(\Delta^I))$, we have

$$\lambda(\varphi_{I,J}(v)) = \lambda(v).$$

Main goal:

Find a binary compliant labeling $\lambda : V(\chi^t(\Delta^n)) \to \{0, 1\}$, so that no n-simplex of $\chi^t(\Delta^n)$ is monochromatic.

When this is possible we shall say that $\chi^t(\Delta^n)$ has a **symmetry breaking** coloring.

Symmetry breaking when $n+1=p^\alpha$

Herlihy-Shavit Theorem.

Assume the number of vertices in the original simplex is a prime power p^α.
For any $t \geq 1$, there does **not** exist a symmetry breaking coloring on the set of vertices of $\chi^t(\Delta^n)$.

The idea is to show that there exist integers $k_1, k_2, \ldots, k_n$, such that the number of monochromatic n-simplices in $\chi^t(\Delta^n)$, counted by orientation, is

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

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.

Symmetry breaking when $n+1$ is not a prime power

Castañeda - Rajsbaum Theorem.

When $n+1$ is not a primer power, there exists $t \geq 1$, such that $\chi^t(\Delta^n)$ has a symmetry breaking coloring.

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.

Minimizing the number of iterations I

It is easy to see that if $\chi^t(\Delta^n)$ has a symmetry breaking coloring, then $\chi^{t+1}(\Delta^n)$ has one as well.

Main question.

What is the minimal value of t such that $\chi^t(\Delta^n)$ has a symmetry breaking coloring?

Let $\varphi(n)$ denote the minimal value t, such that there exists a symmetry breaking coloring on $\chi^t(\Delta^{n-1})$. It is not difficult to show that $\varphi(n) \geq 2$.

Previously known upper bounds (due to Atiyah, Castañeda, Herlihy, Paz):

- $\varphi(n) = O(n^{q+5})$, where q is the largest prime power dividing n,
- $\varphi(6) \leq 17$.

Minimizing the number of iterations II

Theorem. (DK'16)

Assume that we have a binomial identity

$$\binom{n}{0} + \binom{n}{a_1} + \cdots + \binom{n}{a_k} = \binom{n}{1} + \binom{n}{b_1} + \cdots + \binom{n}{b_m},$$

where $2 < a_1 < \cdots < a_k < n$, $1 < b_1 < \cdots < b_m < n$, and $a_i \neq b_j$, for all i, j. Then there exists symmetry breaking coloring on the vertices of $\chi^3(\Delta^{n-1})$, i.e., $\varphi(n) \leq 3$.

Examples:

$$\binom{6}{0} + \binom{6}{3} = \binom{6}{1} + \binom{6}{4}$$

$$\binom{15}{0} + \binom{15}{4} + \binom{15}{6} + \binom{15}{13} = \binom{15}{1} + \binom{15}{3} + \binom{15}{5} + \binom{15}{10}$$

In particular: **there exists infinitely many values of n for which $\varphi(n) \leq 3$.**

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

Constructing paths I

Plan of the proof.

- Start with the binomial identity and use the Sperner theory to find a bijection between two families of subsets of $[n]$, such that one of the matched sets contains the other one.
- Find a familiy of non-intersecting simplex paths in $\chi(\Delta^{n-1})$ corresponding to that matching.
- Construct a certain binary assignment for the vertices on the boundary of $\chi^2(\Delta^{n-1})$. That assignment depends on the binomial identity and makes sure the total assignment is compliant.

Constructing paths II

- Assign value 0 to all internal vertices of $\chi^2(\Delta^{n-1})$. This will produce a lot of monochromatic n-simplices.

- We now look for a perfect matching on the set of monochromatic n-simplices. We start with a standard one matching almost all n-simplices to each other. The rest of the n-simplices are the *critical* ones.

- The technical core of the proof is constructing the extension of that standard matching to include the critical simplices.

- When a complete matching is found it can be used to find symmetry breaking coloring after just one further subdivision.

Constructing paths III

A standard technique of extending matchings: **augmenting paths**.

We link such local paths across the patchwork of simplices to produce global augmenting paths.

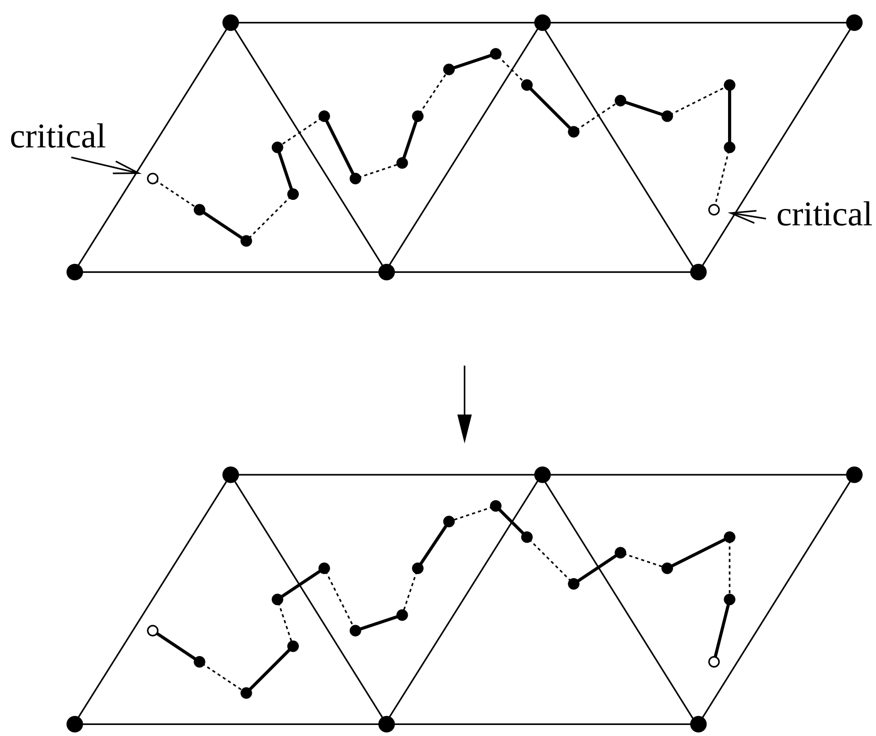

Concatenating augmenting paths.

Perspectives

- It is unclear at the moment for which n such a binomial identity exists.
- We also do not know whether 3 is a sharp bound.
- What is the bound for other n?

SUMMARY

- The study of the solvability of Weak Symmetry Breaking task leads to the rich theory of subdivisions of the simplicial paths where much structure remains mysterious

- Simplicial methods can be used to design fast protocols solving Weak Symmetry Breaking for some values of n

- The explicit protocol solving Weak Symmetry Breaking in general (when it is solvable) remains out of reach