

## A Verified Parallel Scheduler for OCaml 5



Clément  
Allain



Gabriel  
Scherer

## Task Parallelism

Parabs: A Verified Realistic Parallel Scheduler

Chase-Lev Work-Stealing Deque

Future work

## Parallelism in OCaml 5

- ▶ Domains

- ▶ `Domain.spawn` : `(unit -> 'a) -> 'a Domain.t`
- ▶ `Domain.join` : `'a Domain.t -> 'a`

- ▶ Atomic primitives

- ▶ Atomic references (`'a Atomic.t`)
- ▶ Atomic record fields (`'a Atomic.Loc.t`)

## Fibonacci Using Domains

```
let rec fibonacci n =  
  if n <= 1 then  
    1  
  else  
    let dom1 = Domain.spawn (fun () -> fibonacci (n - 1)) in  
    let dom2 = Domain.spawn (fun () -> fibonacci (n - 2)) in  
    Domain.join dom1 + Domain.join dom2
```

## Fibonacci Using Domains

```
let rec fibonacci n =  
  if n  
  1  
  else  
  let  
  let  
  Dom
```

**This implementation is *wrong*:**

- ▶ the number of active domains is limited;
- ▶ performance collapses when spawning more domains than available cores.

## Fibonacci Using a Parallel Scheduler

```
let rec fibonacci ctx n =  
  if n <= 1 then  
    1  
  else  
    let fut1 = Future.async ctx (fun ctx -> fibonacci ctx (n - 1)) in  
    let fut2 = Future.async ctx (fun ctx -> fibonacci ctx (n - 2)) in  
    Future.wait ctx fut1 + Future.wait ctx fut2  
  
let fibonacci num_domain n =  
  let pool = Pool.create num_domain in  
  let res = Pool.run_on pool (fun ctx -> fibonacci ctx n) in  
  Pool.close pool ;  
  res
```

## Existing Parallel Schedulers



Domainslib  
Moonpool



Rayon  
Tokio




Cilk  
TBB  
Taskflow



Goroutines

# Work-Stealing



Task Parallelism

Parabs: A Verified Realistic Parallel Scheduler

Chase-Lev Work-Stealing Deque

Future work

# Overview

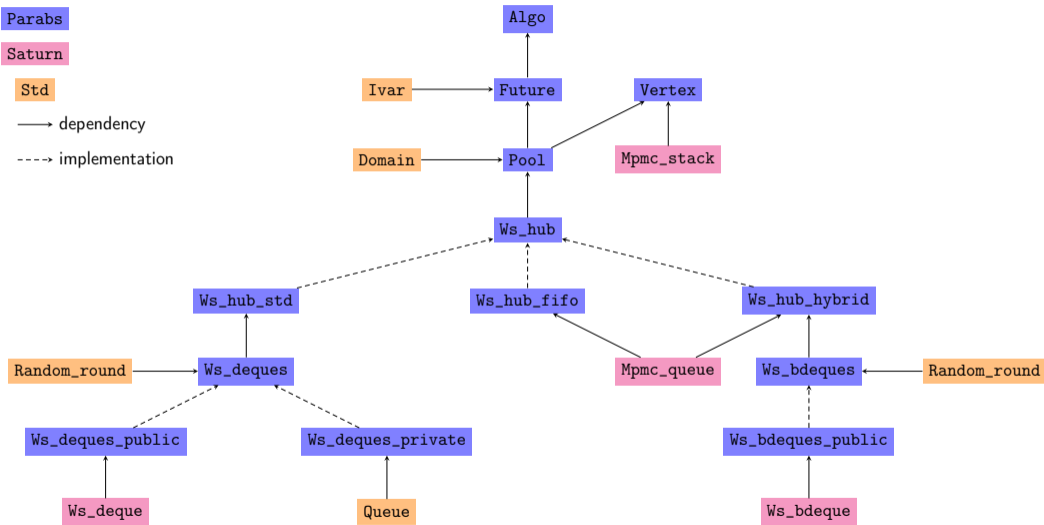
Parabs

Saturn

Std

— dependency

- - - - implementation







# Overview

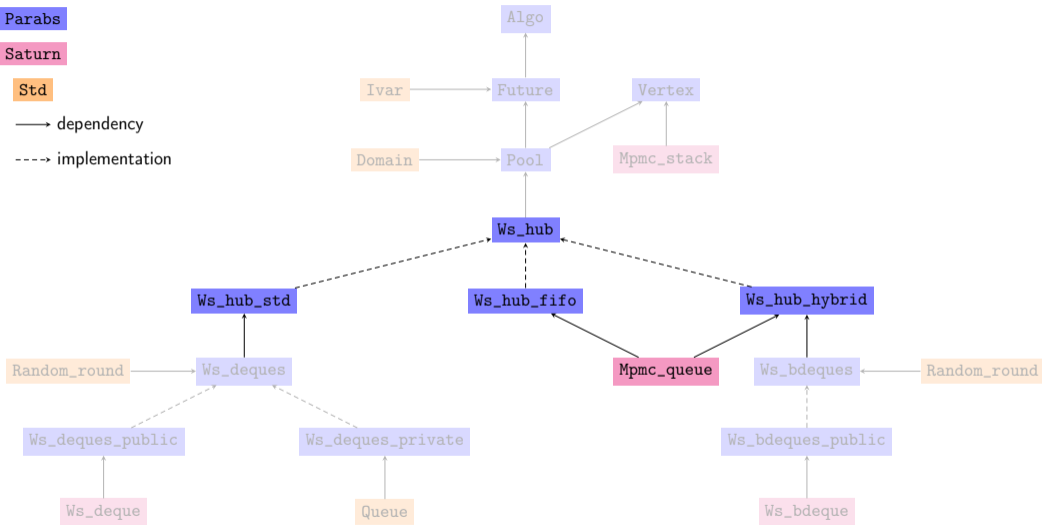
Parabs

Saturn

Std

— dependency

- - - - implementation



# Overview

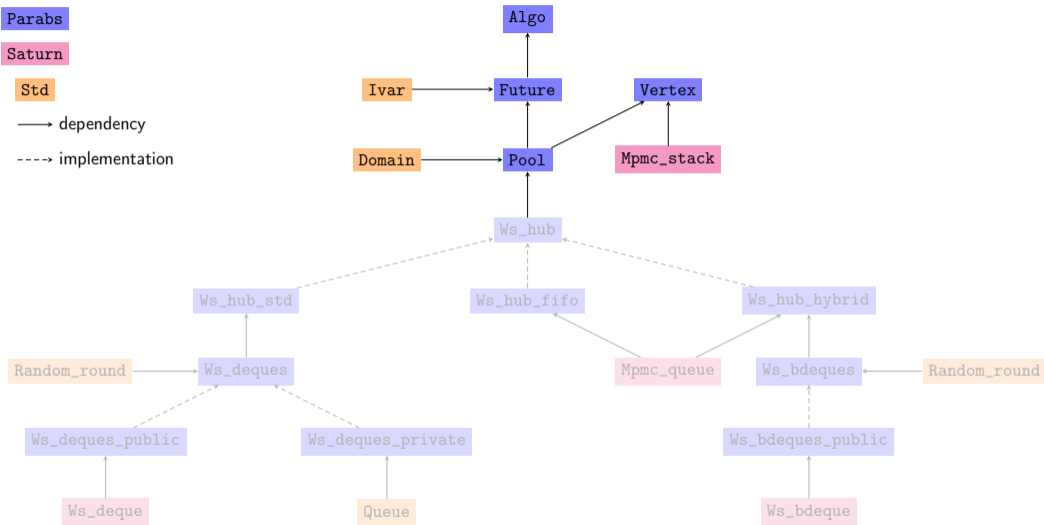
Parabs

Saturn

Std

— dependency

- - - - implementation





**Parabs has been fully verified using the Zoo framework**  
(assuming sequential consistency).

# Futures

ASYNC-SPEC

$$\frac{
 \left\{ \begin{array}{l}
 \text{pool.context } \text{pool } \text{ctx } \text{scope } * \\
 \forall \text{ ctx } \text{scope.} \\
 \text{pool.context } \text{pool } \text{ctx } \text{scope } \rightarrow * \\
 \text{wp task ctx } \left\{ \begin{array}{l}
 v. \text{pool.context } \text{pool } \text{ctx } \text{scope } * \\
 \triangleright \Psi v * \\
 \triangleright \square \Xi v
 \end{array} \right\}
 \end{array} \right\}
 }{
 \text{async ctx task}
 }$$


---


$$\left\{ \begin{array}{l}
 t. \text{pool.context } \text{pool } \text{ctx } \text{scope } * \\
 \text{inv pool } t \Psi \Xi * \\
 \text{consumer } t \Psi
 \end{array} \right\}$$

CONSUMER-DIVIDE

$$\frac{
 \begin{array}{l}
 \text{inv pool } t \Psi \Xi \\
 \text{consumer } t \left( \lambda v. \bigstar_{X \in X_s} X v \right)
 \end{array}
 }{
 \Rightarrow \bigstar_{X \in X_s} \text{consumer } t X
 }$$

# Futures

WAIT-SPEC

$$\frac{\left\{ \begin{array}{l} \text{pool.context } pool \text{ ctx } scope * \\ \text{inv } pool \ t \ \Psi \ \Xi \end{array} \right\}}{\text{wait } ctx \ t} \frac{}{\left\{ \begin{array}{l} v. \mathbb{L} 2 * \\ \text{pool.context } pool \text{ ctx } scope * \\ \text{result } t \ v \end{array} \right\}}$$


INV-RESULT-CONSUMER

$$\frac{\begin{array}{l} \text{inv } pool \ t \ \Psi \ \Xi \\ \text{result } t \ v \\ \text{consumer } t \ X \end{array}}{\Rightarrow \triangleright^2 X \ v}$$


INV-RESULT

$$\frac{\begin{array}{l} \text{inv } pool \ t \ \Psi \ \Xi \\ \text{result } t \ v \end{array}}{\Rightarrow \triangleright \square \Xi \ v}$$


# Benchmarks




(a) fibonacci




(b) for\_irregular



(c) iota



(d) lu



(e) matmul

Parabs has equal or better performance than Domainslib.


Task Parallelism

Parabs: A Verified Realistic Parallel Scheduler

Chase-Lev Work-Stealing Deque

Future work

# Work-Stealing Deque



## Dynamic Circular Work-Stealing Deque

David Chase  
Sun Microsystems Laboratories  
Mailstop UBUR02-311  
1 Network Drive  
Burlington, MA 01803, USA  
dr2chase@sun.com

Yossi Lev  
Brown University & Sun Microsystems Laboratories  
Mailstop UBUR02-311  
1 Network Drive  
Burlington, MA 01803, USA  
yosef.lev@sun.com

### ABSTRACT

The non-blocking work-stealing algorithm of Arora, Blumofe, and Plaxton (henceforth *ABP work-stealing*) is on its way to becoming the multiprocessor load balancing technology of choice in both industry and academia. This highly efficient scheme is based on a collection of array-based double-ended queues (dequeues) with low cost synchronization among local and stealing processes. Unfortunately, the algorithm's synchronization protocol is strongly based on the use of fixed size arrays, which are prone to overflows, especially in the multiprogrammed environments for which they are designed. We present a work-stealing deque that does not have the overflow problem.

The only ABP-style work-stealing algorithm that eliminates the overflow problem is the list-based one presented by Hender, Lev and Shavit. Their algorithm indeed deals with the overflow problem, but it is complicated, and introduces a trade-off between the space and time complexity, due to the extra work required to maintain the list.

Our new algorithm presents a simple lock-free work-stealing deque, which stores the elements in a cyclic array that can grow when it overflows. The algorithm has no limit other than integer overflow (and the system's memory size) on the number of elements that may be on the deque, and the total

### 1. INTRODUCTION

The ABP work-stealing algorithm of Arora, Blumofe, and Plaxton [2] has been gaining popularity as the multiprocessor load-balancing technology of choice in both industry and academia [2, 1, 4, 9]. The scheme implements a provably efficient work-stealing paradigm due to Blumofe and Leiserson [3] that allows each process to maintain a local work deque,<sup>1</sup> and steal an item from others if its deque becomes empty. The deque's owner process pushes and pops local work to and from the deque's bottom end. To minimize synchronization overhead for the deque's owner, stolen elements are taken from the top end of the deque. No elements are added to the top end of the deque. An ABP deque thus presents three methods in its interface:

- `pushBottom(Object o)`: Pushes *o* onto the bottom of the deque.
- `Object popBottom()`: Pops an object from the bottom of the deque if the deque is not empty, otherwise returns *Empty*.
- `Object steal()`: If the deque is empty, returns *Empty*. Otherwise, re-

## Verification of Chase-Lev work-stealing deque (work in progress)

Clément Allain  
François Pottier

Inria Paris

October 15, 2025

## Conclusion

- ▶ Coq mechanization is available on `github` :  
`https://github.com/clef-men/caml5`
- ▶ Simplified Chase-Lev deque (one infinite array) ✓  
Real-life Chase-Lev deque (multiple circular arrays) ⚠
- ▶ Proof looks more complex than the sketch. In particular, transitions between logical states are not really formalized.
- ▶ We plan to verify more primitives (Domainslib, Taskflow) based on Chase-Lev deque. This is thanks to modularity of IRIS specifications.

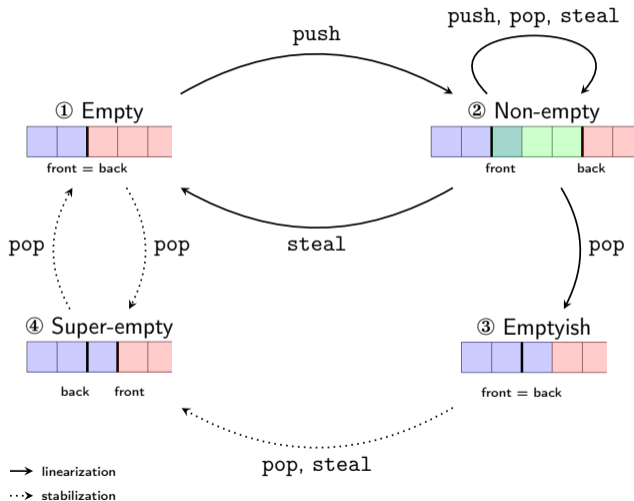
## My First Interaction With Robbert Krebbers



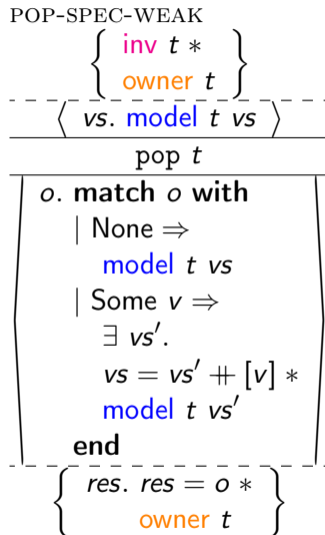

*Already verified in iris/examples  
without prophecy variables!*

Robbert Krebbers, Iris Workshop 2023

# Future-Dependent Linearization Point



# Weak Specification



# Weak Specification

STEAL-SPEC-WEAK

{ inv  $t$  }

POP-SPEC-WEAK

{ inv  $t^*$   
owner  $t$  }

**This specification is *too weak* to prove completion :**  
Once the scheduler terminates, all submitted tasks have finished.

OBLIGATION-FINISHED

obligation pool  $P$

finished pool

$\triangleright \square P$

CONSUMER-FINISHED

consumer pool  $P$

finished pool

$\Rightarrow P$

model  $\tau$  vs

end

{ res. res = 0 }

model  $\tau$  vs

end

{ res. res = 0 \*  
owner  $t$  }

# Strong Specification

STEAL-SPEC

$$\frac{\frac{\{ \text{inv } t \}}{\langle \text{vs. model } t \text{ vs} \rangle}}{\text{steal } t} \quad \frac{\langle \text{model } t \text{ (tail vs)} \rangle}{\{ \text{res. res} = \text{head vs} \}}$$

POP-SPEC

$$\frac{\frac{\{ \text{inv } t * \text{owner } t \text{ ws} \}}{\langle \text{vs. model } t \text{ vs} \rangle}}{\text{pop } t} \quad \left[ \begin{array}{l} o \text{ ws}'. \text{ match } o \text{ with} \\ | \text{ None} \Rightarrow \\ \quad \text{vs} = [] * \text{ws}' = [] * \\ \quad \text{model } t [] \\ | \text{ Some } v \Rightarrow \\ \quad \exists \text{vs}'. \\ \quad \text{vs} = \text{vs}' \uparrow [v] * \text{ws}' = \text{vs}' * \\ \quad \text{model } t \text{vs}' \\ \text{end} \end{array} \right]$$

$$\{ \text{res. res} = o * \text{owner } t \text{ws}' \}$$

# Strong Specification

POP-SPEC

$$\left\{ \begin{array}{l} \text{inv } t * \\ \text{owner } t \text{ } ws \end{array} \right\}$$
$$\left\{ \text{vs. model } t \text{ } vs \right\}$$

pop  $t$

$o \text{ } ws'$ . match  $o$  with

STEAL-SPEC

**This specification is *sufficient* to prove completion :**  
Once the scheduler terminates, all submitted tasks have finished.

$$\left\{ \begin{array}{l} \text{model } t \text{ (tail } vs) \\ \text{res. res = head } vs \end{array} \right\}$$

Some  $v \Rightarrow$

$\exists vs'$ .

$vs = vs' \# [v] * ws' = vs' * \text{model } t \text{ } vs'$

end

$$\left\{ \begin{array}{l} \text{res. res = } o * \\ \text{owner } t \text{ } ws' \end{array} \right\}$$

## Prophecy Variables / Prophets

- ▶ **Primitive prophets**  
Predict the future of the execution.
- ▶ **Typed prophets**  
Prediction belongs to a user-supplied semantic type.
- ▶ **Wise prophets**  
Remember past predictions.  
Eliminate inconsistent branches.
- ▶ **Multiplexed prophets**  
Separate logically independent predictions on a single prophet.

# Multiplexed Prophets

SNAPSHOT-GET

$$\frac{\text{model } pid \ \gamma \ \text{pasts } \text{prophss}}{\text{snapshot } \gamma \ k \ (\text{pasts } k) \ (\text{prophss } k)}$$

SNAPSHOT-VALID

$$\frac{\text{model } pid \ \gamma \ \text{pasts } \text{prophss} \quad \text{snapshot } \gamma \ k \ \text{past } \text{prophs}}{\exists \text{past}' . \text{pasts } k = \text{past} \# \text{past}' * \text{prophs} = \text{past}' \# \text{prophss } k}$$

RESOLVE-SPEC

$$\frac{\begin{array}{l} \text{atomic } e \\ \text{to-val } e = \text{None} \quad v = \text{prophet.to-val } \text{proph} \quad \text{model } pid \ \gamma \ \text{pasts } \text{prophss} \\ \text{wp}_{\mathcal{E}} e \left\{ \begin{array}{l} w. \forall \text{prophs}. \\ \text{prophss } k = \text{proph} :: \text{prophs} \rightarrow * \\ \text{model } pid \ \gamma \ (\text{alter } (\cdot \# [\text{proph}]) \ k \ \text{pasts}) \ (\text{prophss } [k \mapsto \text{prophs}]) \rightarrow * \\ \Phi \ w \end{array} \right. \end{array}}{\text{wp}_{\mathcal{E}} \text{Resolve } e \ pid \ (k, v) \ \{ \Phi \}}$$

Task Parallelism

Parabs: A Verified Realistic Parallel Scheduler

Chase-Lev Work-Stealing Deque

Future work

## Future work

- ▶ Functors
- ▶ Relaxed memory
- ▶ Other scheduling strategies
- ▶ Heterogeneous scheduling (CPUs & GPUs)

Thank you for your attention!