Topological insights on probabilistic agreement

Jules Chouquet

IRIF, Université de Paris

1 Introduction

2 Combinatorial topology

3 Beyond impossibility and determinism

4 The lower bound approach

Trailer

What this talk is about

- Asynchronous computability
- Shared memory
- Fault tolerance
- Impossibility results
- Combinatorial topology
- Randomization

Contribution: use topology to get probability lower bounds on randomized agreement protocols.

Wait-free Asynchronous computability

A will: Compute things with many agents.

Motivations: Efficiency, multiprocessor architectures, economy of energy, networks, IoT,...

Two main paradigms for communication: Shared memory and message-passing.

Difficulties: Asynchrony and Fault tolerance.

In this talk

We consider any number of possible crashes (wait-freedom), and consider only the shared memory model (keeping in mind that there are translations), in order to use topological interpretation.

Definition (Binary consensus)

Specification:

- Each process starts with an initial value 0 or 1
- At the end, every process outputs the same value
- The output value must be one of the inputs of some participating process

Definition (k-set agreement)

Specification:

- Each process starts with an initial proper value (n distinct inputs)
- At the end, there is no more than k distinct outputs
- Thue output values must be among the inputs of the participating process

Immediate atomic snapshot protocols

The model : a set of n processes $p_0, \ldots, p_{n-1}$, with variables for local computations, and SWMR registers r_i for each p_i.

Operations for p_i :

- *update* r_i (u_i)
- *snapshot* (s_i)
- other... (specific to the algorithm we consider)

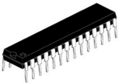

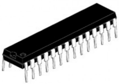

Immediate atomic snapshot protocols

The model : a set of n processes $p_0, \ldots, p_{n-1}$, with variables for local computations, and SWMR registers r_i for each p_i.

Operations for p_i :

- *update* r_i (u_i)
- *snapshot* (s_i)
- other... (specific to the algorithm we consider)

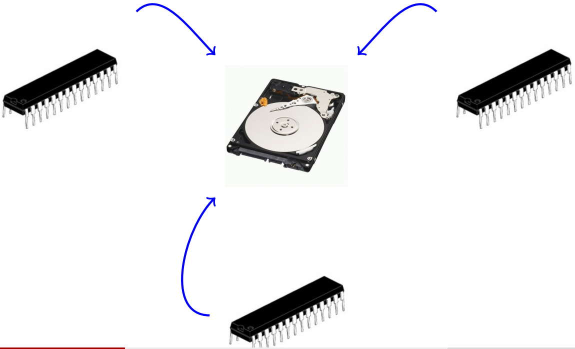

Immediate atomic snapshot protocols

The model : a set of n processes $p_0, \ldots, p_{n-1}$, with variables for local computations, and SWMR registers r_i for each p_i.

Operations for p_i :

- *update* r_i (u_i)
- *snapshot* (s_i)
- other... (specific to the algorithm we consider)

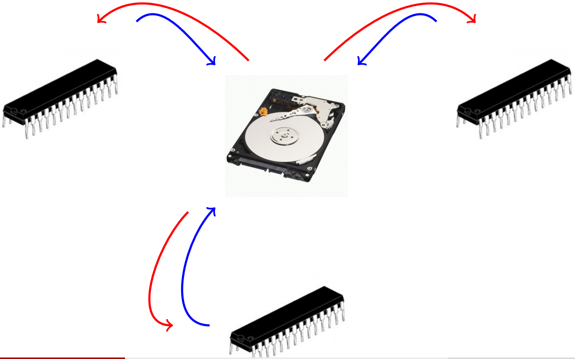

Immediate atomic snapshot protocols

The model : a set of n processes $p_0, \ldots, p_{n-1}$, with variables for local computations, and SWMR registers r_i for each p_i.

Operations for p_i :

- *update* r_i (u_i)
- *snapshot* (s_i)
- other. . . (specific to the algorithm we consider)

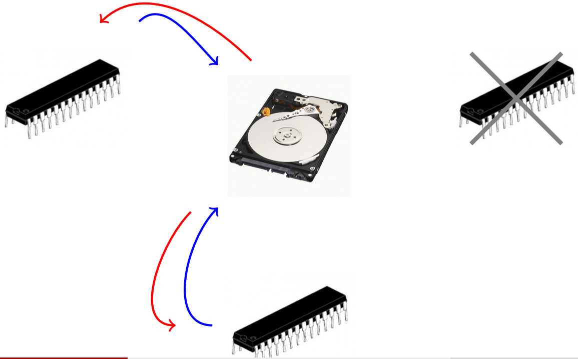

Executions as words (1)

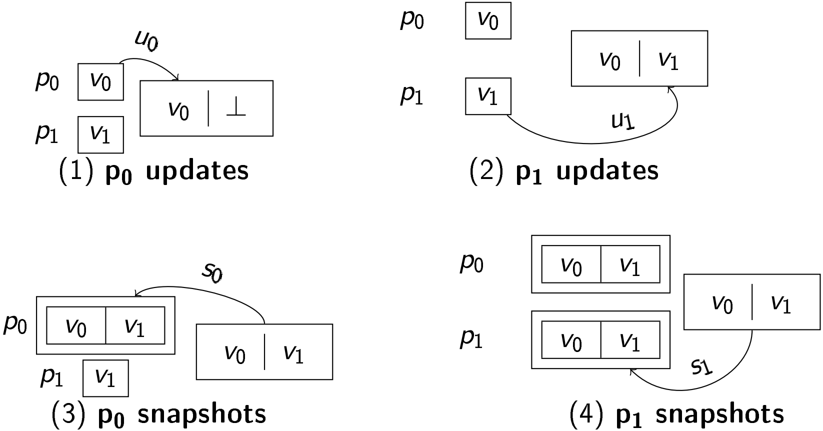

Figure: Execution trace on word $u_0 u_1 s_0 s_1$

Executions as words (2)

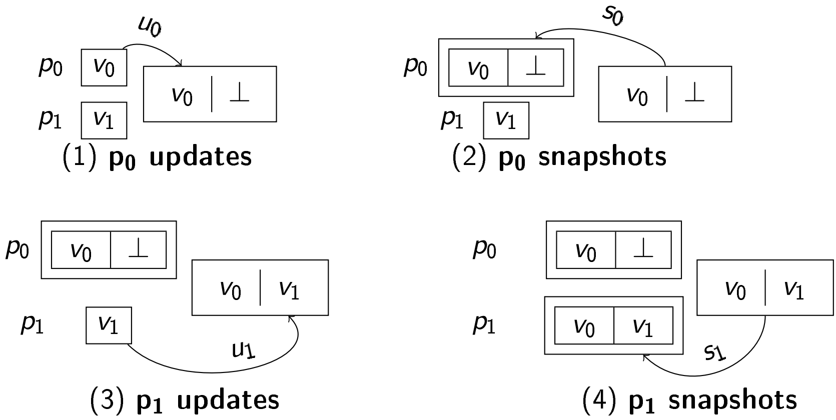

Figure: Execution trace on word $u_0 s_0 u_1 s_1$

Executions as words (3)

Final states :

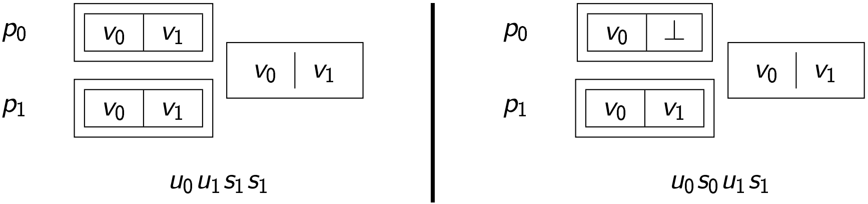

Equivalence on executions

Examples

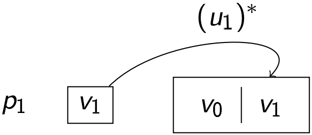

$u_i u_i \approx u_i$

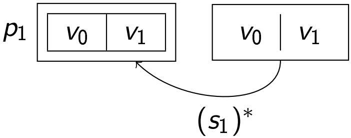

$s_i s_i \approx s_i$

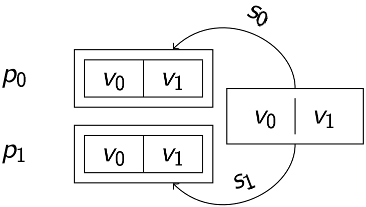

$s_i s_j \approx s_j s_i$

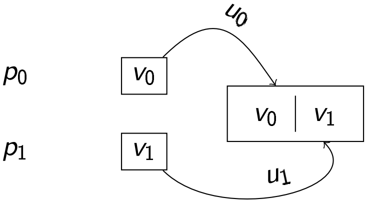

$u_i u_j \approx u_j u_i$

1 Introduction

2 Combinatorial topology

3 Beyond impossibility and determinism

4 The lower bound approach

Simplicial complexes

Definition (Abstract simplicial complexes)

Let S a set, and C a family of subsets of S. C is an *Abstract Simplicial Complex over* S if :

- When $\sigma \in C$ and $\tau \subseteq \sigma$, $\tau \in C$
- For all $x \in S$, $\{x\} \in C$.

Elements of C are called *simplices* (and we write $|C|$ for S).

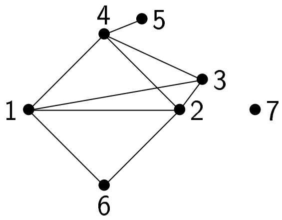

$\{1,2,3,4\}$, $\{1,2,6\}$, $\{4,5\}$, $\{7\}$

(and all subsets)

If σ is a simplex, its dimension is $\mathbf{card}(\sigma) - 1$

Simplicial complexes

Definition (Abstract simplicial complexes)

Let S a set, and C a family of subsets of S. C is an *Abstract Simplicial Complex over* S if :

- When $\sigma \in C$ and $\tau \subseteq \sigma$, $\tau \in C$
- For all $x \in S$, $\{x\} \in C$.

Elements of C are called *simplices* (and we write $|C|$ for S).

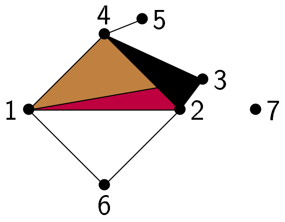

$\{1,2,3,4\}$, $\{1,2,6\}$, $\{4,5\}$, $\{7\}$

(and all subsets)

If σ is a simplex, its dimension is $\mathbf{card}(\sigma) - 1$

Simplicial complexes

Definition (Abstract simplicial complexes)

Let S a set, and C a family of subsets of S. C is an *Abstract Simplicial Complex over* S if :

- When $\sigma \in C$ and $\tau \subseteq \sigma$, $\tau \in C$
- For all $x \in S$, $\{x\} \in C$.

Elements of C are called *simplices* (and we write $|C|$ for S).

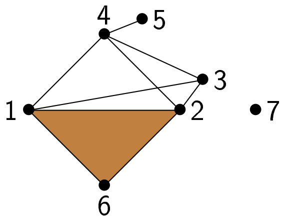

$\{1,2,3,4\}$, $\{1,2,6\}$, $\{4,5\}$, $\{7\}$

(and all subsets)

If σ is a simplex, its dimension is $\mathbf{card}(\sigma) - 1$

Simplicial complexes

Definition (Abstract simplicial complexes)

Let S a set, and C a family of subsets of S. C is an *Abstract Simplicial Complex over S* if :

- When $\sigma \in C$ and $\tau \subseteq \sigma$, $\tau \in C$
- For all $x \in S$, $\{x\} \in C$.

Elements of C are called *simplices* (and we write $|C|$ for S).

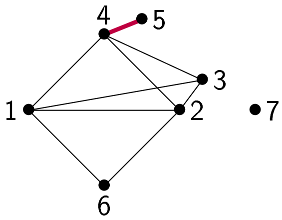

$\{1,2,3,4\}$, $\{1,2,6\}$, $\{4,5\}$, $\{7\}$

(and all subsets)

If σ is a simplex, its dimension is $\mathbf{card}(\sigma) - 1$

Simplicial complexes

Definition (Abstract simplicial complexes)

Let S a set, and C a family of subsets of S. C is an *Abstract Simplicial Complex over* S if :

- When $\sigma \in C$ and $\tau \subseteq \sigma$, $\tau \in C$
- For all $x \in S$, $\{x\} \in C$.

Elements of C are called *simplices* (and we write $|C|$ for S).

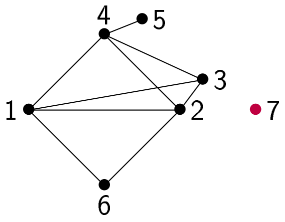

$\{1, 2, 3, 4\}$, $\{1, 2, 6\}$, $\{4, 5\}$, $\{7\}$

(and all subsets)

If σ is a simplex, its dimension is $\mathbf{card}(\sigma) - 1$

Simplicial maps

Definition (Simplicial map)

$f : |C| \to |D|$ is a *simplicial map* if for any simplex $\{x_1, \ldots, x_n\} \in C$, $\{f(x_1), \ldots, f(x_n)\} \in D$.

Simplicial maps

Definition (Simplicial map)

$f : |C| \to |D|$ is a *simplicial map* if for any simplex $\{x_1, \ldots, x_n\} \in C$, $\{f(x_1), \ldots, f(x_n)\} \in D$.

These maps preserve topological properties. For example,

neither

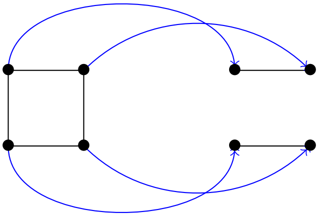

nor

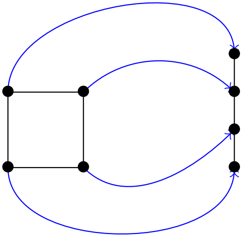

is a simplicial map

Simplicial maps

Definition (Simplicial map)

$f : |C| \to |D|$ is a *simplicial map* if for any simplex $\{x_1, \ldots, x_n\} \in C$, $\{f(x_1), \ldots, f(x_n)\} \in D$.

These maps preserve topological properties. For example,

neither

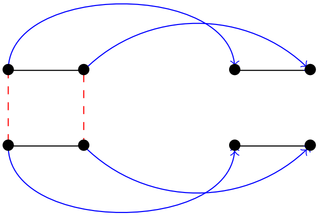

nor

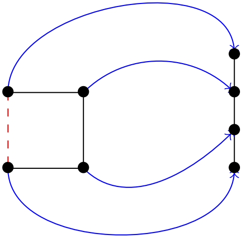

is a simplicial map

A simplicial complex that represents all the possible initial configurations.
Example: binary consensus for $n = 1$

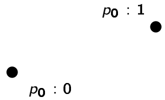

Representing tasks

Input complex

A simplicial complex that represents all the possible initial configurations.
Example: binary consensus for $n = 2$

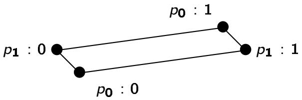

A simplicial complex that represents all the possible initial configurations.
Example: binary consensus for $n = 3$

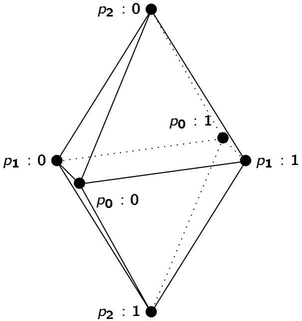

A simplicial complex that represents all the possible final configurations.
Example: binary consensus for $n = 1$

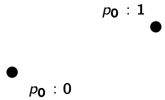

A simplicial complex that represents all the possible final configurations.
Example: binary consensus for $n = 2$

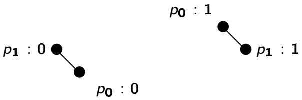

A simplicial complex that represents all the possible final configurations.
Example: binary consensus for $n = 3$

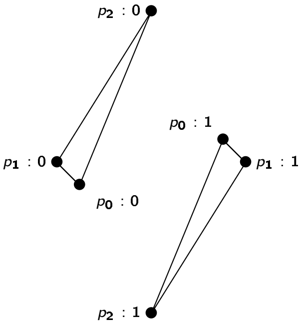

What's next ?

- What is the link between initial and final configurations ?

  $\rightarrow$ Executions
- What is the topological representation of an execution ?

  $\rightarrow$ A specific simplicial complex
- Why is it interesting ?
  - $\rightarrow$ This representation captures exactly the observational equivalence.
  - $\rightarrow$ It brings new techniques for proving impossibility results.
  - $\rightarrow$ It allows a precise quantitative study of the executions, that can be useful in a probabilistic approach.

Executions as simplices

The protocol complex (dimension 1)

Executions as simplices

The protocol complex (dimension 1)

Executions as simplices

The protocol complex (dimension 1)

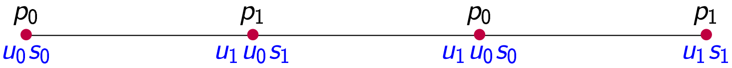

Executions as simplices

The protocol complex (dimension 1)

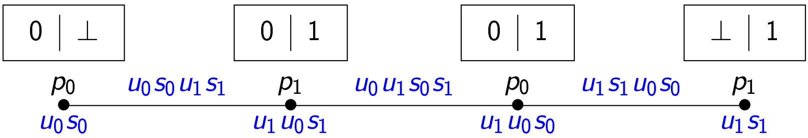

Executions as simplices

The protocol complex (dimension 1)

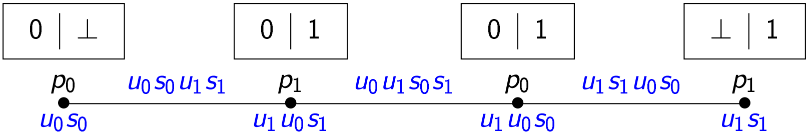

The protocol complex for dimension d is defined as the subdivision of the d-simplex.

Executions as simplices

The protocol complex (dimension 1)

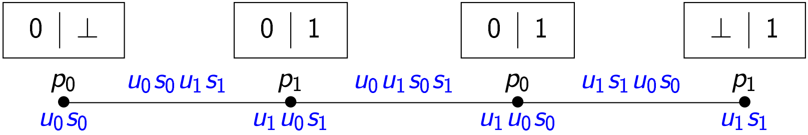

The protocol complex for dimension d is defined as the subdivision of the d-simplex.

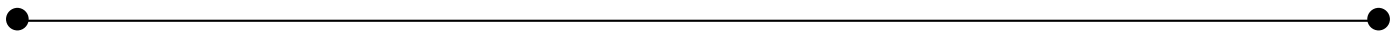

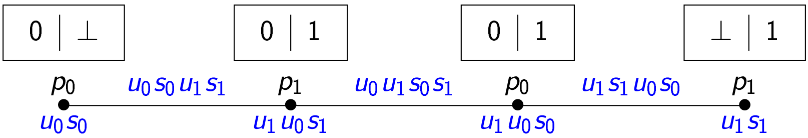

The protocol complex for dimension d is defined as the subdivision of the d-simplex.

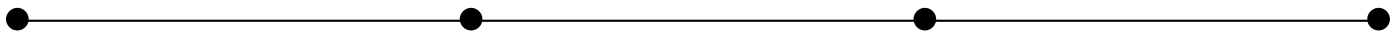

Executions as simplices

The protocol complex (dimension 1)

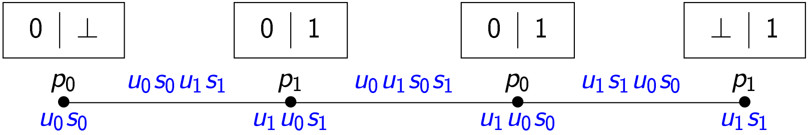

The protocol complex for dimension d is defined as the subdivision of the d-simplex.

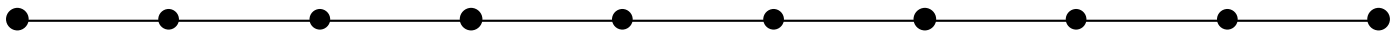

Executions as simplices

The protocol complex (dimension 1)

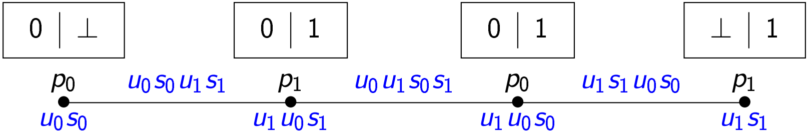

The protocol complex for dimension d is defined as the subdivision of the d-simplex.

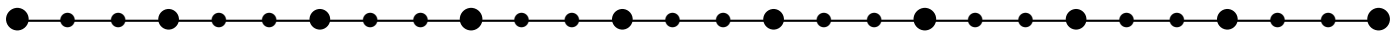

Protocol complex, (dimension 2)

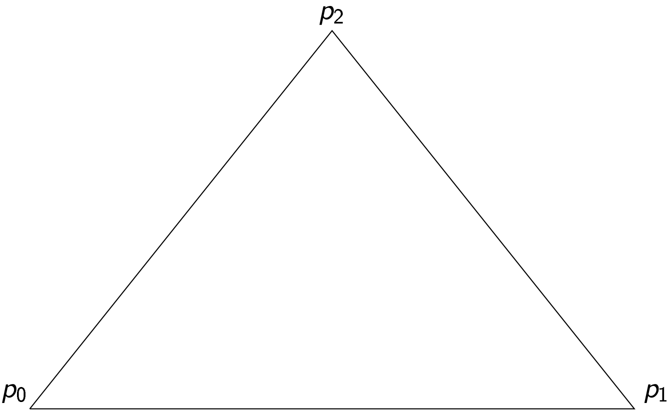

Protocol complex, (dimension 2)

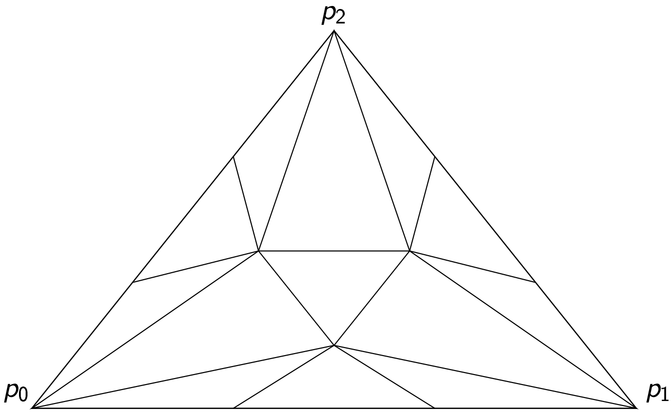

Protocol complex, (dimension 2)

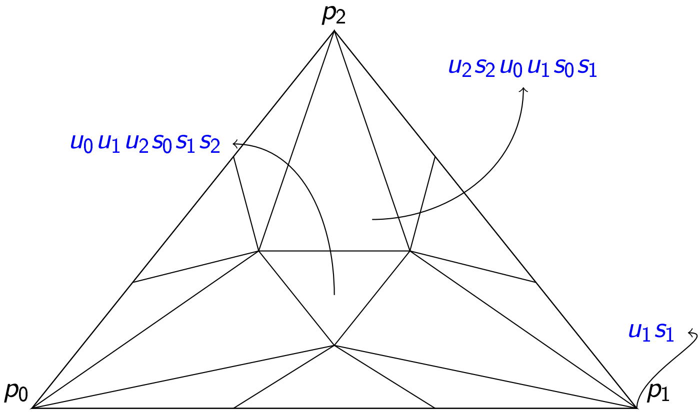

Protocol complex, (dimension 2)

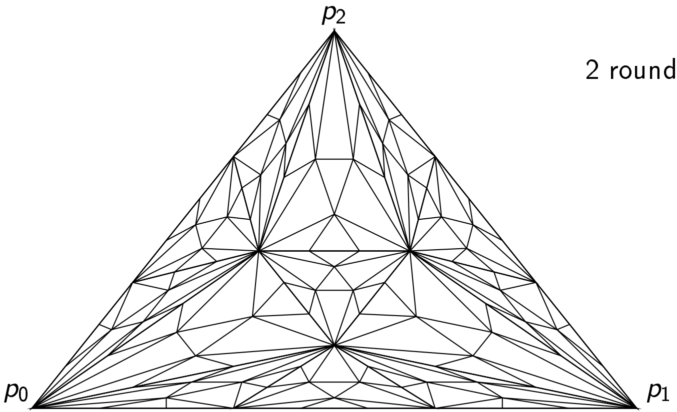

2 rounds

Protocol complex, (dimension 2)

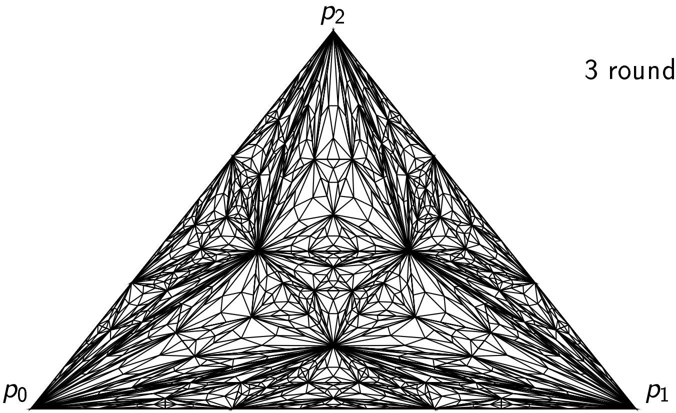

3 rounds

Protocol complex, (dimension 2)

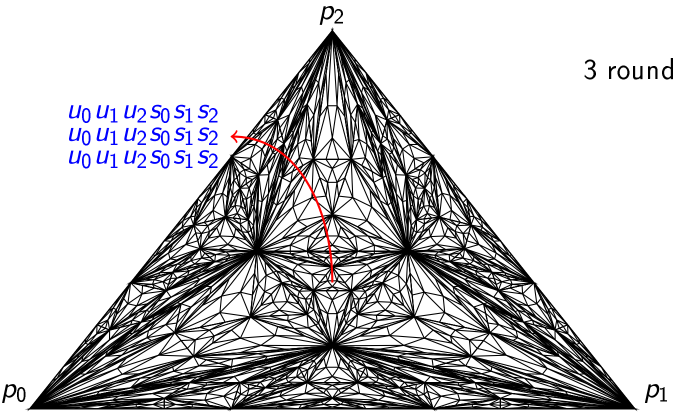

3 rounds

Protocol complex, (dimension 2)

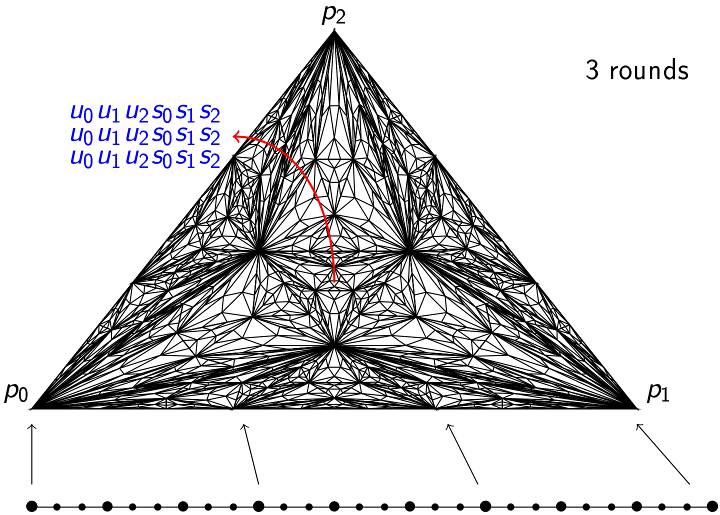

Equivalence of representations

Theorem (Goubault, Mimram, Tasson 2018)

An execution in iterated immediate snapshot protocols is represented equivalently by:

- *an interleaving trace (equivalence class on words $u_i s_j \ldots$)*
- *a dihomotopic dipath*
- *an interval order*
- *a simplex in the protocol complex*

Asynchronous computability theorem

Theorem (Herlihy and Shavit, 1999)

Let T be a distributed task, with I_T and O_T its input and output complexes.
There is a protocol solving task T if and only if there is a color-preserving simplicial map between a subdivision of I_T and O_T.

An application: Impossibility of consensus

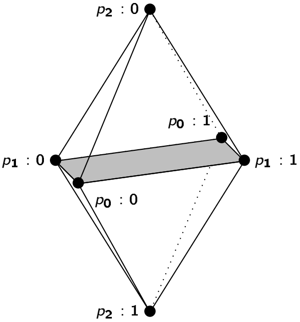

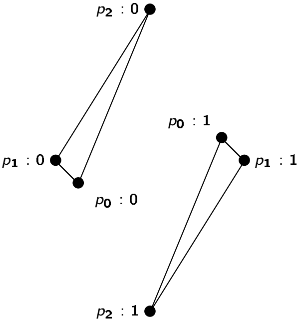

An application: Impossibility of consensus

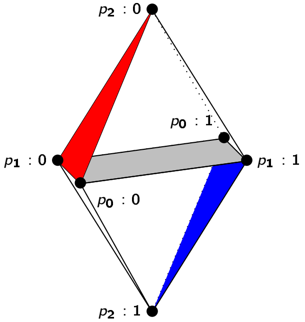

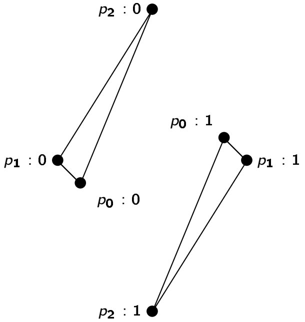

An application: Impossibility of consensus

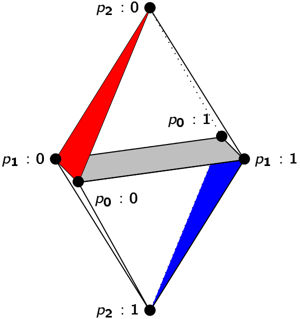

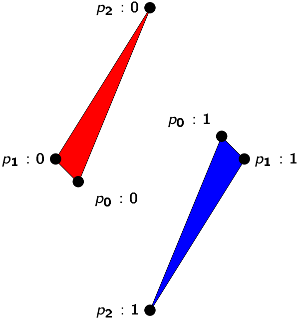

An application: Impossibility of consensus

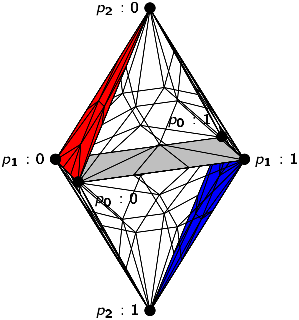

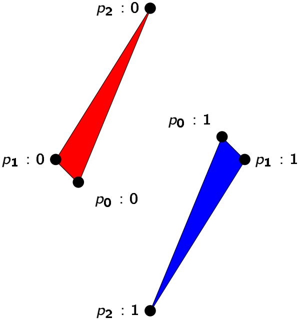

An application: Impossibility of consensus

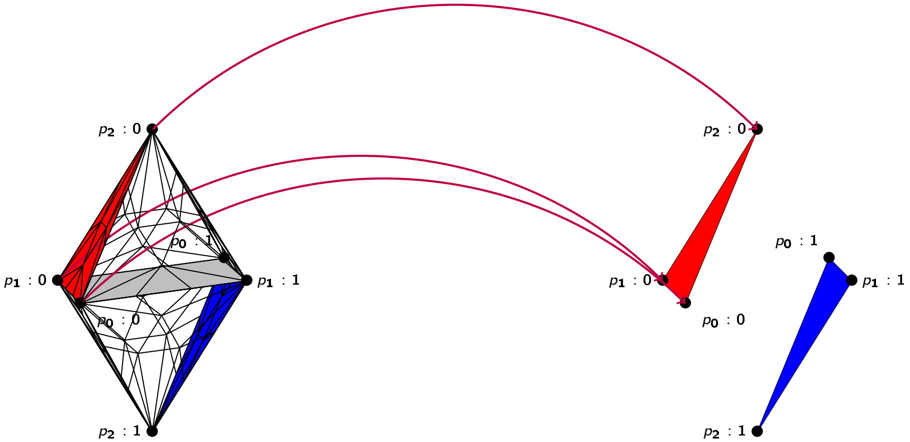

An application: Impossibility of consensus

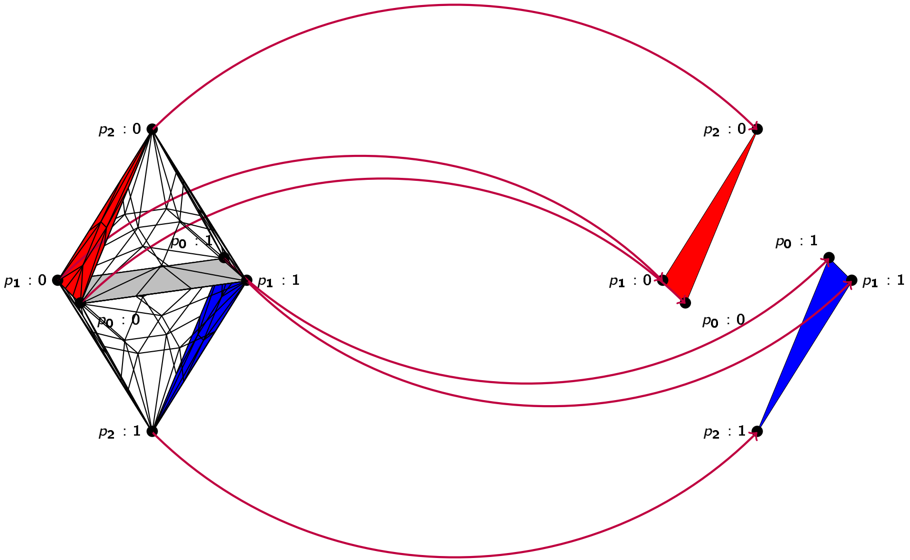

k-set agreement — recall

Definition (k-set agreement)

Specification:

- Each process starts with an initial proper value (n distinct inputs)
- At the end, there is no more than k distinct outputs
- Thue output values must be among the inputs of the participating process

k-set agreement is impossible if $k < n$.

2-set agreement, 3 processes

Input complex and subdivisions

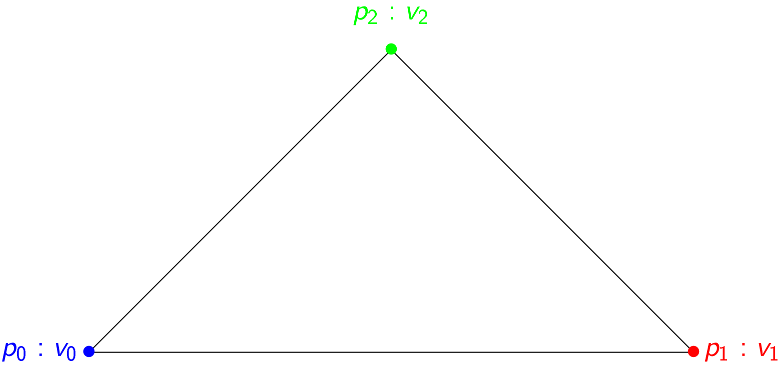

2-set agreement, 3 processes

Input complex and subdivisions

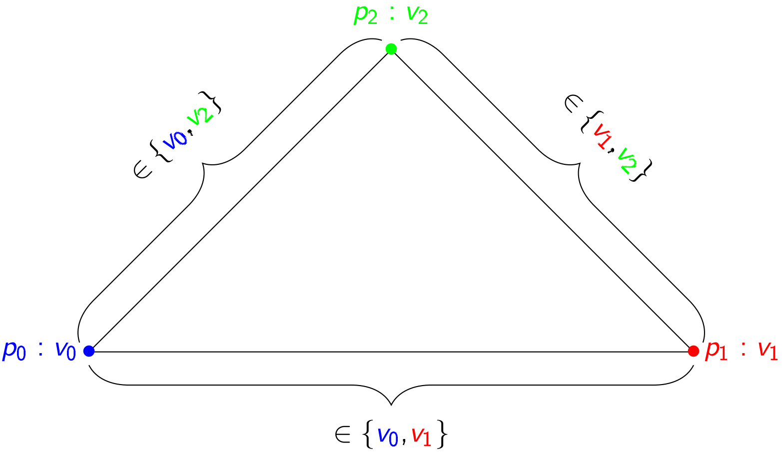

2-set agreement, 3 processes

Input complex and subdivisions

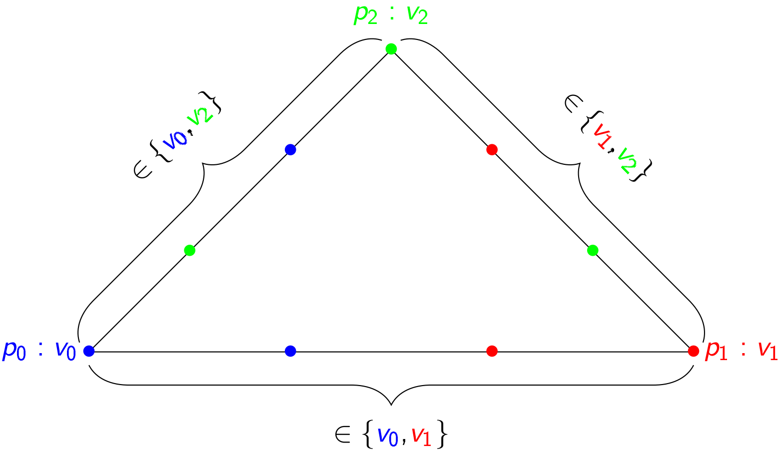

2-set agreement, 3 processes

Input complex and subdivisions

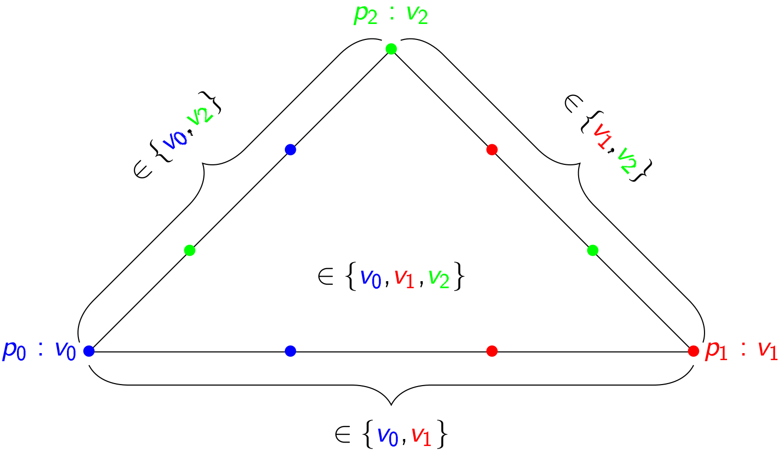

2-set agreement, 3 processes

Input complex and subdivisions

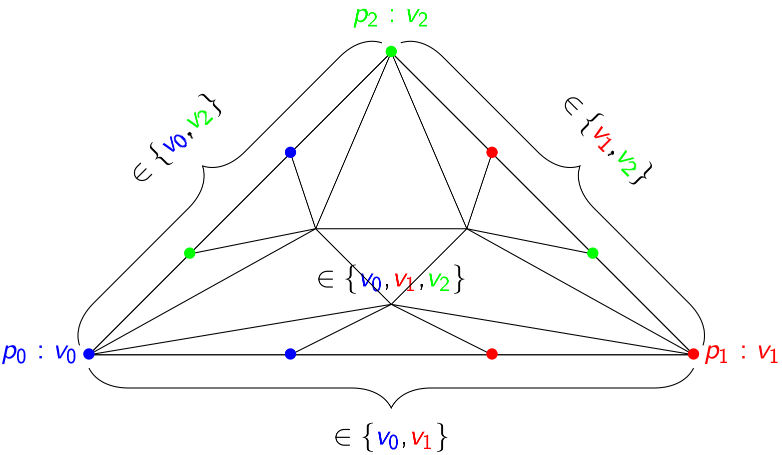

2-set agreement, 3 processes

Input complex and subdivisions

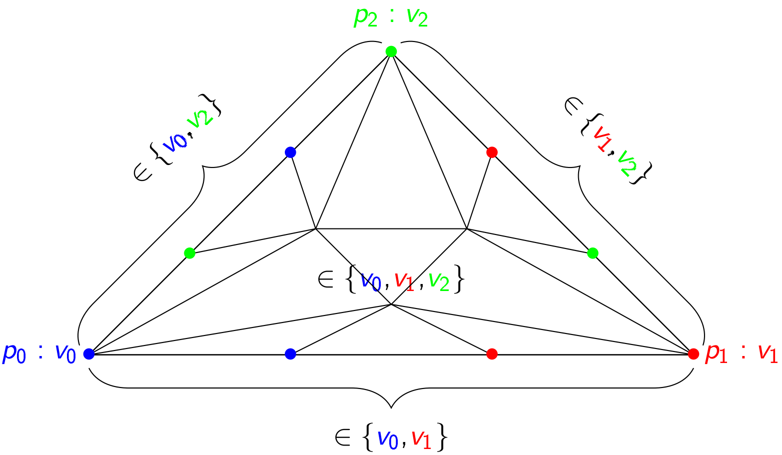

Sperner's Lemma

In any subdivision, there is a simplex thas has 3 colors.

2-set agreement, 3 processes
Output complex

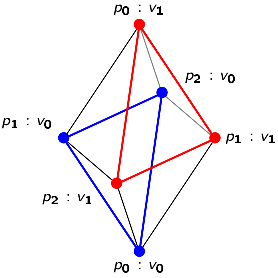

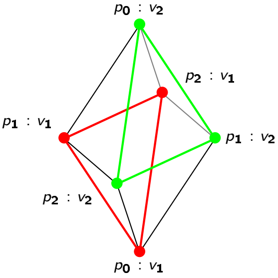

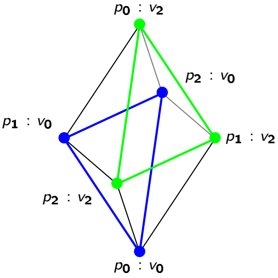

2-set agreement, 3 processes

Output complex

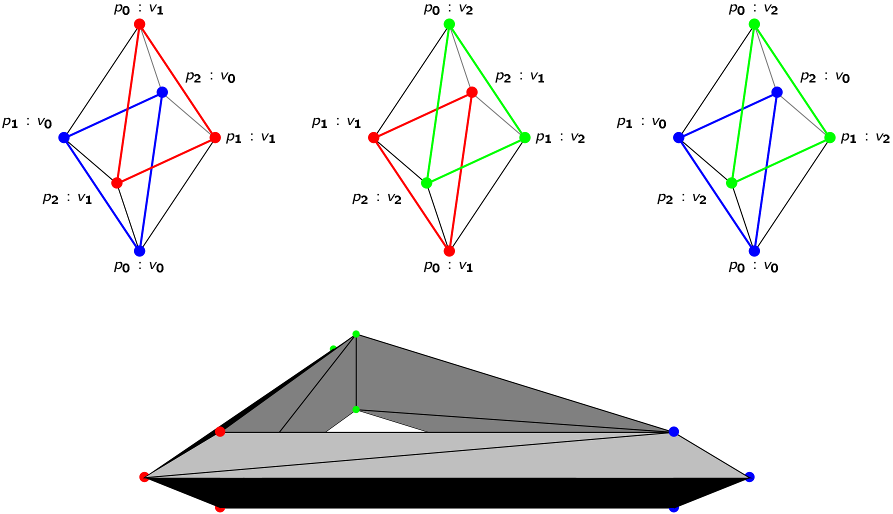

1 Introduction

2 Combinatorial topology

3 Beyond impossibility and determinism

4 The lower bound approach

Tackle the impossibility results

In order to circumvent impossibility, there exist various approaches[1] :

- Time assumption
- Partial synchrony
- Failure detectors
- Randomization

In the last case, we do not talk about termination, but about *probability of success at round r*.

The processes are given an operation `coin()`, that returns a random bit, which is generally assumed not to be known in advance by the adversary.

[1] (see Aspnes, *randomized consensus survey*, 2002)

Randomized protocols

1984 Ben-Or (consensus, t-resilient message-passing)

1990 Aspnes-Herlihy (consensus, wait-free shared memory, shared coin)

1994 Chor-Israeli-Li (multi valued-consensus, wait-free shared memory)

2001 Mostefaoui-Raynal (k-set agreement, message passing)

2010 Censor (k-set agreement, shared memory)

...

Observation

In all these protocols, for any execution, the probability of failure decreases as the number of rounds increases.

The lower bound approach aims to show that this phenomenon is inherent to consensus and agreement.

1 Introduction

2 Combinatorial topology

3 Beyond impossibility and determinism

4 The lower bound approach

Indistinguishability

Definition (Indistinguishability)

Two execution σ and τ are *indistinguishable* if there is at least one process p such that p has the same state after σ and τ.

Definition (Indistinguishability chain)

An indistinguishability chain is a sequence of executions $(\sigma_0, \ldots, \sigma_{n-1})$ s.t σ_i and σ_{i+1} are indistinguishable for all i.

Lower bound for binary consensus

Definition (Probability of failure)

let A be a consensus protocol, and σ an execution of A. $\overline{p}^r_A(\sigma)$ is the probability that A fails on σ at round r.
$\overline{p}^r_A = \max\{\overline{p}^r_A(\sigma) \mid \sigma \text{ is an execution}\}$

Definition

Let C_0 (resp C_1) the initial configuration where every process proposes 0 (resp 1). $f(r)$ is the length of the smallest chain between an execution starting from C_0 and an execution starting from C_1, for r rounds.

Theorem (Attiya-Censor(2010))

For any randomized consensus protocol A, $\overline{p}^k_A \geq \frac{1}{f(r)}$

Computing indistinguishability chains
(binary consensus)

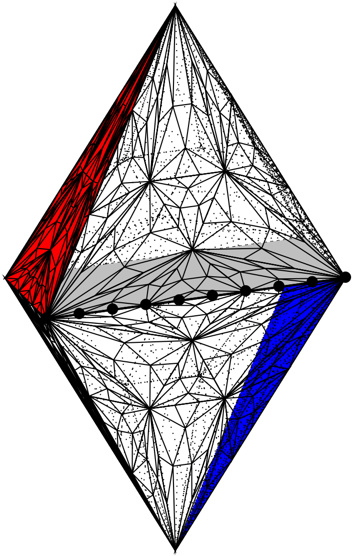

Computing indistinguishability chains

(binary consensus)

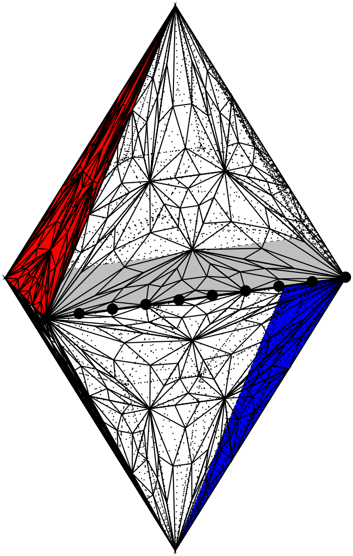

$\longrightarrow f(r) = 3^r$

Lower bound for $n-1$-set agreement

Definition (g(n))

$g(n)$ is the number of maximal simplices in the subdivision of the $n-$simplex[a].

[a] $g(n)$ is the ordered Bell's number of rank n. $g(n) \approx \frac{n!}{2(\log 2)^{n+1}}$

Theorem (Chouquet, Phd thesis, 2019)

For any algorithm A, its probability of failing $n-1$-set agreement at round r $\overline{q}^r_A$ is at least $\frac{1}{g(n)^r}$

Conclusion and perspectives

Résumé:

- Combinatorial topology is a powerful tool for the analysis of communication in snapshot models.
- Randomization can be imported is topological considerations.
- Probability lower bound can be inferred from combinatorial analysis of the protocol complex.

Perspectives:

- Extend this methods to other tasks (renaming, coloring,symmetry breaking. . .)
- Use the lower bound analysis to design agreement algorithms inspired from topology (ongoing work with Pierre Fraigniaud, Ami Paz and Christine Tasson)
- Consider t-resilience, message-passing,. . .

Thank you