

Verifying Wait-Freedom for Concurrent Higher-Order Programs

ECOOP'26

Egor Namakonov, Lars Birkedal, Amin Timany

Aarhus University

June 8, 2026

Wait-freedom: the strongest progress guarantee

Wait-freedom: the strongest progress guarantee

“An [implementation of] operation [on a concurrent data structure] is wait-free if it guarantees that every call finishes its execution in a finite number of steps.”

Herlihy and Shavit (2012)

Wait-freedom: the strongest progress guarantee

“An [implementation of] operation [on a concurrent data structure] is wait-free if it guarantees that every call finishes its execution in a finite number of steps.”

Herlihy and Shavit (2012)

(* Does the increment in the left thread always terminate? *)

```
cnt.incr () || while true do cnt.incr (); done
```

Wait-freedom: the strongest progress guarantee

“An [implementation of] operation [on a concurrent data structure] is wait-free if it guarantees that every call finishes its execution in a finite number of steps.”
Herlihy and Shavit (2012)

```
(* Does the increment in the left thread always terminate? *)  
cnt.incr () || while true do cnt.incr (); done
```

```
let new_cnt_blocking () =  
  let c = ref 0 in  
  let l = new_lock () in  
  let i _ =  
    acquire l;  
    let n = !c in  
    c := n + 1;  
    release l;  
    n  
  in { incr = i }
```

X

Wait-freedom: the strongest progress guarantee

“An [implementation of] operation [on a concurrent data structure] is wait-free if it guarantees that every call finishes its execution in a finite number of steps.”
Herlihy and Shavit (2012)

```
(* Does the increment in the left thread always terminate? *)  
cnt.incr () || while true do cnt.incr (); done
```

```
let new_cnt_blocking () =  
  let c = ref 0 in  
  let l = new_lock () in  
  let i _ =  
    acquire l;  
    let n = !c in  
    c := n + 1;  
    release l;  
    n  
  in { incr = i }
```

X

```
let new_cnt_lockfree () =  
  let c = ref 0 in  
  let rec i _ =  
    let n = !c in  
    if CAS(c, n, n + 1)  
    then n  
    else i ()  
  in { incr = i }
```

X

Wait-freedom: the strongest progress guarantee

“An [implementation of] operation [on a concurrent data structure] is wait-free if it guarantees that every call finishes its execution in a finite number of steps.”
Herlihy and Shavit (2012)

(* Does the increment in the left thread always terminate? *)
cnt.incr () || while true do cnt.incr (); done

```
let new_cnt_blocking () =  
  let c = ref 0 in  
  let l = new_lock () in  
  let i _ =  
    acquire l;  
    let n = !c in  
    c := n + 1;  
    release l;  
    n  
  in { incr = i }
```

X

```
let new_cnt_lockfree () =  
  let c = ref 0 in  
  let rec i _ =  
    let n = !c in  
    if CAS(c, n, n + 1)  
    then n  
    else i ()  
  in { incr = i }
```

X

```
let new_cnt_waitfree () =  
  let c = ref 0 in  
  let i _ = FAA(c, 1)  
  in { incr = i }
```

✓

Wait-free implementations are the hardest

Blocking queue, Vindum and Birkedal (2021)

```
Let cg_enqueue lock head x =
  let rec enqueue_go x head =
    match head with
    | None
  -> Some (x, None)
    | Some p -> Some (fst p, enqueue_go x (snd p))
  in
  acquire lock;
  head := enqueue_go x !head;
  release lock

Let cg_dequeue lock head () =
  acquire lock;
  let v = match !head with
  | None
  -> None
    | Some p -> head := snd p; Some (fst p)
  in
  release lock;
  v

Let cg_queue () =
  let head = ref None in
  let lock = newLock () in
  (cg_dequeue lock head, cg_enqueue lock head x)
```

Wait-free implementations are the hardest

Blocking queue, Vindum and Birkedal (2021)

```
Let cg_enqueue lock head x =
  let rec enqueue_go x head =
    match head with
    | None
-> Some (x, None)
    | Some p -> Some (fst p, enqueue_go x (snd p))
  in
  acquire lock;
  head := enqueue_go x !head;
  release lock

let cg_dequeue lock head () =
  acquire lock;
  let v = match !head with
    | None
  | Some p -> head := snd p; Some (fst p)
  in
  release lock;
  v

let cg_queue () =
  let head = ref None in
  let lock = newLock () in
  (cg_dequeue lock head, cg_enqueue lock head x)
```

Lock-free Michael-Scott queue, version by Vindum and Birkedal (2021)

```
Let get_value x =
  match x with
  | None
-> let rec spin () = spin () in spin ()
  | Some v -> v

let ms_dequeue toSent toTail =
  let rec loop () =
    let head = !toSent in
    let tail = !toTail in
    let toNext = snd (get_value !head) in
    let p = NewProp in
    let next = !toNext in
    if head = (Resolve (!toSent) p ()) then
      if head = tail then
        match !next with
        | None -> None
        | Some _ -> CAS toTail tail next; loop ()
      end
      else
        if CAS toSent head next then
          Some (get_value (fst (get_value !next)))
        else loop ()
    else loop ()
  in loop ()

let ms_enqueue toTail x =
  let node = ref (Some (Some x, ref (ref None))) in
  let rec loop () =
    let tail = !toTail in
    let toNext = snd (get_value !tail) in
    let next = !toNext in
    if tail = !toTail then
      match !next with
      | None ->
        if CAS toNext next node then
          (CAS toTail tail node; ())
        else loop ()
      | Some _ -> CAS toTail tail next; loop ()
    end
  else loop ()
  in loop ()

let ms_queue () =
  let node
= ref (Some (None, ref (ref None))) in
  let toTail = ref node in
  let toSent = ref node in
  ((ms_dequeue toSent toTail), (ms_enqueue toTail))
```


Wait-freedom is hard to even define

“An operation is wait-free if it guarantees that every call finishes its execution in a finite number of steps.” [Herlihy and Shavit (2012)]

Wait-freedom is hard to even define

Is the data structure initialized?

What if operation was stored in heap?

How to define formally in an ML-like language?

"An **operation** is wait-free if it **guarantees** that **every** **call** **finishes** its **execution** in a finite number of steps." [Herlihy and Shavit (2012)]

What about safety?

Who calls it? Is it an arbitrary client program?

How the execution is scheduled?

Wait-freedom is hard to even define

Is the data structure initialized?

What if operation was stored in heap?

How to define formally in an ML-like language?

“An **operation** is wait-free if it **guarantees** that **every** **call** **finishes** its **execution** in a finite number of steps.” [Herlihy and Shavit (2012)]

What about safety?

Who calls it? Is it an arbitrary client program?

How the execution is scheduled?

$waitFree(op) : Prop$

Wait-freedom is hard to even define

Is the data structure initialized?

What if operation was stored in heap?

How to define formally in an ML-like language?

“An **operation** is wait-free if it **guarantees** that **every** **call** **finishes** its **execution** in a finite number of steps.” [Herlihy and Shavit (2012)]

What about safety?

Who calls it? Is it an arbitrary client program?

How the execution is scheduled?

$waitFree_c(op) : Prop$

need to account for initialization;
omit this parameter for simplicity

Wait-freedom is hard to even define, harder to prove.

 $\vdash \text{waitFree}(op)$

Wait-freedom is hard to even define, harder to prove.
Solution: verify the operation in Iris instead

$$\boxed{*} \vdash \text{WaitFreeSpec}(op) \Rightarrow \text{🦉} \vdash \text{waitFree}(op)$$

Wait-freedom is hard to even define, harder to prove.
Solution: verify the operation in Iris instead

$$\boxed{\star} \vdash \text{WaitFreeSpec}(op) \Rightarrow \text{🐙} \vdash \text{waitFree}(op)$$

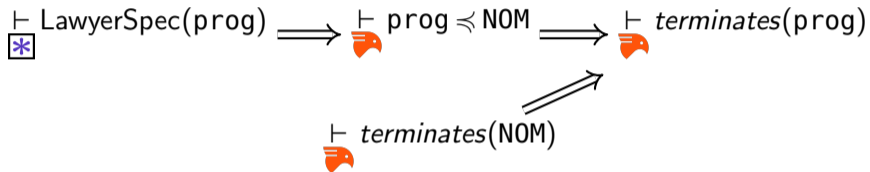
Our contributions, all mechanized in Rocq:

- Definition $\text{waitFree}(op)$ of op 's wait-freedom in higher-order setting;
- Iris-based specification $\text{WaitFreeSpec}(op)$ and the corresponding adequacy theorem;
- Case studies for simple algorithms;
- Adapting the approach for verifying realistic, *not exactly wait-free* algorithms

Iris Workshop'25: termination proofs with Lawyer logic

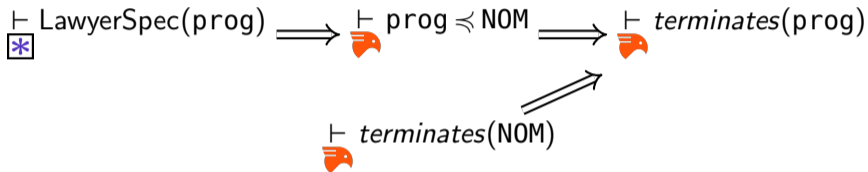
Iris Workshop'25: termination proofs with Lawyer logic

Lawyer approach, restricted to non-blocking concurrency:



Iris Workshop'25: termination proofs with Lawyer logic

Lawyer approach, restricted to non-blocking concurrency:



LTS	Language	NOM (“No-Obligations Model”)
State	thread pool and heap	finite <i>fuel</i> assigned to each thread
Labels	thread ids	thread ids
Transitions	small-step semantics	besides others, <i>burning fuel on every step</i>

Not done yet: specs imply weaker property for open programs

Not done yet: specs imply weaker property for open programs

Key property for refining closed prog to NOM

$$\boxed{*} \vdash \text{LawyerSpec}(\text{prog}) \implies \boxed{*} \vdash \begin{array}{l} \text{execution of prog preserves invariants} \\ \text{and refinement to NOM} \end{array}$$

Not done yet: specs imply weaker property for open programs

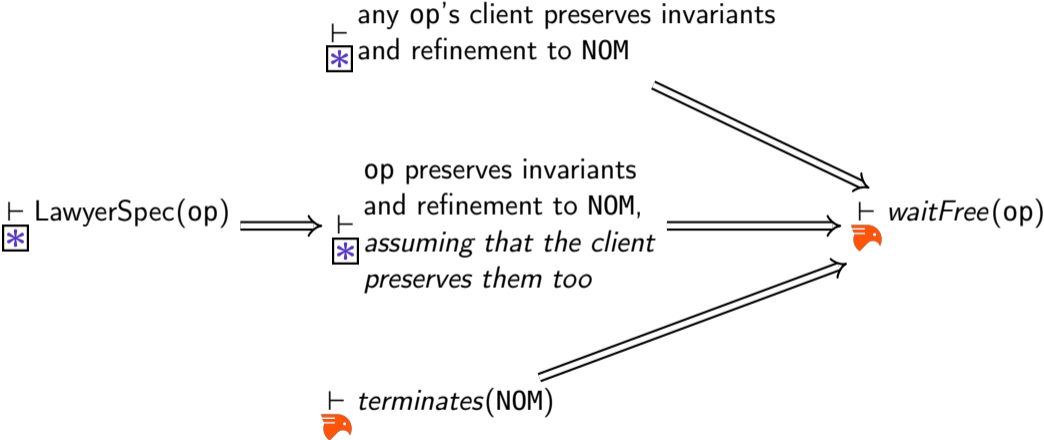
Key property for refining closed prog to NOM

$\boxed{*} \vdash \text{LawyerSpec}(\text{prog}) \implies \boxed{*} \vdash \begin{array}{l} \text{execution of prog preserves invariants} \\ \text{and refinement to NOM} \end{array}$

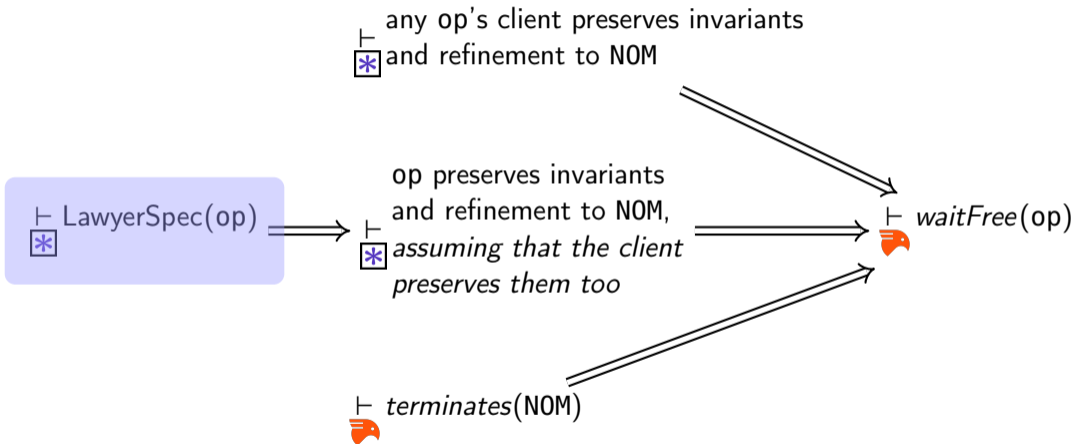
turns into a weaker property for standalone wait-free op:

$\boxed{*} \vdash \text{LawyerSpec}(\text{op}) \implies \boxed{*} \vdash \begin{array}{l} \text{execution of op preserves invariants} \\ \text{and refinement to NOM,} \\ \textit{assuming that client of op preserves them too} \end{array}$

User proves spec and applies our new adequacy theorem



User proves spec and applies our new adequacy theorem



Reusing Lawyer logic to prove wait-freedom specification

```
let incr x =
```

```
  FAA(c, 1)
```

Reusing Lawyer logic to prove wait-freedom specification

Key part of proving $\vdash \text{LawyerSpec}(\text{incr})$:
*

For every $\tau : \text{Thread}$
and $x : \text{Val}$:

$\{\text{⊖}_\tau * \text{⊖}_\tau * \boxed{\exists n : \mathbb{Z}. c \mapsto n}^{\mathcal{N}}\}^\tau$

let incr x =

FAA(c, 1)

{True}

Reusing Lawyer logic to prove wait-freedom specification

Key part of proving $\vdash \text{LawyerSpec}(\text{incr})$:
 \star

For every $\tau : \text{Thread}$
and $x : \text{Val}$:

$\{\text{fuel}_\tau * \text{fuel}_\tau * \exists n : \mathbb{Z}. c \mapsto n\}^\tau$

let incr x =

FAA(c, 1)

{True}

Amount of fuel and
concrete invariant
depend on operation

“phase” resource omitted
for simplicity

Reusing Lawyer logic to prove wait-freedom specification

Key part of proving $\vdash \text{LawyerSpec}(\text{incr})$:
*

For every $\tau : \text{Thread}$
and $x : \text{Val}$:

$\{ \text{[]}_\tau * \text{[]}_\tau * \boxed{\exists n : \mathbb{Z}. c \mapsto n} \}^\tau$

let incr x =

FAA(c, 1)

{True}

Reusing Lawyer logic to prove wait-freedom specification

Key part of proving $\vdash \text{LawyerSpec}(\text{incr})$:
 $\boxed{*}$

For every $\tau : \text{Thread}$
and $x : \text{Val}$:

$\{\text{⊖}_\tau * \text{⊖}_\tau * \boxed{\exists n : \mathbb{Z}. c \mapsto n}^{\mathcal{N}}\}^\tau$

let incr $x =$

$\{\text{⊖}_\tau\}$

FAA(c, 1)

$\{\text{True}\}$

$\{\text{True}\}$

$$\frac{\text{LAWYER-BETA-FUEL} \quad \{P\} e[v/x] \{Q\}_{\mathcal{E}}^\tau}{\{\text{⊖}_\tau * P\} (\lambda x. e) v \{Q\}_{\mathcal{E}}^\tau}$$

Reusing Lawyer logic to prove wait-freedom specification

Key part of proving $\vdash \text{LawyerSpec}(\text{incr})$:
 $\boxed{*}$

For every $\tau : \text{Thread}$
 and $x : \text{Val}$:

$\{\text{Irr}_\tau * \text{Irr}_\tau * \boxed{\exists n : \mathbb{Z}. c \mapsto n}^{\mathcal{N}}\}^\tau$

let incr $x =$

$\{\text{Irr}_\tau\}$

$\{\text{Irr}_\tau * \exists n. c \mapsto n\}^{\tau \setminus \mathcal{N}}$

FAA($c, 1$)

$\{\exists n. c \mapsto n\}^{\tau \setminus \mathcal{N}}$

$\{\text{True}\}$

$\{\text{True}\}$

LAWYER-BETA-FUEL

$\{P\} e[v/x] \{Q\}_{\mathcal{E}}^\tau$

$\frac{\{P\} e[v/x] \{Q\}_{\mathcal{E}}^\tau}{\{\text{Irr}_\tau * P\} (\lambda x. e) v \{Q\}_{\mathcal{E}}^\tau}$

IRIS-INV-ACC-TL

$\{P * Q\} e \{P * R\}_{\mathcal{E} \setminus \mathcal{N}}^\tau$ e is atomic $\text{timeless}(P)$ $\mathcal{N} \subseteq \mathcal{E}$

$\frac{\{P * Q\} e \{P * R\}_{\mathcal{E} \setminus \mathcal{N}}^\tau \quad e \text{ is atomic} \quad \text{timeless}(P) \quad \mathcal{N} \subseteq \mathcal{E}}{\{\boxed{P}^{\mathcal{N}} * Q\} e \{R\}_{\mathcal{E}}^\tau}$

Reusing Lawyer logic to prove wait-freedom specification

Key part of proving $\vdash \text{LawyerSpec}(\text{incr})$:

$\boxed{*}$

For every $\tau : \text{Thread}$
and $x : \text{Val}$:

$\{\text{m}_\tau * \text{m}_\tau * \boxed{\exists n : \mathbb{Z}. c \mapsto n}^{\mathcal{N}}\}^\tau$

let incr $x =$

$\{\text{m}_\tau\}$

$\{\text{m}_\tau * \exists n. c \mapsto n\}^{\tau \setminus \mathcal{N}}$

FAA($c, 1$)

$\{\exists n. c \mapsto n\}^{\tau \setminus \mathcal{N}}$

$\{\text{True}\}$

$\{\text{True}\}$

LAWYER-BETA-FUEL

$\{P\} e[v/x] \{Q\}_{\mathcal{E}}^\tau$

$\frac{\{P\} e[v/x] \{Q\}_{\mathcal{E}}^\tau}{\{\text{m}_\tau * P\} (\lambda x. e) v \{Q\}_{\mathcal{E}}^\tau}$

IRIS-INV-ACC-TL

$\{P * Q\} e \{P * R\}_{\mathcal{E} \setminus \mathcal{N}}^\tau$ e is atomic timeless(P) $\mathcal{N} \subseteq \mathcal{E}$

$\frac{\{P * Q\} e \{P * R\}_{\mathcal{E} \setminus \mathcal{N}}^\tau \quad e \text{ is atomic} \quad \text{timeless}(P) \quad \mathcal{N} \subseteq \mathcal{E}}{\{\boxed{P}^{\mathcal{N}} * Q\} e \{R\}_{\mathcal{E}}^\tau}$

LAWYER-FAA-FUEL

$\{\text{m}_\tau * \ell \mapsto n\} \text{FAA}(\ell, k) \{m. m = n * \ell \mapsto (n + k)\}_{\mathcal{E}}^\tau$

Reusing Lawyer logic to prove wait-freedom specification

Key part of proving $\vdash \text{LawyerSpec}(\text{incr})$:

\star

For every $\tau : \text{Thread}$
and $x : \text{Val}$:

$\{\text{m}_\tau * \text{m}_\tau * \boxed{\exists n : \mathbb{Z}. c \mapsto n}^{\mathcal{N}}\}^\tau$

let incr x =

$\{\text{m}_\tau\}$

$\{\text{m}_\tau * \exists n. c \mapsto n\}^{\tau \setminus \mathcal{N}}$

FAA(c, 1)

$\{\exists n. c \mapsto n\}^{\tau \setminus \mathcal{N}}$

$\{\text{True}\}$

$\{\text{True}\}$



LAWYER-BETA-FUEL

$\{P\} e[v/x] \{Q\}_{\mathcal{E}}^\tau$

$\frac{\{P\} e[v/x] \{Q\}_{\mathcal{E}}^\tau}{\{\text{m}_\tau * P\} (\lambda x. e) v \{Q\}_{\mathcal{E}}^\tau}$

IRIS-INV-ACC-TL

$\{P * Q\} e \{P * R\}_{\mathcal{E} \setminus \mathcal{N}}^\tau$ e is atomic timeless(P) $\mathcal{N} \subseteq \mathcal{E}$

$\frac{\{P * Q\} e \{P * R\}_{\mathcal{E} \setminus \mathcal{N}}^\tau \text{ e is atomic timeless}(P) \ \mathcal{N} \subseteq \mathcal{E}}{\{\boxed{P}^{\mathcal{N}} * Q\} e \{R\}_{\mathcal{E}}^\tau}$

LAWYER-FAA-FUEL

$\{\text{m}_\tau * \ell \mapsto n\} \text{FAA}(\ell, k) \{m. m = n * \ell \mapsto (n + k)\}_{\mathcal{E}}^\tau$

User proves spec and applies our new adequacy theorem

\vdash any op's client preserves invariants
and refinement to NOM

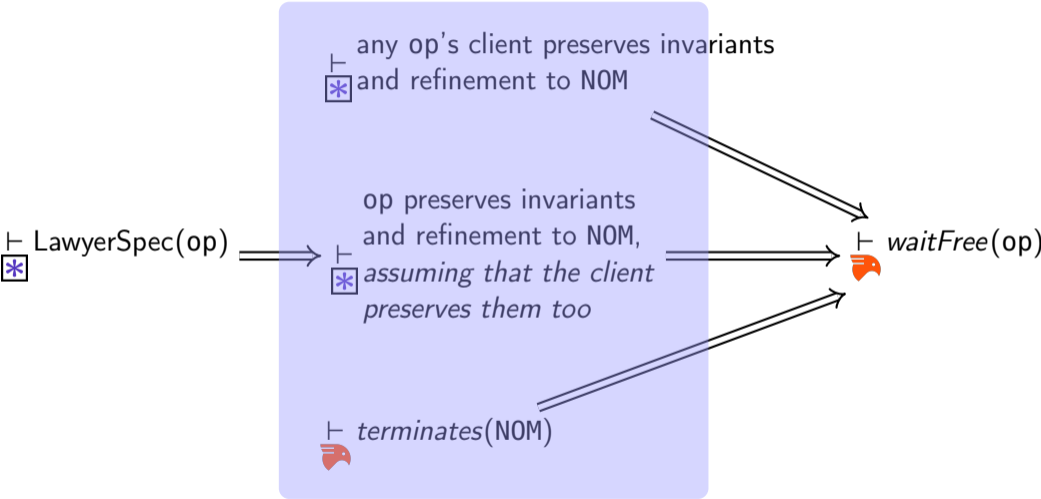
\vdash LawyerSpec(op)

\vdash op preserves invariants
and refinement to NOM,
*assuming that the client
preserves them too*

\vdash terminates(NOM)

\vdash waitFree(op)

User proves spec and applies our new adequacy theorem



User proves spec and applies our new adequacy theorem

any op's client preserves invariants

\vdash
 $\boxed{*}$ and refinement to NOM


\vdash
 $\boxed{*}$ LawyerSpec(op)



\vdash
 $\boxed{*}$ op preserves invariants
and refinement to NOM,
*assuming that the client
preserves them too*



\vdash waitFree(op)

\vdash
 terminates(NOM)



Invariants of the wait-free operation must be preserved

E.g. safe termination of **let** incr x = **FAA**(c, 1)
relies on $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ being preserved by its client

Invariants of the wait-free operation must be preserved, *i.e.* the operation should be “robustly safe”

E.g. safe termination of **let** incr x = **FAA**(c, 1)
relies on $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ being preserved by its client

Robust safety of verified op, more formally:

$\vdash \forall \text{prog} : \text{Expr}. \text{LawyerSpec}(\text{op}) \dashv^* \text{PresInv}(\text{prog}[\text{op}/x])$
 $\boxed{*}$

op's client

Invariants of the wait-free operation must be preserved, *i.e.* the operation should be “robustly safe”

E.g. safe termination of **let** incr x = **FAA**(c, 1)
relies on $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ being preserved by its client

Robust safety of verified op, more formally:

$\vdash \forall \text{prog} : \text{Expr}. \text{LawyerSpec}(\text{op}) \multimap \text{PresInv}(\text{prog}[\text{op}/x])$
 $\boxed{*}$

op's client

How to prove it for any prog?

In fact, most of clients preserve operation's invariants

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

In fact, most of clients preserve operation's invariants

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

In fact, most of clients preserve operation's invariants

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

In fact, most of clients preserve operation's invariants

`incr ()`

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

Calling the verified operation



In fact, most of clients preserve operation's invariants

```
inc () || let n = 5 + 5 in inc ()  
inc () || inc n; || inc ()
```

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

Calling the verified operation (+ in multiple threads, with any argument) ✓

In fact, most of clients preserve operation's invariants

```
inc () || let n = 5 + 5 in || while true do
      || incr n;           ||   incr ()
      ||                   ||   done
```

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

Calling the verified operation (+ in multiple threads, with any argument) ✓

Running infinitely ✓

In fact, most of clients preserve operation's invariants

```

() ()      ||      ||      let n = 5 + 5 in      ||      while true do
           ||  incr ()  ||  incr n;                ||      incr ()
           ||      ||      ||      ||      done
```

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

Calling the verified operation (+ in multiple threads, with any argument) ✓

Running infinitely ✓

Getting stuck outside of the operation ✓

In fact, most of clients preserve operation's invariants

```
let l = ref 0 in ||           || let n = 5 + 5 in || while true do
l := true;           || incr () || incr n;           ||   incr ()
() ()                ||           ||                   ||   done
```

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

Calling the verified operation (+ in multiple threads, with any argument) ✓

Running infinitely ✓

Getting stuck outside of the operation ✓

Writing to *fresh* locations ✓

In fact, most of clients preserve operation's invariants

```
let l = ref 0 in || || let n = 5 + 5 in || while true do
l := true; || incr () || incr n; || incr ()
() () || || [n] := true || done
```

What preserves the invariant $\boxed{\exists n : \mathbb{Z}. c \mapsto n}$ of `incr`?

Calling the verified operation (+ in multiple threads, with any argument) ✓

Running infinitely ✓

Getting stuck outside of the operation ✓

Writing to *fresh* locations ✓

Writing to *arbitrary* locations ✗

Prove robust safety with logical relations

Prove robust safety with logical relations

$RS_E(\text{prog} : Expr) \approx$

execution of prog preserves invariants,
might get stuck,
and the returned value (if any) is “safe”

$RS_V(v : Val) \approx$

v can be used in safe programs
(e.g. location allocated by the client)

Prove robust safety with logical relations

$RS_E(\text{prog} : Expr) \approx$

execution of prog preserves invariants,
might get stuck,
and the returned value (if any) is “safe”

$RS_V(v : Val) \approx$

v can be used in safe programs
(e.g. location allocated by the client)

Break this circularity by defining RS_E and RS_V with more primitive Iris propositions!

Prove robust safety with logical relations

$RS_E(\text{prog} : Expr) \approx$

execution of prog preserves invariants,
might get stuck,
and the returned value (if any) is “safe”

$RS_V(v : Val) \approx$

v can be used in safe programs
(e.g. location allocated by the client)

Break this circularity by defining RS_E and RS_V with more primitive Iris propositions!
With that, show a general property:

$$\boxed{*} \vdash \forall \text{prog}, \vec{v}. \text{NoPtrArith}(\text{prog}) \multimap RS_V(\vec{v}) \multimap RS_E(\text{prog}[\vec{v} / \overline{\text{freeVars}(\text{prog})}])$$

Prove robust safety with logical relations

$RS_E(\text{prog} : Expr) \approx$

execution of prog preserves invariants,
might get stuck,
and the returned value (if any) is “safe”

$RS_V(v : Val) \approx$

v can be used in safe programs
(e.g. location allocated by the client)

Break this circularity by defining RS_E and RS_V with more primitive Iris propositions!
With that, show a general property:

$$\boxed{*} \vdash \forall \text{prog}, \vec{v}. \text{NoPtrArith}(\text{prog}) \multimap RS_V(\vec{v}) \multimap RS_E(\text{prog}[\vec{v} / \overline{\text{freeVars}(\text{prog})}])$$

and conclude robust safety of verified op:

$$\boxed{*} \vdash \forall \text{prog}. \text{NoPtrArith}(\text{prog}) \multimap RS_V(\text{op}) \multimap RS_E(\text{prog}[\text{op}/x])$$

by adapting the proof of `LawyerSpec(op)`

User proves spec and applies our new adequacy theorem

any op's client preserves invariants

\vdash
 $\boxed{*}$ and refinement to NOM


\vdash
 $\boxed{*}$ LawyerSpec(op)



\vdash
 $\boxed{*}$ op preserves invariants
and refinement to NOM,
*assuming that the client
preserves them too*

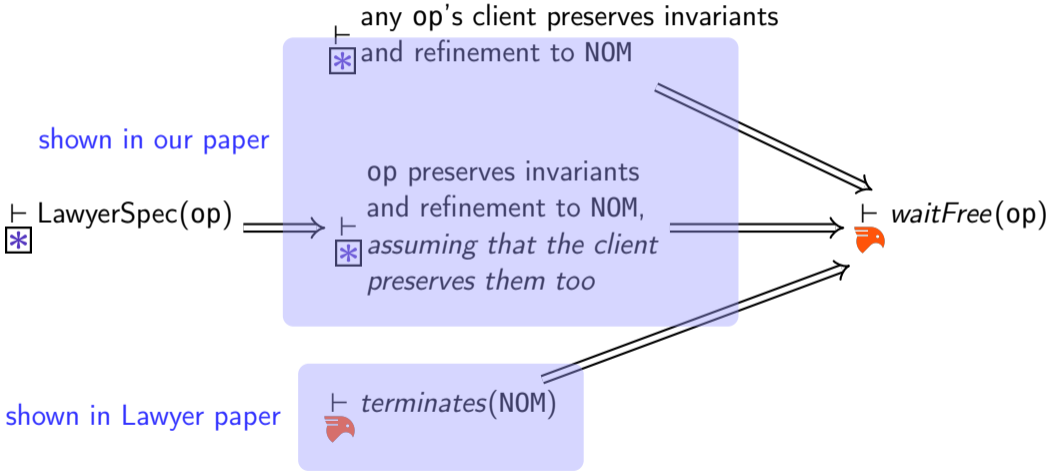


\vdash waitFree(op)

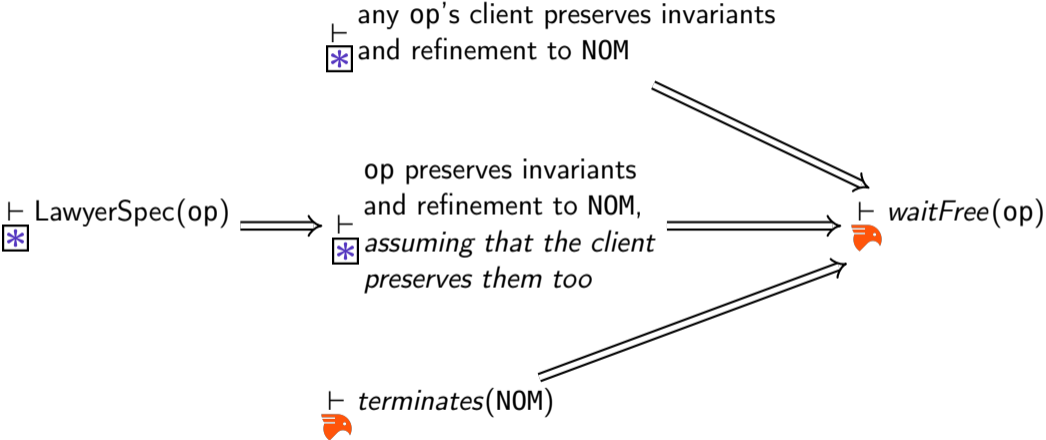
\vdash
 terminates(NOM)



User proves spec and applies our new adequacy theorem



User proves spec and applies our new adequacy theorem



More complicated algorithms are not textbook-wait-free

More complicated algorithms are not textbook-wait-free

- Wait-free operation might get stuck if it actually inspects the input argument (e.g., `list_map f l` with wait-free `f` when `l` is not a list)

More complicated algorithms are not textbook-wait-free

- Wait-free operation might get stuck if it actually inspects the input argument (e.g., `list_map f l` with wait-free `f` when `l` is not a list)
- Realistic wait-free data structures restrict concurrency ...

Realistic wait-free data structures restrict concurrency

```
void enq(int value) {  
    long phase = maxPhase() + 1;  
    state.set(TID, new  
        OpDesc(phase, true, true, new Node(value, TID)));  
    help(phase);  
    help_finish_enqueue();  
}
```

1. Post its own operation
2. Help previously posted operations
3. Complete its own operation

```
AtomicReferenceArray<OpDesc> state;
```

```
WFQueue () {  
    state = new AtomicReferenceArray<OpDesc>(NUM_THRDS);  
    ■ ■ ■
```

Adapting our approach to restricted wait-freedom

Adapting our approach to restricted wait-freedom

Result for 1-enqueuer, 1-dequeuer queue by Jayanti and Petrovic (2005):

$\vdash \text{waitFreeRestr}([\text{enqueuer}, \text{dequeuer}])$



Adapting our approach to restricted wait-freedom

Result for 1-enqueuer, 1-dequeuer queue by Jayanti and Petrovic (2005):

$\vdash \text{waitFreeRestr}([\text{enqueuer}, \text{dequeuer}])$



similar to $\text{waitFree}(\text{ops}[i])$ for every i ,
 $\text{waitFreeRestr}(\text{ops} : \text{List Val}) \approx$ but i th thread can only call $\text{ops}[i]$;
forks are prohibited, but $\text{ops}[i] = \text{ops}[j]$ is allowed

Adapting our approach to restricted wait-freedom

Result for 1-enqueuer, 1-dequeuer queue by Jayanti and Petrovic (2005):

$\vdash \text{waitFreeRestr}([\text{enqueuer}, \text{dequeuer}])$



similar to $\text{waitFree}(\text{ops}[i])$ for every i ,
 $\text{waitFreeRestr}(\text{ops} : \text{List Val}) \approx$ but i th thread can only call $\text{ops}[i]$;
forks are prohibited, but $\text{ops}[i] = \text{ops}[j]$ is allowed

$\text{WaitFreeRestrSpec}(\text{ops}) \approx$ for each $\text{op} \in \text{ops}$, show $\forall a, \tau. \{\exists_{\tau} * \dots * \exists_{\tau} * \text{tok}_{\text{op}}\} \text{op } a \{\text{tok}_{\text{op}}\}^{\tau}$
assuming there are only $\text{count}(\text{op}, \text{ops})$ many tok_{op}

$\boxed{*} \vdash \text{WaitFreeRestrSpec}(\text{ops}) \Rightarrow \text{waitFreeRestr}(\text{ops})$



Adapting our approach to restricted wait-freedom

Result for 1-enqueuer, 1-dequeuer queue by Jayanti and Petrovic (2005):

$\vdash \text{waitFreeRestr}([\text{enqueuer}, \text{dequeuer}])$



similar to $\text{waitFree}(\text{ops}[i])$ for every i ,
 $\text{waitFreeRestr}(\text{ops} : \text{List Val}) \approx$ but i th thread can only call $\text{ops}[i]$;
forks are prohibited, but $\text{ops}[i] = \text{ops}[j]$ is allowed

$\text{WaitFreeRestrSpec}(\text{ops}) \approx$ for each $\text{op} \in \text{ops}$, show $\forall a, \tau. \{\exists_{\tau} * \dots * \exists_{\tau} * \text{tok}_{\text{op}}\} \text{op } a \{\text{tok}_{\text{op}}\}^{\tau}$
assuming there are only $\text{count}(\text{op}, \text{ops})$ many tok_{op}

$\vdash \text{WaitFreeRestrSpec}(\text{ops}) \Rightarrow \vdash \text{waitFreeRestr}(\text{ops})$



Logical relation and adequacy proof change, but the program logic stays the same!

Future work

- Linearizability in the same framework (requires prophecy variables in Trillium)
- Verification of universal constructions for wait-freedom
- Minor changes, *e.g.*, being more flexible with amount of fuel

Wait-freedom is hard to even define, harder to prove.
Solution: verify the operation in Iris instead

$$\boxed{\star} \vdash \text{WaitFreeSpec}(op) \Rightarrow \text{🐙} \vdash \text{waitFree}(op)$$

Our contributions, all mechanized in Rocq:

- Definition $\text{waitFree}(op)$ of op 's wait-freedom in higher-order setting;
- Iris-based specification $\text{WaitFreeSpec}(op)$ and the corresponding adequacy theorem;
- Case studies for simple algorithms;
- Adapting the approach for verifying realistic, *not exactly wait-free* algorithms

Wait-freedom is hard to even define, harder to prove. Solution: verify the operation in Iris instead

$$\boxed{\star} \vdash \text{WaitFreeSpec}(op) \Rightarrow \text{🦉} \vdash \text{waitFree}(op)$$

Our contributions, all mechanized in Rocq:

- Definition $\text{waitFree}(op)$ of op 's wait-freedom in higher-order setting;
- Iris-based specification $\text{WaitFreeSpec}(op)$ and the corresponding adequacy theorem;
- Case studies for simple algorithms;
- Adapting the approach for verifying realistic, *not exactly wait-free* algorithms

Thank you!

Backup slides

Logical relation

$$RS_E(e) \triangleq \forall \tau. \{True\}^{\downarrow} e \{v. RS_V(v)\}^{\tau}$$

$$RS_V(()) \triangleq RS_V(b : \mathbb{B}) \triangleq \dots \triangleq True$$

...

$$RS_V((v_1, v_2)) \triangleq RS_V(v_1) * RS_V(v_2)$$






$$RS_V(\ell : Loc) \triangleq \boxed{(\exists v. \ell \mapsto v * RS_V(v)) \vee freed(\ell)}$$

$$RS_V(\lambda x. e) \triangleq \forall \tau, v. \{RS_V(v)\}^{\downarrow} (\lambda x. e)v \{r. RS_V(r)\}^{\tau}$$

The \downarrow triples are different from those used in Lawyer proofs:

- they don't enforce model refinement, *i.e.* only consider safety
- they allow to get stuck

Bibliography

-  Herlihy, Maurice and Nir Shavit (2012). *The Art of Multiprocessor Programming, Revised Reprint*.
-  Jayanti, Prasad and Srdjan Petrovic (2005). “Logarithmic-time single deleter, multiple inserter wait-free queues and stacks”. In: DOI: [10.1007/11590156_33](https://doi.org/10.1007/11590156_33).
-  Kogan, Alex and Erez Petrank (2011). “Wait-free queues with multiple enqueueers and dequeuers”. In: DOI: [10.1145/1941553.1941585](https://doi.org/10.1145/1941553.1941585).
-  Namakonov, Egor, Lars Birkedal, and Amin Timany (June 2026). “Verifying wait-freedom for concurrent higher-order programs”. In: DOI: doi.org/10.4230/LIPIcs.EC00P.2026.30.
-  Namakonov, Egor, Justus Fasse, et al. (Apr. 2026). “Lawyer: Modular Obligations-Based Liveness Reasoning in Higher-Order Impredicative Concurrent Separation Logic”. In: DOI: [10.1145/3798240](https://doi.org/10.1145/3798240).
-  Vindum, Simon Friis and Lars Birkedal (2021). “Contextual refinement of the Michael-Scott queue (proof pearl)”. In: DOI: [10.1145/3437992.3439930](https://doi.org/10.1145/3437992.3439930).