Extending the Wait-free Hierarchy to Multi-Threaded Systems

Matthieu Perrin   Achour Mostefaoui   Grégoire Bonin

LS2N, Université de Nantes

PODC 2020

ACM Symposium on Principles of Distributed Computing

Virtual Conference

August 3-6, 2020

Introduction

Special instructions

- Atomic operations on registers that are added in the hardware
- Examples: test-and-set, fetch-and-add, compare-and-swap, load-link/store conditional...

Universality

O (shared object or register with a special instruction) is universal if:

- Any object with a sequential specification has a linearizable implementation using read/write registers and instances of O

Wait-freedom

- Every operation terminates, no matter crashes or asynchrony

Introduction

Special instructions

- Atomic operations on registers that are added in the hardware
- Examples: test-and-set, fetch-and-add, compare-and-swap, load-link/store conditional...

Universality

O (shared object or register with a special instruction) is universal if:

- Any object with a sequential specification has a linearizable implementation using read/write registers and instances of O

Wait-freedom

- Every operation terminates, no matter crashes or asynchrony

Introduction

Special instructions

- Atomic operations on registers that are added in the hardware
- Examples: test-and-set, fetch-and-add, compare-and-swap, load-link/store conditional...

Universality

O (shared object or register with a special instruction) is universal if:

- Any object with a sequential specification has a linearizable implementation using read/write registers and instances of O

Wait-freedom

- Every operation terminates, no matter crashes or asynchrony

The Wait-free Hierarchy

Consensus number

- Consensus number $k \in \mathbb{N}$:
  - Wait-free universel for k processes
  - Not wait-free universel for $k+1$ processes
- Consensus number ∞:
  - Wait-free universel for k processes, for all k

Significance

- Objects with CN x cannot implement objects with CN $y > x$

The Wait-free Hierarchy in Multi-Threaded Systems

Multi-threaded systems

- Threads can be created dynamically
  - No bound on the number of threads in an execution
- Allocation of unbounded but finite arrays
  - How to allocate one shared register to each thread?

The iterator stack:

- Infinite consensus number
- Not universal in multi-threaded systems

Problem statement

How to compare the synchronization power of shared objects in multi-threaded systems?

The Wait-free Hierarchy in Multi-Threaded Systems

Multi-threaded systems

- Threads can be created dynamically
  - No bound on the number of threads in an execution
- Allocation of unbounded but finite arrays
  - How to allocate one shared register to each thread?

The iterator stack:

- Infinite consensus number
- Not universal in multi-threaded systems

Problem statement

How to compare the synchronization power of shared objects in multi-threaded systems?

The Wait-free Hierarchy in Multi-Threaded Systems

Multi-threaded systems

- Threads can be created dynamically
  - No bound on the number of threads in an execution
- Allocation of unbounded but finite arrays
  - How to allocate one shared register to each thread?

The iterator stack:

- Infinite consensus number
- Not universal in multi-threaded systems

Problem statement

How to compare the synchronization power of shared objects in multi-threaded systems?

Arrival Models

Maximal number of processes in an execution:

M_1^n Classical model
- At most n processes (known to the developper)

M_1 Bounded arrival model
- The bound is known at initialization

M_2 Finite arrival model
- After some time, no new thread is started

M_3 Infinite arrival model
- New threads may keep arriving

Starting an extended hierarchy

	Universal in the arrival model?			
Infinite	✗	✗	✗	✓
Finite	✗	✗	✓	✓
Bounded	✗	✓	✓	✓
Objects	$O_1 < O_2 \cdots$	$< O$	$< O'$	$< O''$

Finite consensus number — Infinite consensus number

Arrival Models

Maximal number of processes in an execution:

M_1^n Classical model

- At most n processes (known to the developper)

M_1 Bounded arrival model

- The bound is known at initialization

M_2 Finite arrival model

- After some time, no new thread is started

M_3 Infinite arrival model

- New threads may keep arriving

Starting an extended hierarchy

	Universal in the arrival model?			
Infinite	✗	✗	✗	✓
Finite	✗	✗	✓	✓
Bounded	✗	✓	✓	✓
Objects	$O_1 < O_2 \cdots$	$< O$	$< O'$	$< O''$

Finite consensus number (first column) — Infinite consensus number (other three columns)

Extended Wait-free Hierarchy

	Arrival Models			Universal without infinite allocation?			
			Infinite	✗	✗	✗	✓
			Finite	✗	✗	✓	✓
	Infinite	Finite	Bounded	✗	✓	✓	✓
Universal with infinite allocation?	✓	✓	✓				
	✗	✓	✓				
	✗	✗	✓				
	✗	✗	✗				

Extended Wait-free Hierarchy

				Universal without infinite allocation?			
	Arrival Models		Infinite	✗	✗	✗	✓
			Finite	✗	✗	✓	✓
	Infinite	Finite	Bounded	✗	✓	✓	✓
Universal with infinite allocation?	✓	✓	✓				
	✗	✓	✓				
	✗	✗	✓				
	✗	✗	✗				

Extended Wait-free Hierarchy

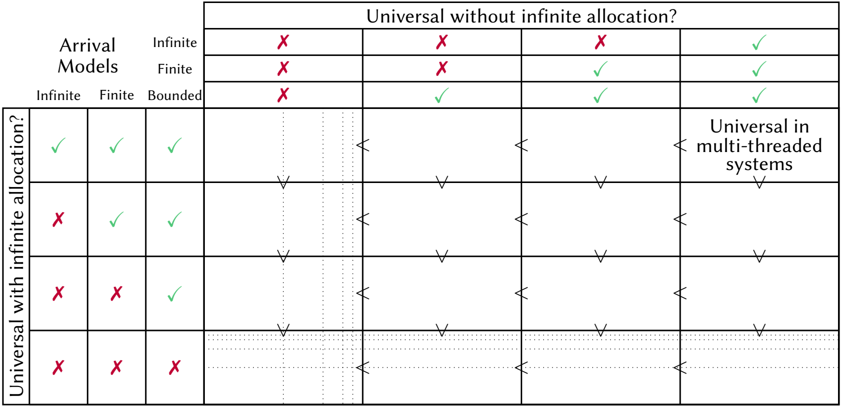

Extended Wait-free Hierarchy

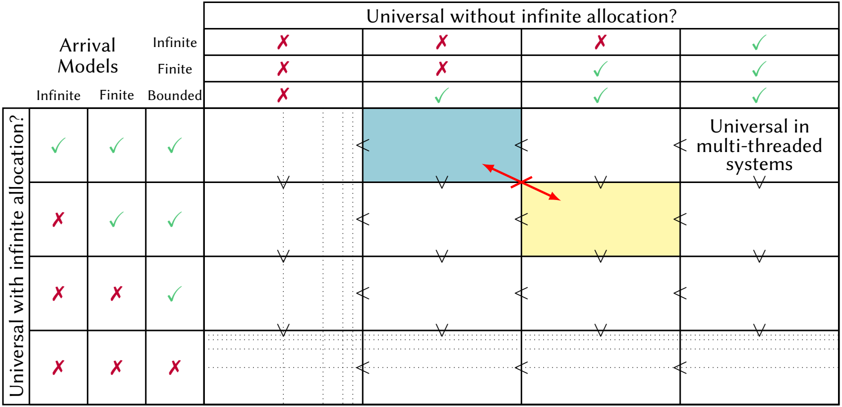

Filling the Hierarchy: State of the Art

				Universal without infinite allocation?			
	Arrival Models		Infinite	✗	✗	✗	✓
			Finite	✗	✗	✓	✓
	Infinite	Finite	Bounded	✗	✓	✓	✓
Universal with infinite allocation?	✓	✓	✓	?	?	?	?
	✗	✓	✓	?	?	?	?
	✗	✗	✓	?	?	?	?
	✗	✗	✗	?	?	?	?

Filling the Hierarchy: State of the Art

Arrival Models				Universal without infinite allocation?			
			Infinite	✗	✗	✗	✓
			Finite	✗	✗	✓	✓
	Infinite	Finite	Bounded	✗	✓	✓	✓
Universal with infinite allocation?	✓	✓	✓	?	?	?	?
	✗	✓	✓	?	?	?	empty
	✗	✗	✓	?	?	empty	empty
	✗	✗	✗	?	empty	empty	empty

empty
(if universal without infinite allocation, still universal with infinite allocation)

Filling the Hierarchy: State of the Art

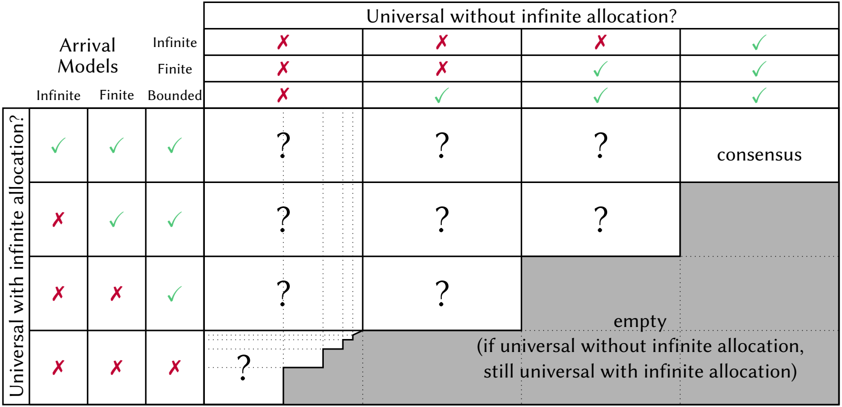

Filling the Hierarchy: State of the Art

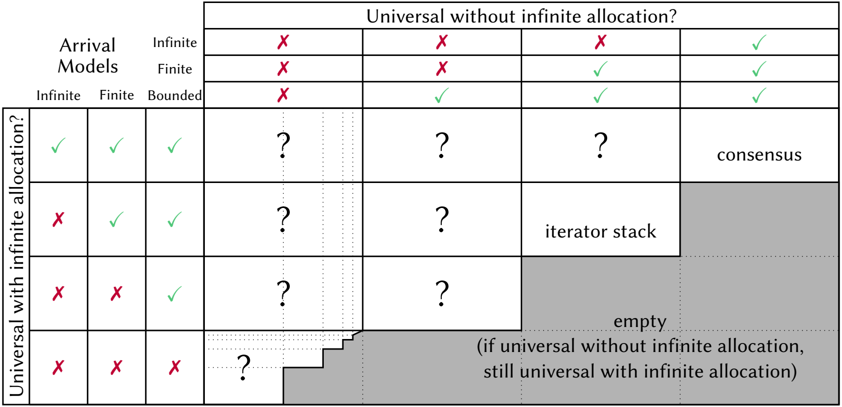

Filling the Hierarchy: State of the Art

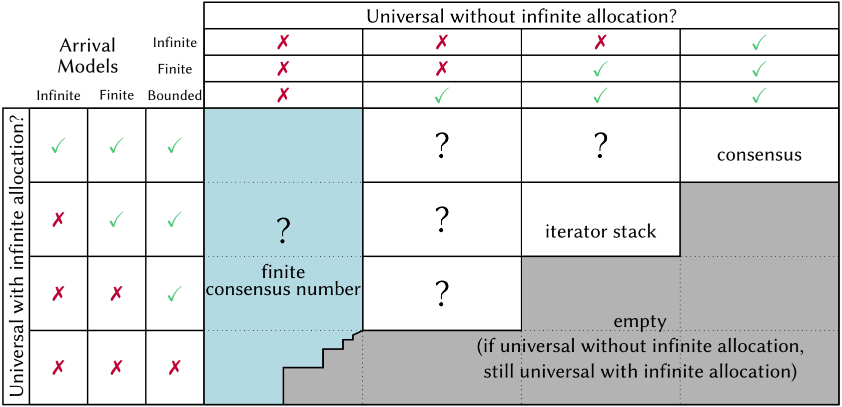

Filling the Hierarchy: State of the Art

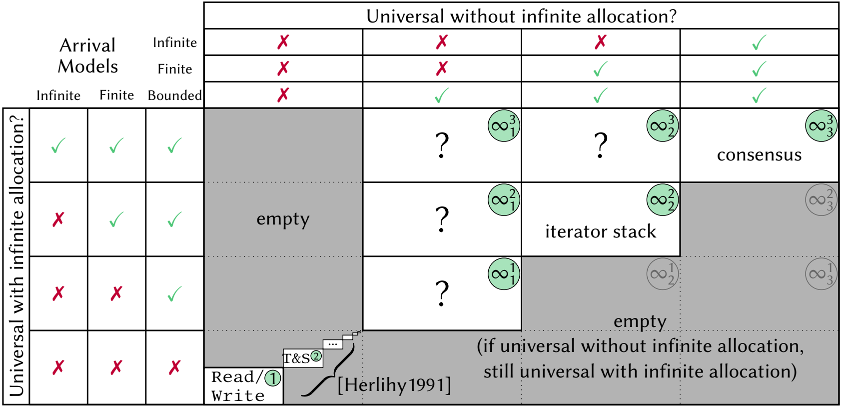

Objects with a Finite Consensus Number

Proposition

If an object O is universal in the classical model M_1^n enriched with infinite arrays, then O is universal in the classical model M_1^n.

Remarks

- The same proof works for the bounded arrival model
- This does not prove that the two models were equivalent

Sketch of the proof

- An algorithm A uses O to solve binary consensus
  - A has a bounded number of inputs (thanks to renaming)
  - For each input, A accesses a bounded number of objects
  - Therefore, A does not need infinite arrays
- Binary consensus is universal in the classical model

Filling the Hierarchy: Consensus Number ∞_1^1

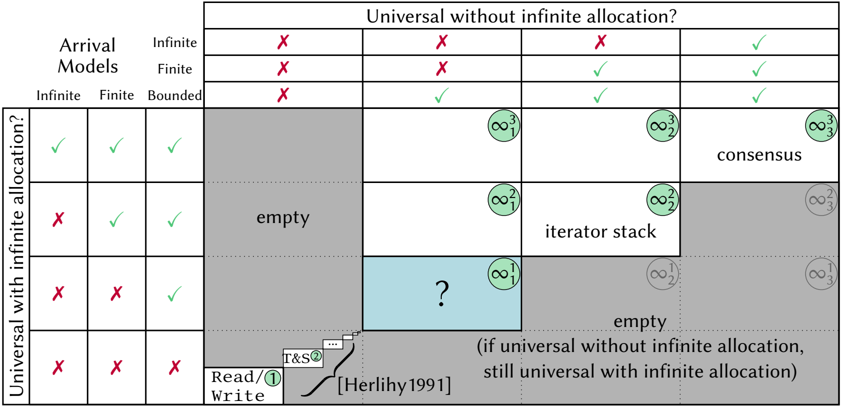

Filling the Hierarchy: Consensus Number ∞_1^1

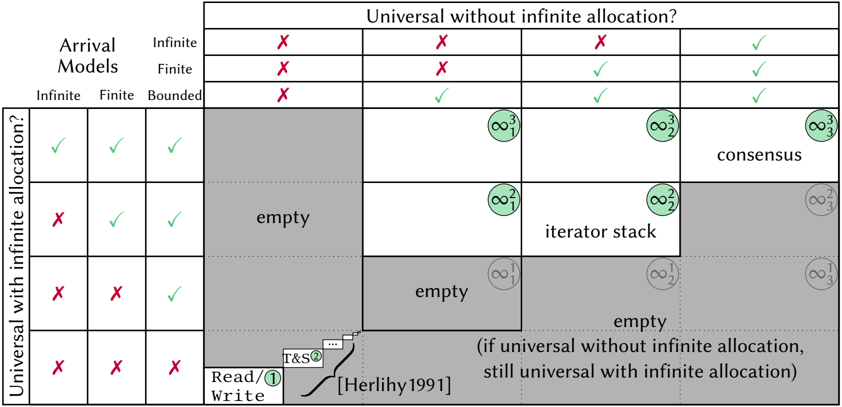

No Object has Consensus Number ∞_1^1

Proposition

If an object O is universal in the bounded arrival model M_1, then O is universal in the finite arrival model M_2 enriched with infinite arrays.

Sketch of the proof

We propose a round-based algorithm. At round k:

- Processes $p_{k+1}, \ldots$
  - Mark `greaterId`$[k]$
  - If `adopt`$[k]$ was written, adopt the value
  - Start a new round
- Processes $p_1, \ldots, p_k$
  - Solve consensus together using O in model M_1^k
  - Write the decided value in `adopt`$[k]$
  - If `greaterId`$[k]$ is marked, start a new round
  - Otherwise, decide

Filling the Hierarchy: Consensus Number ∞_1^2

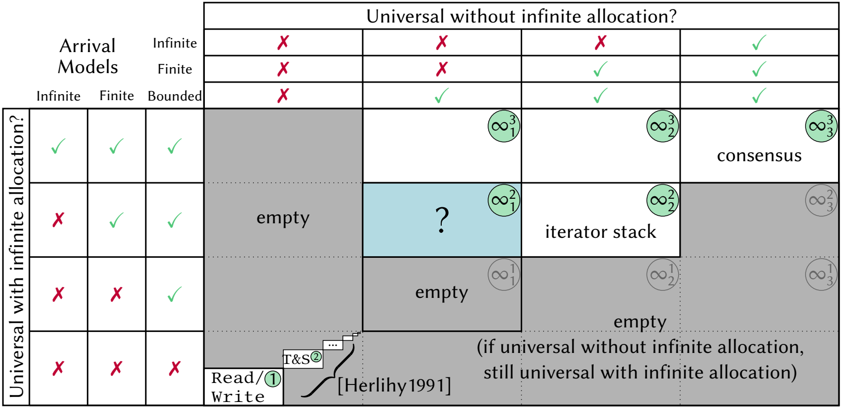

Window register

- Size k chosen at initialization
- A read returns the k last values written

Filling the Hierarchy: Consensus Number ∞_1^2

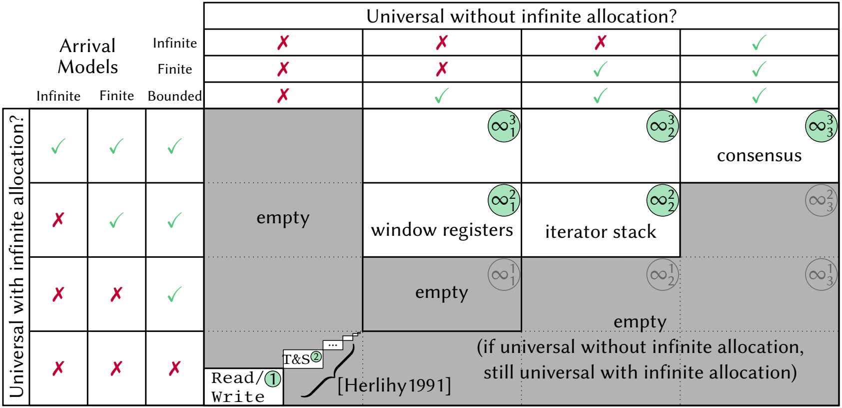

Window register

- Size k chosen at initialization
- A read returns the k last values written

Filling the Hierarchy: Consensus Number ∞_1^2

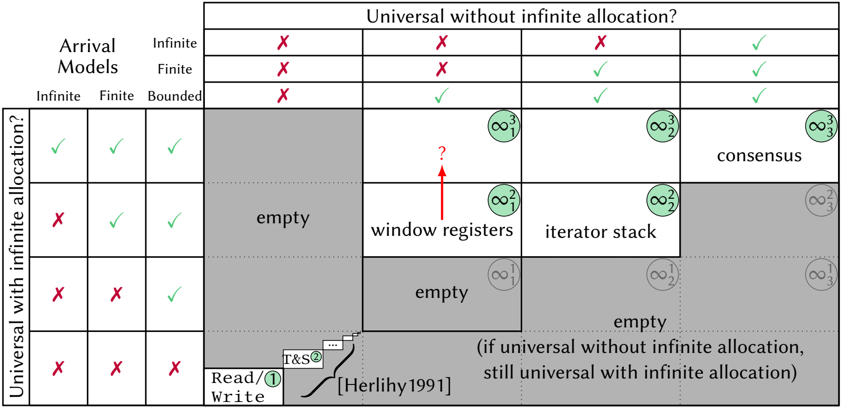

Window register

- Size k chosen at initialization
- A read returns the k last values written

Filling the Hierarchy: Consensus Number ∞_1^2

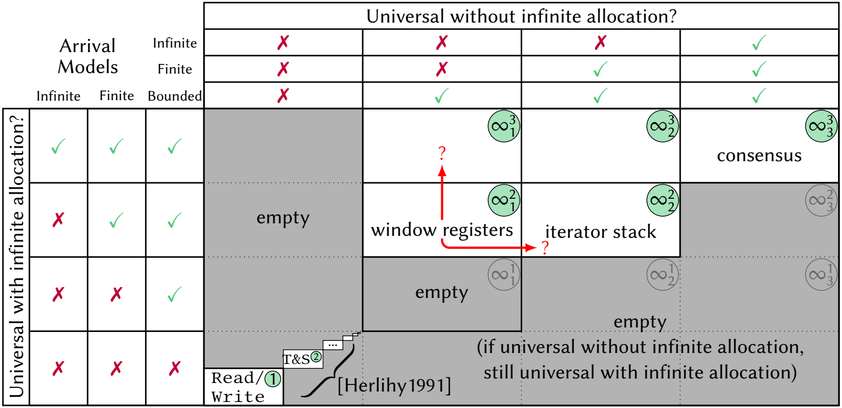

Window register

- Size k chosen at initialization
- A read returns the k last values written

Window Registers are not Universal in M_3

Proposition

It is impossible to solve consensus in the infinite arrival model enriched with infinite arrays, using only window registers

Sketch of the proof

n-critical configuration: If only $p_1, \ldots, p_n$ take steps, consensus is decided by the next step of a process

- There are n-critical configurations for all n
- In an n-critical configurations, $p_1, \ldots, p_n$ are about to write in a window register of size at least k
- We build an infinite execution in which some process passes through n-critical configurations for all n

Window Registers need Infinite Arrays in M_2

Proposition

It is impossible to solve consensus in the finite arrival model, using only window registers and binary consensus objects

Sketch of the proof

- Suppose m binary consensus objects and window registers of size at most l are allocated at initialization
- A large but finite number of processes ($m! \times (2l)^m$) arrive
- We build an execution where two partitions can never communicate
  - Values written on window registers are overwritten
  - All values proposed to binary consensus are the same
- Finally, two different values are decided

The Complete Extended Wait-free Hierarchy

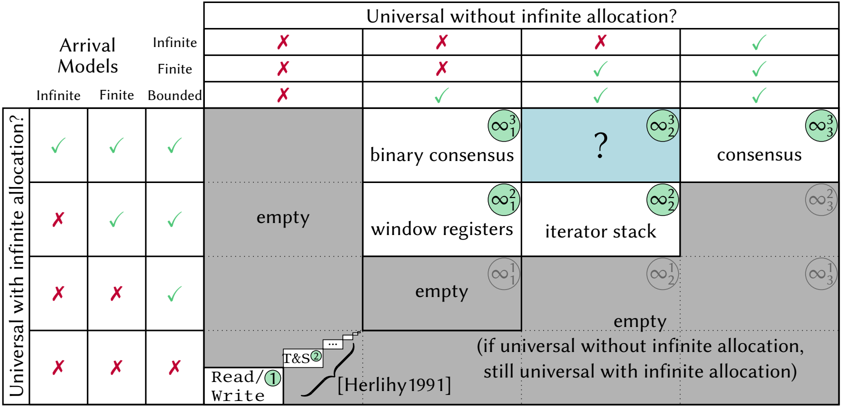

The Complete Extended Wait-free Hierarchy

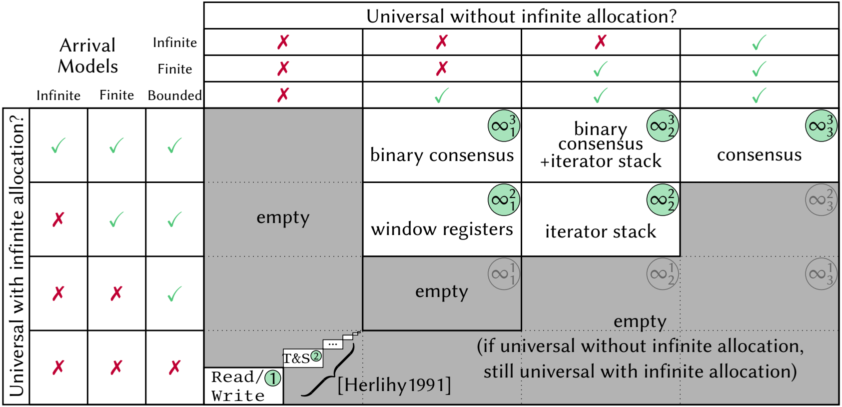

Open Problems

Relevance of consensus number ∞_1^3

- k binary consensus objects synchronize 2^k processes
- Are there objects of consensus number ∞_1^3 that do not allow a polylogarithmic implementation of consensus?

Identification of simpler objects

- Window registers and iterator stacks are complex objects
  - Are there equivalent special instructions?
- We only know a composition for consensus number ∞_2^3
  - Are there not-composed objects of consensus number ∞_2^3?
  - Are there objects of consensus number ∞_1^3 and ∞_2^2 whose composition has consensus number ∞_3^3?

Open Problems

Relevance of consensus number ∞_1^3

- k binary consensus objects synchronize 2^k processes
- Are there objects of consensus number ∞_1^3 that do not allow a polylogarithmic implementation of consensus?

Identification of simpler objects

- Window registers and iterator stacks are complex objects
  - Are there equivalent special instructions?
- We only know a composition for consensus number ∞_2^3
  - Are there not-composed objects of consensus number ∞_2^3?
  - Are there objects of consensus number ∞_1^3 and ∞_2^2 whose composition has consensus number ∞_3^3?