Computing in the Presence of Concurrent Solo Executions

Sergio Rajsbaum, Michel Raynal, Julien Stainer

Outline

Remaining Issues and Perspectives

Wait-free Algorithms and Solo Executions

- Wait-free models in both meanings:
  - as a progress condition: each process makes progress in a finite number of steps, whatever the level of concurrence;
  - as a resiliency condition: the computation has to be valid even if all processes but one crash.

- These wait-free criteria and the fact that slow and crashed processes are undistinguishable entail that some processes
  - may have to behave as if they were alone;
  - do not have access to other processes inputs.

Shared Memory, Message passing and Solo Executions

- If processes share a memory, then at most one of them can be in that situation for a given execution.
- If processes exchange asynchronous messages, then all of them may have to behave as if they were alone.

What could be computed in intermediate models in which up to d processes may run solo?

d-Solo Models

- An iterated model generalizing the iterated immediate snapshot model.
- A one-shot communication object for each round.
- The accesses to a round object are set-linearizable but the first set of concurrent accesses can miss each other.
  - If they do, then this set contains at most d processes.

A spectrum of models that spans from message-passing ($d = n$) to shared memory ($d = 1$).

From the Immediate Snapshot Object...

- Each process p provides a value v_p to the object and retrieves a set of values (a view).
- As with the immediate snapshot object, any ordered partition $(\pi_1, \ldots, \pi_x)$ of the set of the processes accessing the object describe a valid behavior for the object:
  - the view of any process belonging to π_i is $\bigcup_{j \leq i} \{(p, v_p), p \in \pi_j\}$.

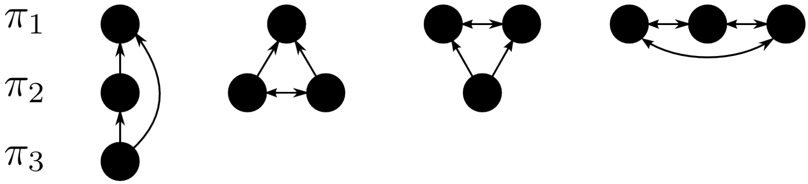

From the Immediate Snapshot Object…

- Each process p provides a value v_p to the object and retrieves a set of values (a view).
- As with the immediate snapshot object, any ordered partition $(\pi_1, \ldots, \pi_x)$ of the set of the processes accessing the object describe a valid behavior for the object:
  - the view of any process belonging to π_i is $\bigcup_{j \leq i} \{(p, v_p), p \in \pi_j\}$.

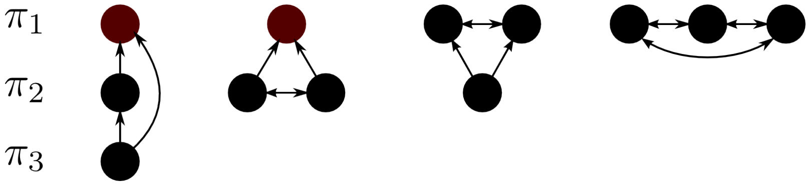

. . . to the CO^d Communication Object

- Additionally, any ordered partition $(\rho_1, \ldots, \rho_x)$ of the set of processes accessing the object describe another authorized behavior for the object if $|\rho_1| \leq d$:
  - if $i > 1$, then the view of any process belonging to π_i is $\bigcup_{j \leq i} \{(p, v_p), p \in \pi_j\}$;
  - the view of a process p of ρ_1 is $\{v_p\}$.

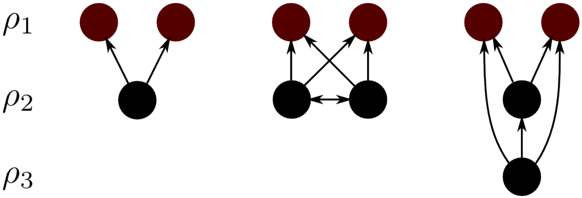

The Colorless Algorithm in the d-solo model

- We consider the case of a colorless algorithm:
  - processes do not use their identities during the computation;
  - they use the object as a set: during each round a process writes the last view it retrieved (initially its input value) ignoring writers identities and multiple occurences of the same view;
  - they compute their output from their view after R rounds.

- It allows us to describe all the possible states of the system after the execution of R rounds by a subdivided complex without coloring vertices with process identities.

The Complex of d-solo executions

- The usual behavior of the immediate snapshot object being still allowed, the usual barycentric subdivision of the (colorless) input complex represent a part of the possible executions.
- At each step of subdivision, we have to consider the additional behaviors where more than one process retrieves only its own value.

- We have to add the simplices built by inserting a barycenter and building the cone over the boundary only in simplices of dimension larger than $d' \leq d$.

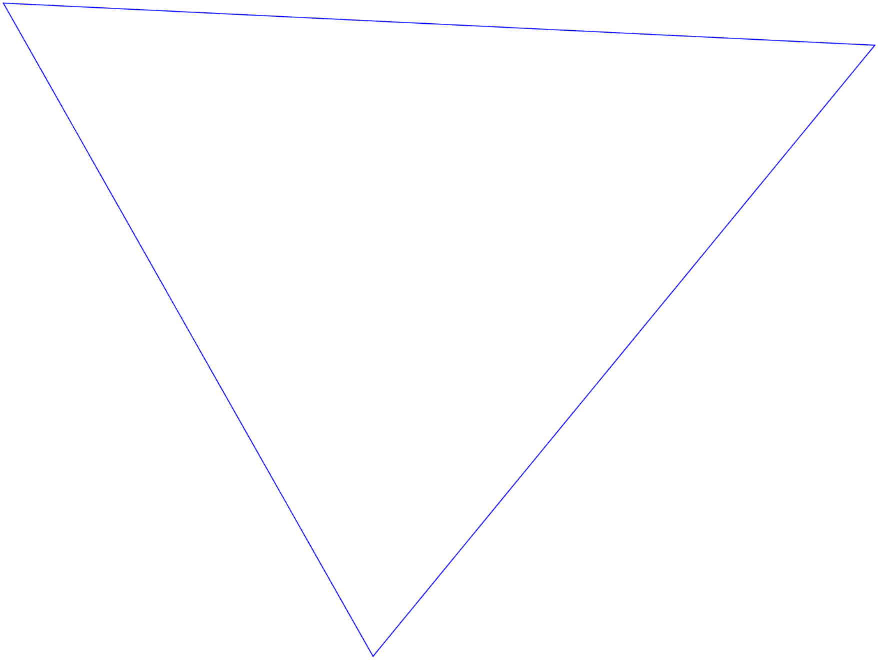

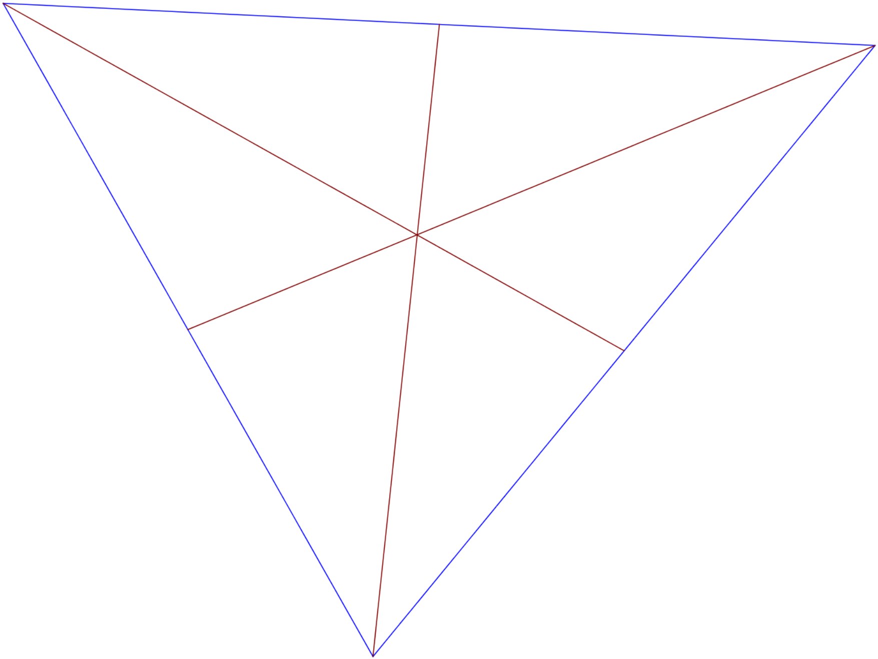

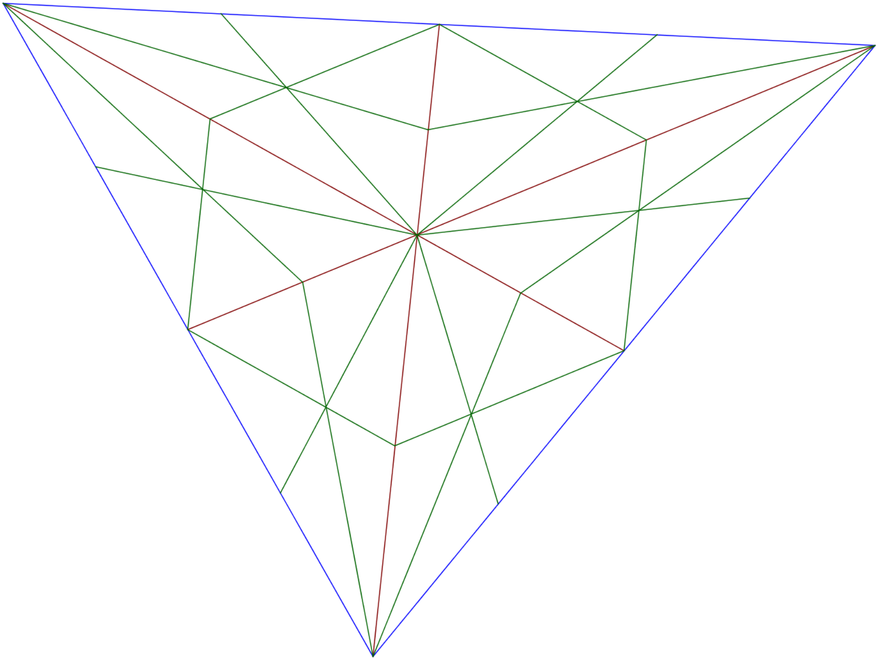

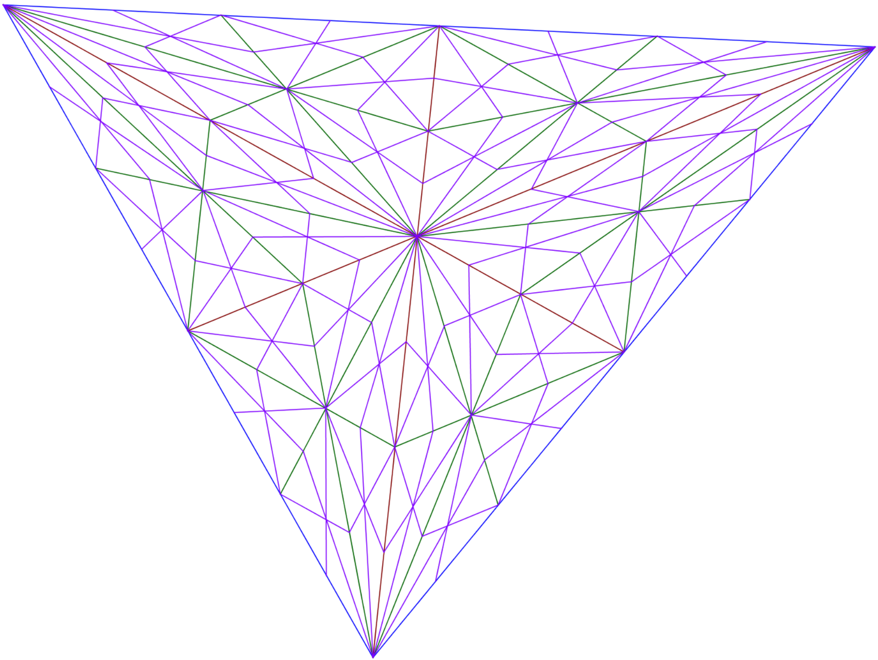

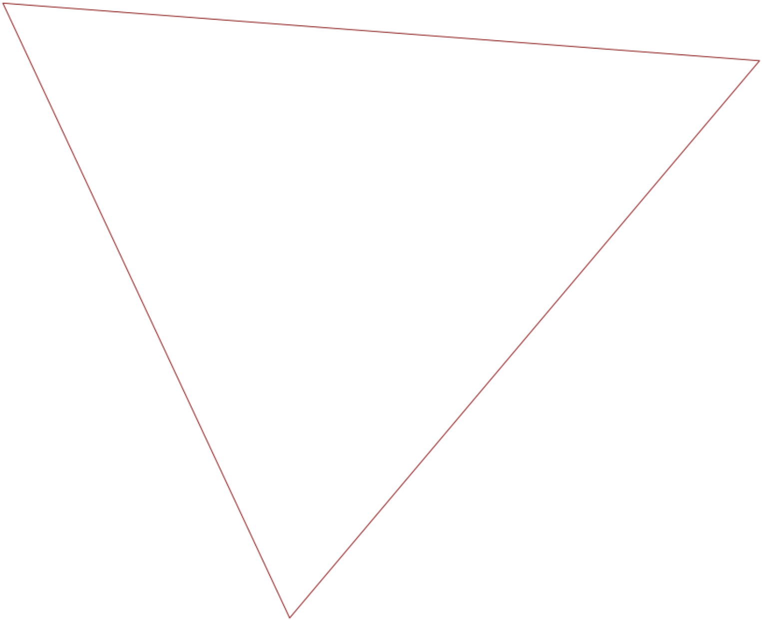

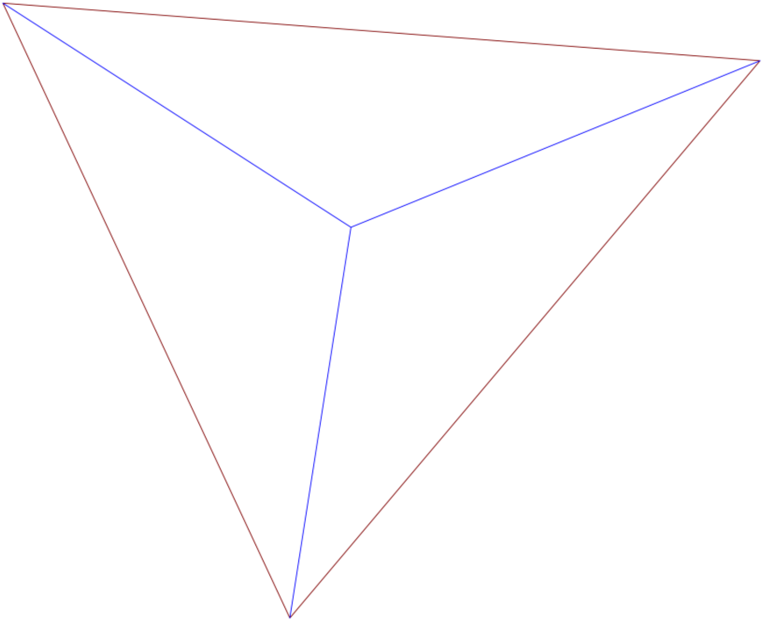

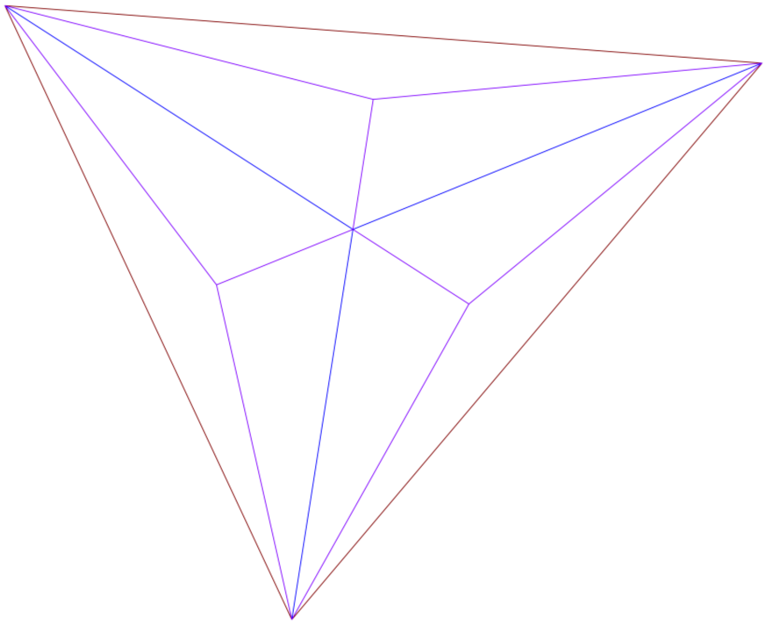

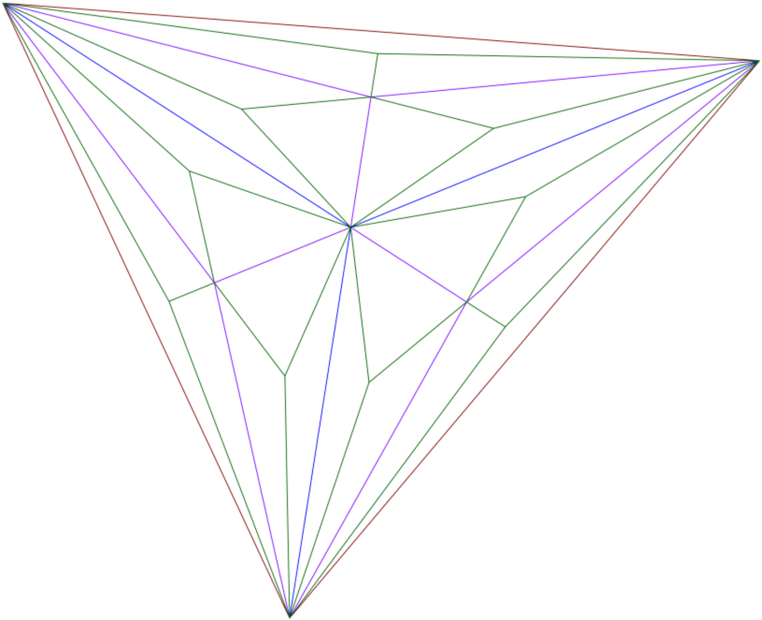

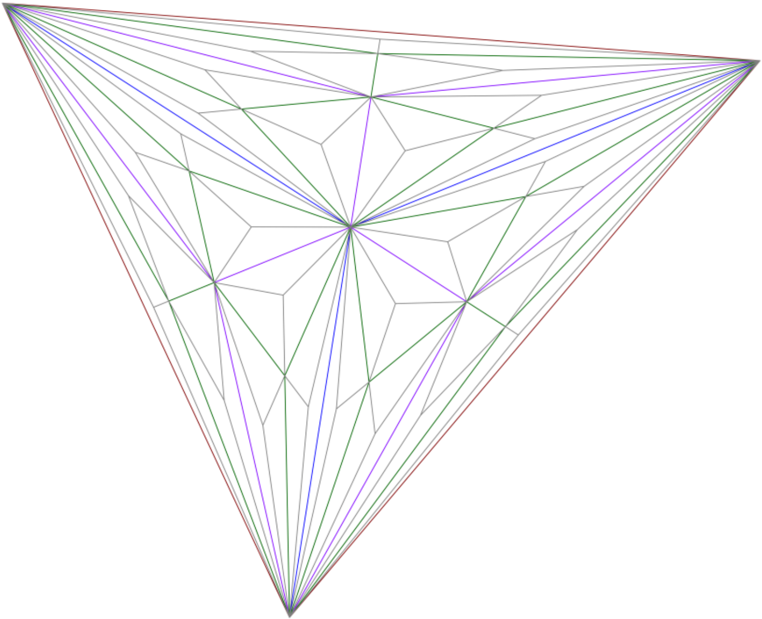

Task Solvability in the d-Solo Model

- A Colorless Task is specified by:
  - the (colorless) complex of all possible input configurations;
  - the (colorless) complex of output configurations;
  - a monotonic carrier map associating each input configuration to a set of allowed output configurations.

Theorem

A colorless task is solvable by a colorless algorithm in the d-solo model with n processes if and only if there is a number of rounds $R \geq 0$ and a simplicial map from the R-iterated d-subdivision of the $n-1$ skeleton of the (colorless) input complex to the (colorless) output complex that is carried by the colorless task carrier map.

The (d, ϵ)-Approximate Agreement Problem

- Generalizing the ϵ-Approximate Agreement that is universal for the Shared Memory Model

- Each process proposes a value from an Euclidian space.
- Termination: all correct processes decide in a finite number of steps.
- Validity: all the decided values belong to the convex hull of the set of proposed values.
- Agreement: there is a set S of up to d processes that can decide any valid value while other processes have to decide within a distance of ϵ from the convex hull of the values decided by processes of S.

Why not a stronger agreement property?

- Proposition 1: there is a set S of up to d processes that can decide any valid value while other processes have to decide within a distance of ϵ from the barycenter of the values decided by processes of S
- Proposition 2: there is a set S of up to d processes that can decide any valid value while other processes have to decide within a distance of ϵ from the barycenter of a (non empty) subset of the values decided by processes of S

Why not a stronger agreement property?

- Proposition 1: there is a set S of up to d processes that can decide any valid value while other processes have to decide within a distance of ϵ from the barycenter of the values decided by processes of S
- Proposition 2: there is a set S of up to d processes that can decide any valid value while other processes have to decide within a distance of ϵ from the barycenter of a (non empty) subset of the values decided by processes of S

- None of these conditions is verified in all runs after a finite number of rounds by the colorless algorithm

(d, ϵ)-Approximate Agreement in the d-Solo Models

- For any ϵ, any d and any n, if the volume of the d-faces of the input complex is bounded, there is a number of round R such that the colorless algorithm solves the (d, ϵ)-approximate agreement problem in the d-solo model.

- For any ϵ, any d and any $n, n > d$, if the input complex contains a large enough regular d-simplex, then the (d, ϵ)-approximate agreement problem is impossible to solve in the $(d + 1)$-solo model.

Since these conditions are compatible, the hierarchy of the d-solo models is strict.

(d, ϵ)-Approximate Agreement is solvable in the d-Solo Model

- Starting from an input complex such that the volume of the d-faces is upper bounded by V
- During each subdivision step, in any d-face, a cone is built with apex at the barycenter
- The volume of the d-faces after the subdivision is then divided by $d+1$
- Taking $R > \frac{\log(V)+\log(d!)-d\log(\epsilon)}{\log(d+1)}$, we have $V \cdot (\frac{1}{d+1})^R < \frac{\epsilon^d}{d!}$
- After R subdivisions, the volume of any d-face is then strictly less than $\frac{\epsilon^d}{d!}$

(d, ϵ)-Approximate Agreement is solvable in the d-Solo Model

- Considering a d-face σ after R subdivisions, let us consider its smallest height h^d_{min}
- It is the distance between a vertex v_d of σ and the $(d-1)$-face of σ with the largest volume V^{d-1}_{max}
- The volume of σ is $\frac{1}{d} \cdot h^d_{min} \cdot V^{d-1}_{max} < \frac{\epsilon^d}{d!}$
- We then have $h^d_{min} < \epsilon$ or $V^{d-1}_{max} < \frac{\epsilon^{d-1}}{(d-1)!}$
- In the first case we are done, in the second case we can iterate

(d, ϵ)-Approximate Agreement is not solvable in the $d+1$-Solo Model

- Suppose that the input complex contains a regular simplex of dimension d whose edge length is strictly larger than $\alpha = 2\epsilon d\sqrt{\frac{2d}{d+1}}$
- Consider a run in which $d+1$ processes start with the vertices of that simplex as inputs
- Suppose that the other processes crash from the beginning and that our $d+1$ run solo forever
- These $d+1$ processes have no choice but outputing their own input values
- The contradiction comes from the impossibility to find a space of dimension $d-1$ distant of less than ϵ from any vertex of our simplex

Relating (d, ϵ)-Approximate Agreement and k-Set Agreement

- A solution to the d-set agreement problem is directly a solution to the (d, ϵ)-approximate agreement.
- It is in general impossible to solve the $(d-1)$-approximate agreement in the d-model enriched with a solution to the d-set agreement.

The weakened memory provided by the d-solo models may give insights on the "weakest memory requirements" needed to solve the k-set agreement.

- We initially thought that a stronger agreement property was possible to fulfill with the colorless algorithm. Considering colored algorithms and/or different termination predicates may be interesting to see what becomes possible.
- The simplicial approximation theorem from which several results are derived in the shared memory model does not apply to the general d-solo model (the diameter of simplices does not tend to zero).
- We would like to have results on decidability, since tasks solvable in message-passing are decidable while those solvable in shared memory are not. Where is the boundary?

- The allowed behaviors for the CO^d object could be changed to authorize partitioning (groups running in isolation) or to evolve during the execution (eventual properties).
- We could investigate further how to enrich the d-solo model with a form of eventual leader allowing d-set agreement to be solved.

Thank you for your attention!