

Neurosymbolic C++ Verification using Rocq

Lennart Beringer Jasper Haag Rodolphe Lepigre **Gregory Malecha** Rayan Soban Ehtesham Zahoor

SkyLabs AI

What if you had to verify this?

- C++
- Low-level code
- Fine-grained concurrency
- Hardware interaction

```
/*
 * Lookup OBJ capability for the specified selector
 *
 * @param sel Selector whose capability is being looked up
 * @return Object Capability (if slot is non-empty) or Null Capability (otherwise)
 */
Capability Space_obj::lookup (unsigned long sel) const
{
    auto l { lev }; uintptr_t cte;

    // Walk down the capability tables from the root, computing the slot index at each level
    for (auto ptr { &root }; ptr = &table (cte)->slot[(sel >> --l * bpl) % Captable::entries] ) {

        // Return capability upon reaching the last existing or leaf level
        if (!(cte = *ptr) || !l)
            return Capability { cte };
    }
}
```

```
static constexpr auto entries { BIT (bpl) };

using entry_t = Atomic<uintptr_t, __ATOMIC_ACQUIRE>;
entry_t slot[entries] {};
```

What if you had to verify this?

- C++
- Low-level code
- Fine-grained concurrency
- Hardware interaction
- **Strong specifications**

```
/*
 * Lookup OBJ capability for the specified selector
 */

(** The continuation [K] lets the caller relate the capability
    returned by [lookup] to the capability [cap] known (by the
    caller's proof) to be bound to [sel]. *)
Definition do_lookup' (E : coPset) (g : spc.obj.gname) (sel : sel.t)
  (K : capability.t → mpred) : mpred :=
  AR1 << 33 q cap, space_obj.entry g sel q cap >> (** current entry for [sel]
    @ top \ E, empty
    << K cap >>.
Definition do_lookup := do_lookup' (↑tableN).
cpp.spec "Space_obj::lookup(unsigned long) const" as lookup from (module) with
  \this this
  \spec@{WpSpec_cpp}
  \arg{sel} "sel" (sel.V sel) (** selector to look up *)
  \pre{g} this |-> space_obj.I g (** space invariant *)
  \prepost{q pd} this |-> space_obj.R g q pd (** space representation *)
  \prepost const.all
  \pre{K} do_lookup g sel K
  \post{v (cap : capability.t)}[Vref v]
  [| published cap |] **
  K cap ** (** found capability [cap] *)
  v |-> capability.R 1$m cap (** and returned it in value [v] *)
).
```

What if you had to verify this?

- C++
- Low-level code
- Fine-grained concurrency
- Hardware interaction
- Strong specifications

```
/*
 * Lookup OBJ capability for the specified selector
 */
(** Lemma lookup_ok : verify[module] "Space_obj::lookup(unsigned long) const".
    Proof using All.
    verify_spec.
    go.

Defi
  wp_for (λ p ⇒
  AR
    Exists p (n : Z),
    _local p "ptr" |-> ptrR<entryT> 1$m p **
    p |-> type_ptrR entryT **
  Defi
  cpp.
    [| (0 ≤ n ≤ sel.lev)%Z |] **
    \t
    _local p "cte" |-> anyR Tuintptr_t 1$m **
    \s
    atomic.do_load (t:=Tuintptr_t) (temp_args:=atomic.default_args)
    \a
    p (do_load_K cns g q sel K n (base_idx n sel))
    \p
    ).
    \p
    go.
    \p
    rewrite {1}/do_load_K /read_case.
    \p
    vc_split.
    { (* return with nullptr *) go. }

    go.
    rewrite /captableNonNull /peeked_slot_type_ptr.
    wp_if.
    { (* return with l = 0 *) go. }
    (* continue to loop *)
    go.
Qed.
```

<https://bluerocksec.gitlab.io/formal-methods/blogs/2026-03-31-NOVA-relaxed-memory/>

NOVA is getting relaxed-memory ready (26.08.0)

March 31, 2026

Tags: [NOVA-proofs](#), [NOVA-status](#)

From the February 2026 release ([26.08.0](#)), the implementation of NOVA has been fixed with more correct uses of relaxed memory accesses. More fixes and documentation will come in the next releases. Accordingly, the formal spec ([nova_interface](#)) has been updated to expose relaxed memory effects that cannot be hidden to client code. The formal proof ([nova_proof](#)) has also been updated to the new spec, with several remaining TODOs concerning the incomplete specification of RCU (garbage collection).

Release-Acquire is now the default in the implementation

Previous releases of NOVA have been using either relaxed accesses ([memory_order_relaxed](#)) or sequentially consistent accesses (SC, [memory_order_seq_cst](#)), which are either too weak or too strong in their use cases. From the February 2026 release, NOVA has switched to use release-acquire ([memory_order_acquire](#), [memory_order_release](#), [memory_order_acq_rel](#)) as the default.

NOVA now uses relaxed accesses in case no synchronization is required, such as updates to an `Ec`'s continuation field when in sequential mode, or external synchronization is required from clients, such as in [ctrl_pt](#). Meanwhile, NOVA uses sequentially consistent accesses to forbid undesirable reordering behaviors, such as when the per-CPU `current` `Ec` field is updated (context-switching) concurrently with other CPUs reading that field to check if a TLB invalidation or [a recall](#) is needed.

The formal spec exposes relaxed behaviors

The uses of weaker-than-SC accesses now expose weaker behaviors to NOVA clients. For example, if clients now use [ctrl_pt](#) concurrently, stale values of a `Pt`'s PID or MTD can be observed, and clients need to provide sufficient external synchronization to observe the latest values. The [updated nova_spec](#) for `ctrl_pt` now exposes a history of readable values for both PID and MTD. From this spec, clients (such as the BlueRock virtualization user code) can derive a sequential spec where the latest value will be observed.

The formal proof uses iRC11 relaxed memory logic

To verify code using relaxed memory, the `nova_proof` [axiomatizes](#) the rules of the [iRC11](#) concurrent separation logic for relaxed memory. The logic allows us to reason about relaxed and release/acquire accesses in terms of observable histories of atomic

What if you verify this?

- C++
- Low-level code
- Fine-grained
- Hardware inter
- Strong specific

NOVA Microhypervisor:

const

Just the beginning...

Many examples in this talk are deliberately small for illustration purposes.

However, these tools are built for working with production code.

- NOVA Microhypervisor: <https://gitlab.com/bluerocksec/NOVA>
- Blockchain: <https://github.com/category-labs/monad>
- Device Drivers: <https://sel4.systems/Summit/2025/slides/verified-zynqmp.pdf>
- C++ Standard Library: <https://github.com/skylabsai/brick-libc++>

* https://bluerocksec.gitlab.io/formal-methods/tech_reports/

Modular, Full-System Verification *

Gregory Malecha
BlueRock Security, Inc
USA
gregory@bluerock.io

Simon Hudon
BlueRock Security, Inc
USA
simon@bluerock.io

Hoang-Hai Dang
BlueRock Security, Inc
Germany
hai@bluerock.io

Jan-Oliver Kaiser
BlueRock Security, Inc
Germany
janno@bluerock.io

Paolo G. Giarrusso
BlueRock Security, Inc
Germany
paolo@bluerock.io

David Swasey*
Riverside Research, Inc
USA
pswasey@riversideresearch.org

ABSTRACT

We present an approach to specifying operating systems that is both highly modular and supports deriving various properties necessary to modern operating systems. The approach combines a machine semantics decomposed by privilege levels with a separation logic specification of the operating system API. We describe how the specification style enables natural proofs of robust safety in addition to deep behavioral refinements of user-mode applications running atop the OS. This approach enables simple and flexible concurrent specifications and unlocks new opportunities for whole-system verification.

KEYWORDS

Operating system specification, Formal verification, Multi-core verification, Separation logic, Iris

ACM Reference Format:

Gregory Malecha, Hoang-Hai Dang, Paolo G. Giarrusso, Simon Hudon, Jan-Oliver Kaiser, and David Swasey. 2025. Modular, Full-System Verification. In *Workshop in Hot Topics in Operating Systems (HOTOS 25)*, May 14–16, 2025, Banff, AB, Canada. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3713082.3730387>

1 INTRODUCTION

Systems-software verification has been an active area of research for almost two decades. Existing work lays a founda-

pieces of software. The verifications are built around monolithic (whole-program) proof principles [11] or rely on highly stylized coding patterns [8]. They are also limited in their ability to capture concurrency and rely on either completely eliminating it [11] or require all specifications to be linearizable [8]. While these approaches simplify reasoning, they do not account for the complexities of modern systems.

We aim for an operating system *specification* that is

- (1) **Realistic**. Modern OS APIs are rarely atomic, and concurrently executing system calls can lead to subtle behaviors that the specification must account for. These include weak memory behaviors that are pervasive on modern platforms [18, 21, 20].
- (2) **Two-sided** [1]. It must be usable not only to verify the operating system itself, but also to support user-mode reasoning on top of the OS. In particular, clients should be able to prove strong functional correctness properties of user-mode code, including behavioral refinements, using only the OS specification.
- (3) **Modular**. It must scale to support the large and complex APIs present in modern systems. This demands both thread- and component-local specifications that are capable of encapsulating implementation details.
- (4) **Robustly Safe** [7]. It must support running *arbitrary, untrusted* user-mode code and enable user-mode applications to use the OS specification to establish their

Virtual Machine
a full computer
target for attack.
s. This VMM is
naturally aligns
within separation
g a combination

e monitor

ning multiple
lications. The
hines—makes
re of the code
by proving the
o ensuring the
ure-compliant

etal property™
or of the VMM
metal machine
not just a CPU
ll as hardware
ecture, but the
s).
strength VMM.
g the program-
eparation logic

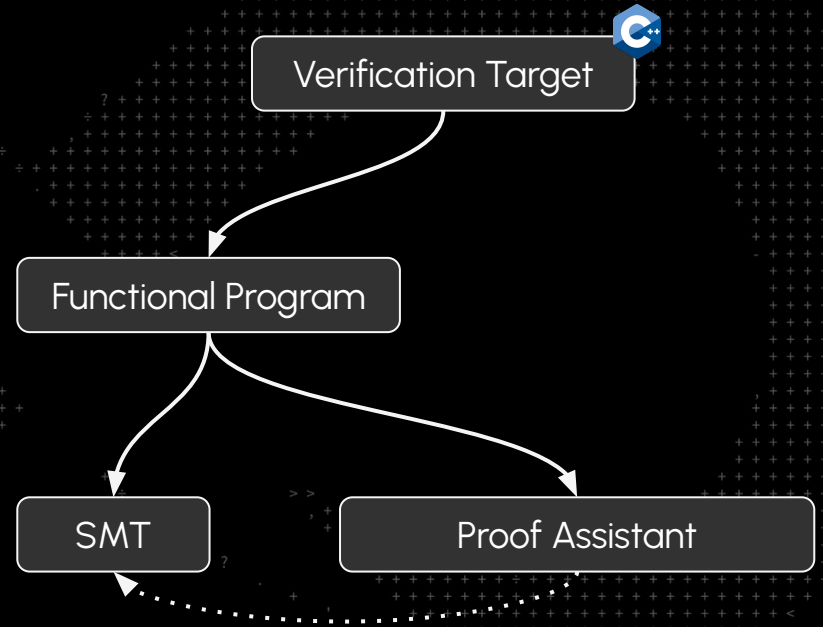
Program Verification Strategies

Most verification methodologies go through functional programs.

- Clean semantics
- Great tools for automatic and auto-active verification, e.g. SMT, Dafny, etc.

Functional programs are not an ideal verification IR.

- Translation strategy is often fixed ahead of time.
- Concurrency is difficult to model and reason about.



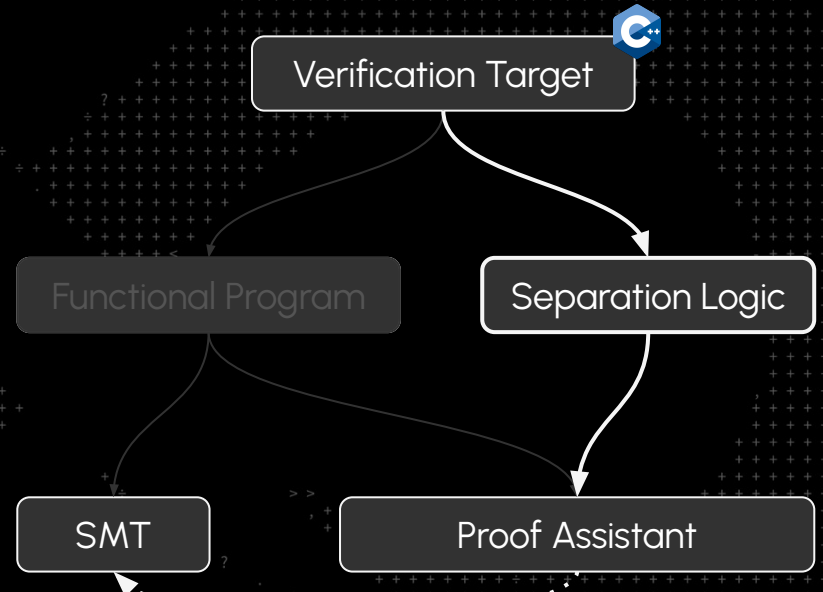
Program Verification Strategies

Separation logic is the logic independence.

- A (universal) **ownership discipline** on program state.
- Dynamic semantics by consuming & producing resources.
- **Language-agnostic** formalism

Tools built on separation logic:

- VST, Viper, RefinedC, RefinedRust, etc.
- BRiCk – more on this later



A Foundation for Software Verification

Proof assistants provide

- Highly expressive language for specifications and proofs
- A way to **integrate evidence** with a small TCB
- Domain-relevant theories, e.g. floating point, labeled transition systems, etc.

Proof Assistant
(Rocq)

A Foundation for Software Verification

Iris is the state-of-the-art in separation logic

- Foundational library for separation logic
- Step-indexed logic for impredicative reasoning
- Powerful support for custom CMRAs
- Proof mode (IPM) for doing “manual” proofs

Separation Logic
(Iris)

Proof Assistant
(Rocq)

A Foundation for Software Verification

Define a logic for the programming language.

- A (universal) **ownership discipline** on program state.
- **Predicate transformers** to express the behavior of **language constructs**.
- Highly modular. Can easily support language subsets.



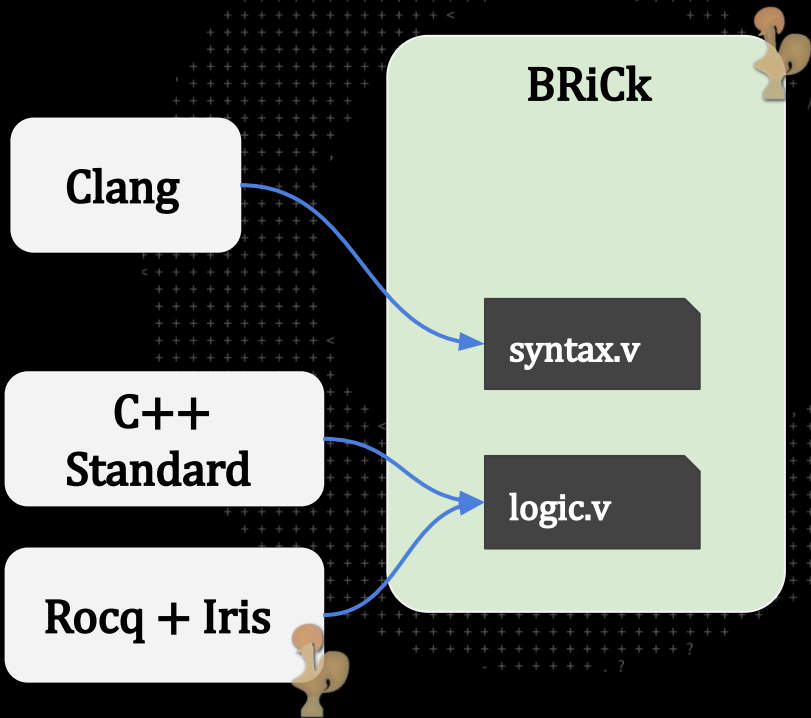
Program Logic
(Language)

↓
Separation Logic
(Iris)

Proof Assistant
(Rocq)

BRiCk C++ Infrastructure

<https://github.com/skylabsai/brick>



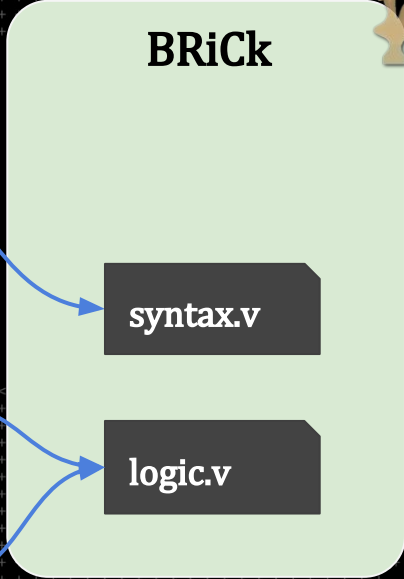
BRiCK C++ Infrastruct

<https://github.com/skylabsai/brick>

Clang

C++
Standard

Rocq + Iris



```
(** [Cl2r] represents reads of locations.  
This counts as an _access_, so it must happen at one of the types listed in  
https://eel.is/c++draft/basic.lval#11.  
*)  
Axiom wp_operand_cast_l2r :  $\forall e Q,$   
  (letI* a, free := wp_glval e in  
    $\exists v,$   
    ( $\exists q, a \mid\rightarrow$  tptsto_fuzzyR (erase_qualifiers $ type_of e) q v ** T)  $\wedge$   
    Q v free)  
   $\vdash$  wp_operand (Ecast Cl2r e) Q.
```

Program state reflected in separation logic.

```
(** The '*' operator is an lvalue, but this does not perform an  
access (see [wp_operand_cast_l2r] instead).
```

```
We check pointer [p] is strictly valid and aligned.  
The standard says (<https://eel.is/c++draft/expr.unary.op#1>):
```

- > The unary * operator performs indirection: the expression
- > to which it is applied shall be a pointer to an object type,
- > or a pointer to a function type and the result is an lvalue
- > referring to the object or function to which the expression
- > points. If the type of the expression is "pointer to T", the
- > type of the result is "T".

Object lifetime and
bounds checking

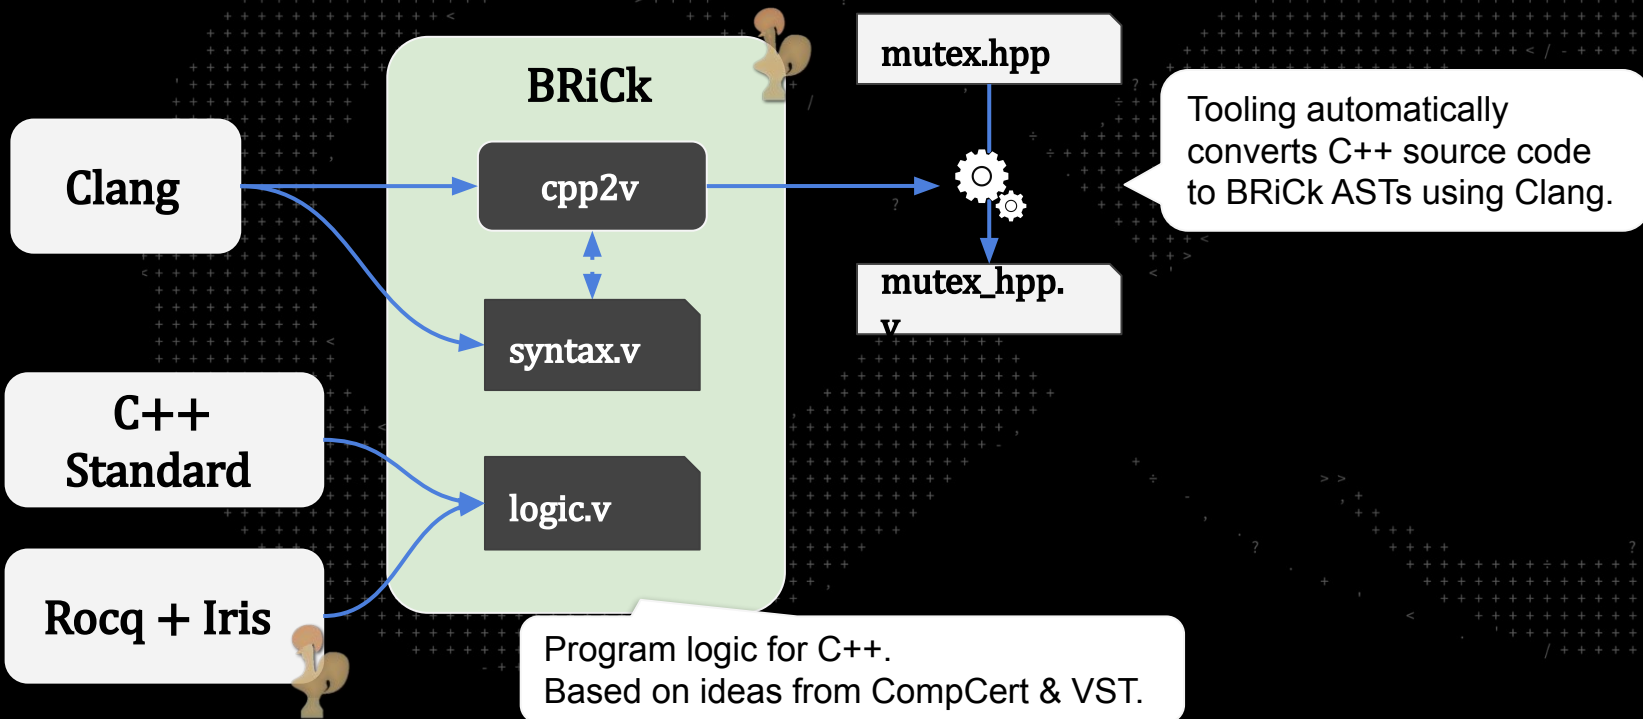
```
*)  
Axiom wp_lval_deref :  $\forall ty e Q,$   
  letI* v , free := wp_operand e in  
   $\exists p, [ | v = Vptr p | ]$  **  
  reference_to (erase_qualifiers ty) p **  
  Q p free  
   $\vdash$  wp_lval (Ederof e ty) Q.
```

Evaluation order

```
(** Evaluation of binary operators. *)  
Axiom wp_operand_binop :  $\forall o e_1 e_2 ty Q,$   
  letI* '(v1,v2), free := nd_seq (wp_operand e1) (wp_operand e2) in  
   $\exists v',$   
  (eval_binop tu o (type_of e1) (type_of e2) ty v1 v2 v' ** T)  $\wedge$   
  Q v' free)  
   $\vdash$  wp_operand (Ebinop o e1 e2 ty) Q.
```

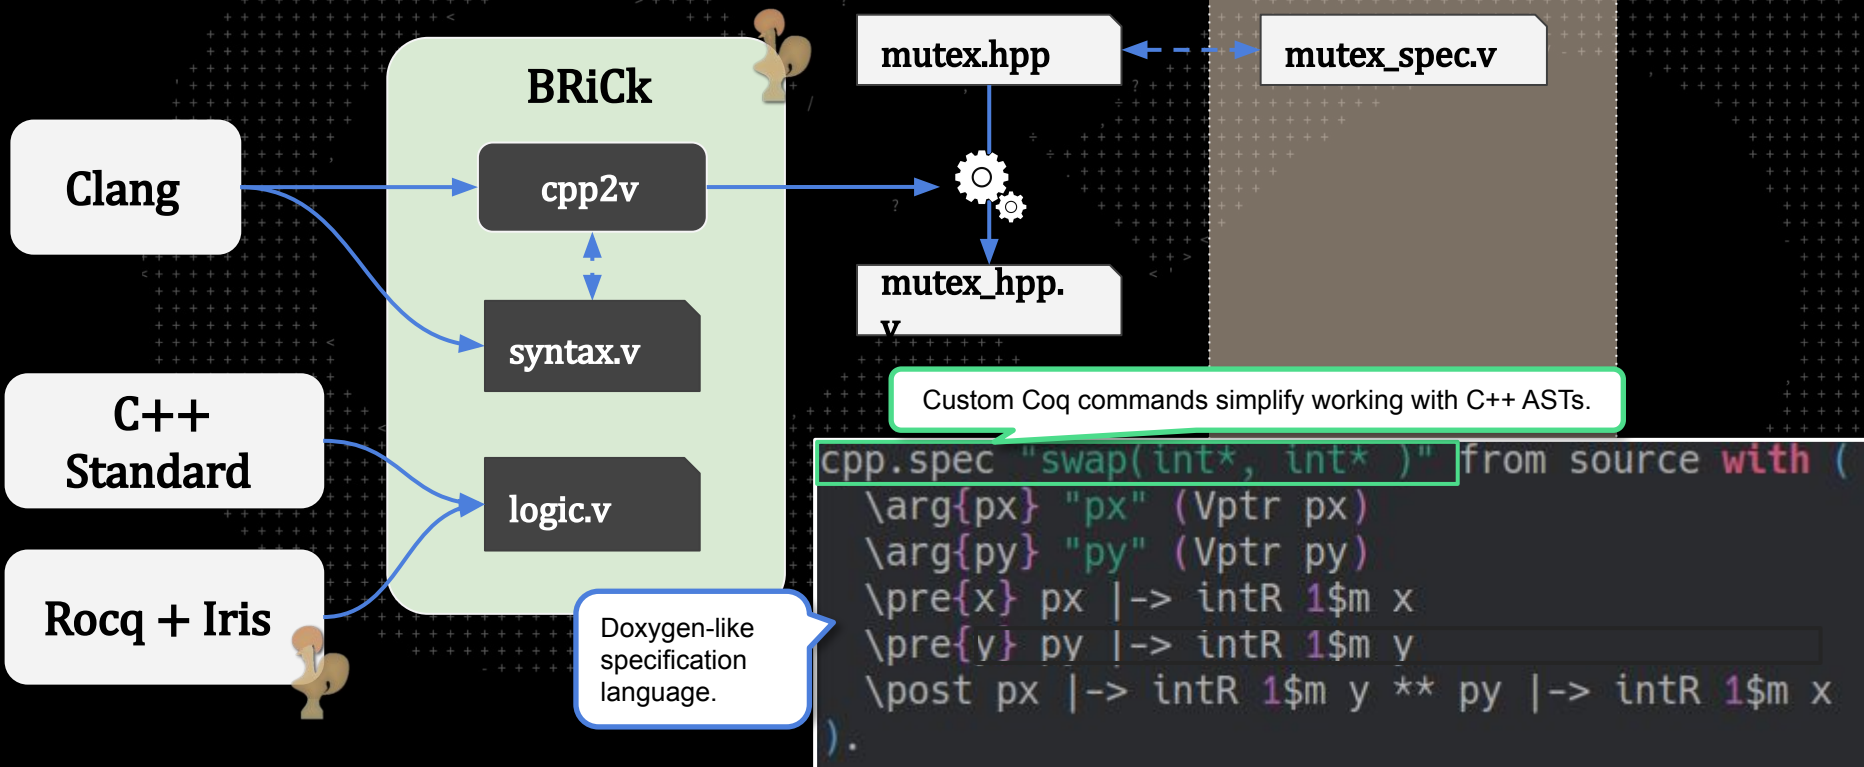
BRiCk C++ Infrastructure

<https://github.com/skylabsai/brick>



BRiCk C++ Infrastructure

<https://github.com/skylabsai/brick>



Doxygen-like specification language.

Custom Coq commands simplify working with C++ ASTs.

```
cpp.spec "swap(int*, int* )" from source with (
  \arg{px} "px" (Vptr px)
  \arg{py} "py" (Vptr py)
  \pre{x} px |-> intR 1$m x
  \pre{y} py |-> intR 1$m y
  \post px |-> intR 1$m y ** py |-> intR 1$m x
  )
.)
```

Mutex Specifications

Specification captures the ownership transfer.

- Mutex protects an invariant.
- `enter()` transfers resources to the caller.
- `exit()` transfers resources to the mutex.

```
cpp.spec "Sys::Mutex::enter()" as enter
  with (λ (this : ptr)
    \nova pd ec
    \persist{γ P} mutex.I (Some γ) P
    \prepost{q} this |-> mutex.R pd (Some γ) q
    \pre mutex.contender γ 1
    \post[Vbool true] P ** mutex.locked γ).

cpp.spec "Sys::Mutex::exit()" as exit
  with (λ (this : ptr) =>
    \nova pd ec
    \persist{γ P} mutex.I (Some γ) P
    \prepost{q} this |-> mutex.R pd (Some γ) q
    \pre |> P ** mutex.locked γ
    \post[Vbool true] mutex.contender γ 1).
```

User-defined representation predicates encapsulate implementation details.

Trade contender for lock ownership

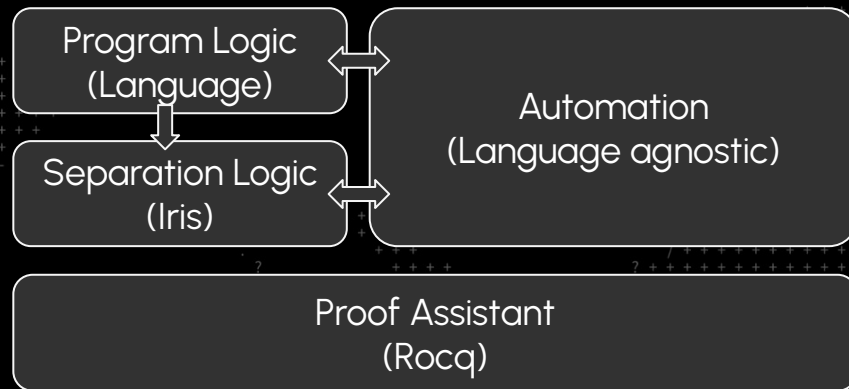
Return lock resources to get contender

Unrestricted use of Iris features for reasoning.

A Foundation for Software Verification

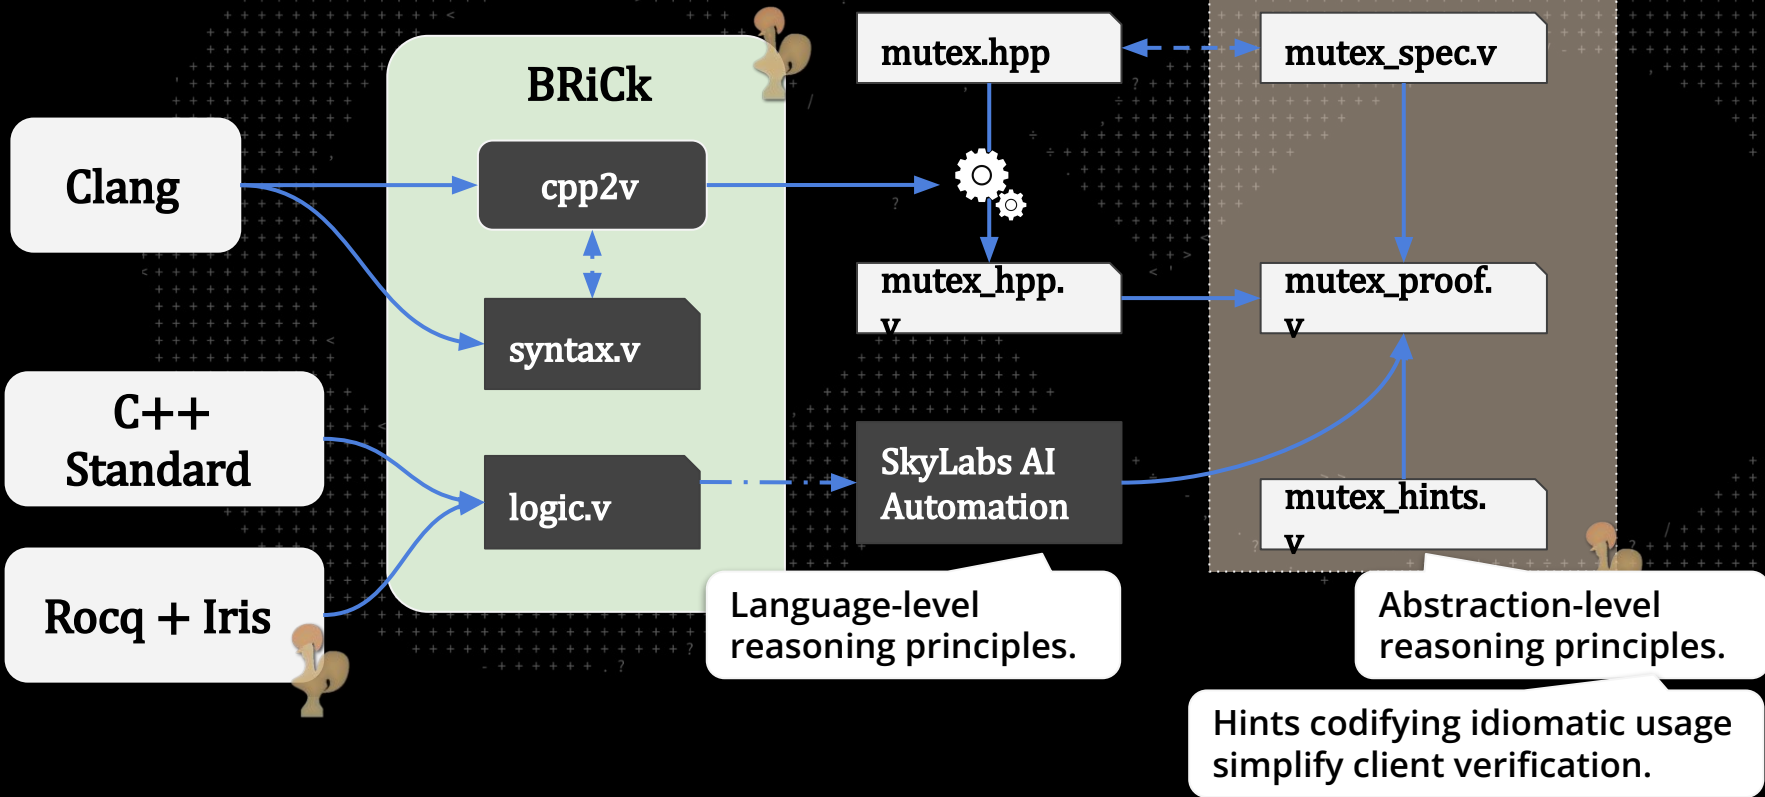
Automation is crucial for scale

- Iris ships with good proof support via Iris Proof Mode (IPM).
 - Excellent for doing sophisticated low-level reasoning
- **Powerful "follow your nose" automation** to scale to more code.
 - Handles the "administrative stuff" so we can focus on the meat of the proof.
 - Instantiating existentials, blasting through WPs, etc.



BRiCk C++ Infrastructure

<https://github.com/skylabsai/brick>



Verification with BRiCk

Language-level Reasoning

Hints capture reasoning principles for the automation to **apply automatically**.

- "Forward" hints work on premises
- "Backward" hints work in the conclusion
- "Cancellation" hints connect premises and conclusions.
- "Learnable" hints learn pure facts from conjunctions.

```
Axiom wp_lval_deref : ∀ ty e Q,  
  letI* v , free := wp_operand e in  
  ∃ p, [| v = Vptr p |] **  
    reference_to (erase_qualifiers ty) p **  
    Q p free  
  ⊢ wp_lval (Ederef e ty) Q.
```

"One-liners"
convert lemmas
into a form usable
by the core
automation.

```
Definition wp_operand_addr_of_B := [BWD] @wp_operand_addr_of.
```

Extend the automation with
BRiCk's language-level
reasoning principles.

SkyLabs AI
Automation

Verification with SL Automation

Reusable hints and hint infrastructure tames the verbosity of low-level proofs.

- Reduce overhead for “follow your nose” proofs.
- Make it tractable for non-experts to **maintain** & (sometimes) **develop** complex proofs.

600 tokens

```
Lemma create_pd_raw2safe_ipm : spec_entails raw.create_pd_spec safe.create_pd_spec.
Proof.
  iIntros (vs K) * /=".
  iDestruct 1 as
    (y_caller_ec y_caller_obj y_caller_hst op sel_target sel_owner_pd y_owner_pd pd_rights q_pd ty) "R".
  iDestruct "R" as (→) "(ECobj & Echst & HpdRights & CapOwn & CapTgt & %tys & SP & Tail)".
  iExists _, _, (λ s, K (status.to_V s)), _, _, _; iFrame "ECobj Echst".
  iSplit; first done.
  iSplit; last by iIntros (?) "S".
  iApply (resolve_sel_rights_intro_weak_sequential HpdRights with "CapOwn [-]"); iIntros "CapOwn /=".
  rewrite tys.
  rewrite /create_pd.create_obj /create_pd.create_hst
    /create_pd.create_first_space /create_pd.create_pio /=.
  destruct ty.
  1, 4, 5 : (← PD; GST; DMA ←)
  iApply (create_sel_intro_weak_sequential with "CapTgt");
  repeat iSplit; iIntros; iApply "Tail" ← /; eauto with iFrame.
  1, 2 : (← OBJ; HST ←)
  iApply (do_reserve_spc_success_intro_weak_sequential with "SPM");
  iIntros "SPM";
  iApply (create_sel_cap_intro_weak_sequential with "CapTgt");
  (repeat iSplit; [..] iIntros (y_spc)]; iIntros "Cap";
  iApply (do_reserve_spc_commit_or_cancel_intro_weak_sequential with "SPM");
  iIntros "SPM"; iApply "Tail" ← /; eauto with iFrame.
  (← PIO ←)
  - iDestruct (x86_only_elim with "SPM") as (IS_Xgg y_pio y_hst) "[SP SH]".
  iApply (do_reserve_spc_failure_intro_weak_sequential with "SPM").
  iIntros "SPM".
  iApply x86_only_intro; [done].
  iApply (do_read_pd_hst_intro_weak_sequential with "SH"). iIntros "SH".
  iApply (create_sel_cap_intro_weak_sequential with "CapTgt").
  repeat iSplit; [..] iIntros (y_spc)]; iIntros "Cap";
  iApply "Tail" ← /;
  rewrite -(? (x86_only_intro _ _ _ IS_Xgg));
  eauto 10 with iFrame.
  (← MSR ←)
  - iDestruct (x86_only_elim with "SPM") as (IS_Xgg) "_".
  iApply x86_only_intro; [done].
  iApply (create_sel_intro_weak_sequential with "CapTgt");
  repeat iSplit; iIntros; iApply "Tail" ← /;
  rewrite -(? (x86_only_intro _ _ _ IS_Xgg));
  eauto with iFrame.
Qed.
```



Several reusable hints later...

83 tokens

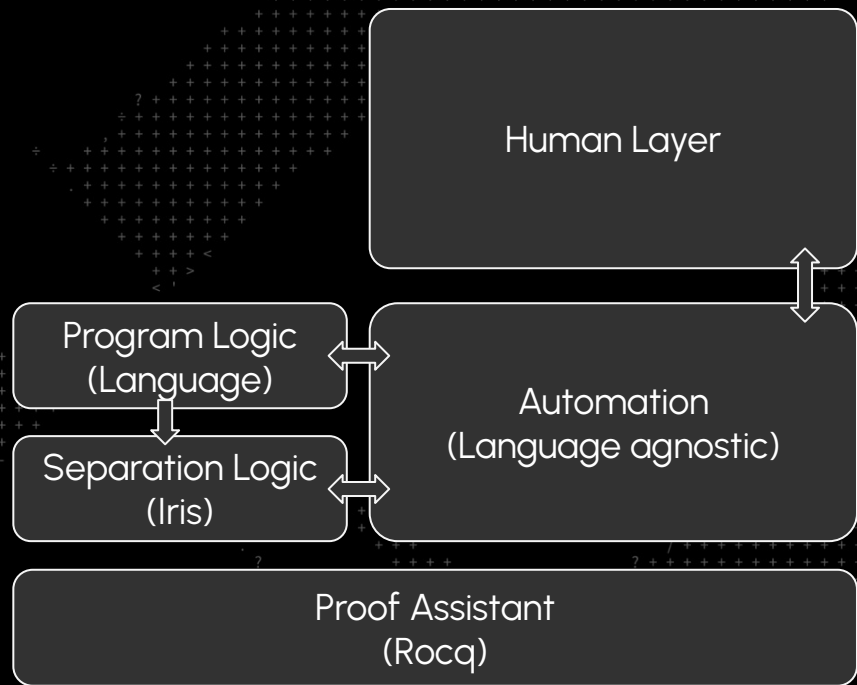
```
Lemma create_pd_raw2safe : spec_entails raw.create_pd_spec safe.create_pd_spec.
Proof.
  iIntros (vs K) "/=" .
  rewrite /create_pd.spec /create_pd.success
    /create_pd.create_obj /create_pd.create_hst
    /create_pd.create_first_space /create_pd.create_pio /=.
  work.
  ren_hyp ty spc.type.t.
  ren_hyp tyS (spc.type.of_N _ = Some ty).
  rewrite tyS.
  case_match; work.
  all: eauto.
Qed.
```



A Foundation for Software Verification

Humans drive the verification by:

- Providing insights when automation gets stuck, e.g. tactics
- Extending automation to handle repetitive reasoning.



Manual Verification w/ Separation Logic

```
void swap(int* px, int* py) {  
  int t = *px;  
  *px = *py;  
  *py = t;  
}
```

```
cpp.spec "swap(int*, int*)" from s  
  \arg{px} "px" (Vptr px)  
  \arg{py} "py" (Vptr py)  
  \pre{x} px |-> intrR 1$m x  
  \pre{y} py |-> intrR 1$m y  
  \post px |-> intrR 1$m y ** py |->  
  ).
```

```
- : PostCond  
- : px |-> intrR 1$m x  
- : py |-> intrR 1$m y  
- : px_addr |-> ptrR<"int"> 1$m px  
- : py_addr |-> ptrR<"int"> 1$m py  
-----*
```

```
- : PostCond  
- : py |-> intrR 1$m y  
- : py_addr |-> ptrR<"int"> 1$m py  
- : px_addr |-> ptrR<"int"> 1$m px  
- : px |-> intrR 1$m x  
- : t_addr |-> intrR 1$m x  
-----*
```

::wpS

```
- : PostCond  
- : t_addr |-> intrR 1$m x  
- : py_addr |-> ptrR<"int"> 1$m py  
- : py |-> intrR 1$m y  
- : px_addr |-> ptrR<"int"> 1$m px  
- : px |-> intrR 1$m y  
-----*
```

::wpS

```
[region: "t" @ t_addr; "py" @ py_addr; "px" @ px_addr; re  
{s: {?: Eassign {e: *$"py"} {e: $"t"} "int"};}
```

Algorithmic

Algorithmics are enough to solve many problems.

- Most code is not very complex, but it is big.

Features such as:

- Straight-line code
- First-order function calls, even indirect
- Reconciling abstractions
- Many array access patterns

```
void swap(int* px, int* py) {  
  int t = *px;  
  *px = *py;  
  *py = t;  
}
```

```
cpp.spec "swap(int*, int*)" from source with (  
  \arg{px} "px" (Vptr px)  
  \arg{py} "py" (Vptr py)  
  \pre{x} px |-> intR 1$m x  
  \pre{y} py |-> intR 1$m y  
  \post px |-> intR 1$m y ** py |-> intR 1$m x  
).
```

```
Lemma swap_ok : verify[source] "swap(int*, int*)".  
Proof. verify_spec. go. Qed.
```

The "bread and butter" of
business logic.

Algorithmic

Purely algorithmic verification is not complete. 🤖

Classically difficult problems

- Loop invariant inference
- Concurrent invariants

SkyLabs' automation is designed to be conservative & predictable.

When we need smarts and intuition, we move up the stack.

```
int loop() {
  int i = 0;
  while (i < 10) {
    i++;
  }
  return i;
}
```

```
cpp.spec "loop()" from source with (
  \post[Vint 10] emp
).
```

```
Lemma loop_ok : verify[source] "loop()".
Proof using MOD.
  verify_spec; go.
(* specify the loop invariant *)
```

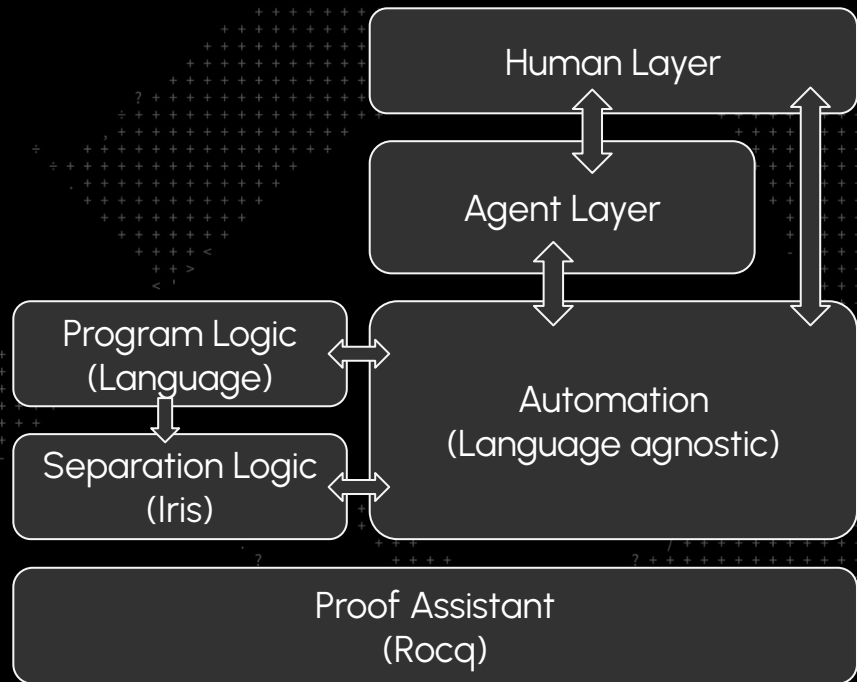
```
- : PostCond
- : i_addr |-> intR 1$m 0
-----*
::wpS
[region: "i" @ i_addr; return {?: "int"}]
{s: while (($"i" < #10)) {
  $"i"++;
  // end block
}}
```

The AI Layer

Human ↔ AI collaboration

- Tight loop integration, e.g. co-pilot.
- Shifting towards agentic delegation.

Need hardened tools for
Agent ↔ Rocq interaction

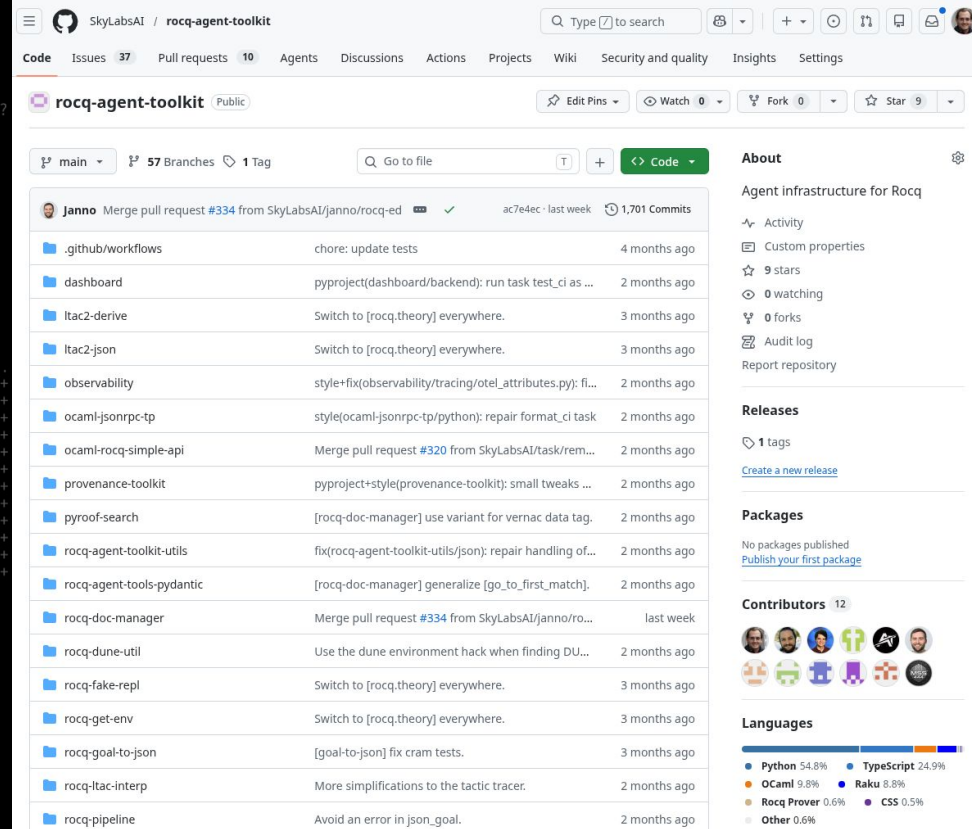


Aside: Rocq Agent Toolkit

Provides tools and libraries for agents to interact with Rocq.

- Tools for inspecting the Rocq environment
- `rocq-ed` – command line interface for Rocq document interaction
- `rocq-doc-manager` – JsonRPC2 interface to Rocq w/ python bindings
- `pyproof-search` – algorithmic proof search built around "strategies" (e.g. Beam, MTS, etc)
- `rocq-agent-tools-pydantic` – Rocq tools for pydantic-ai
- `rocq-pipeline` – infrastructure for *narrow* agent evaluation, e.g. verify this lemma

<https://github.com/skylabsai/rocq-agent-toolkit>



The screenshot shows the GitHub repository page for `rocq-agent-toolkit` by SkylabsAI. The repository is public and has 37 issues, 10 pull requests, 0 agents, 0 discussions, 0 actions, 0 projects, 0 wikis, 0 security and quality issues, 0 insights, and 0 settings. It has 0 edit pins, 0 watch, 0 forks, and 9 stars. The repository is on the `main` branch and has 57 branches and 1 tag. The repository is owned by `Janno` and has 1,701 commits. The repository contains the following files and folders:

File/Folder	Description	Last Commit
<code>.github/workflows</code>	chore: update tests	4 months ago
<code>dashboard</code>	pyproject(dashboard/backend): run task test_ci as ...	2 months ago
<code>ltac2-derive</code>	Switch to [rocq.theory] everywhere.	3 months ago
<code>ltac2-json</code>	Switch to [rocq.theory] everywhere.	3 months ago
<code>observability</code>	style+fix(observability/tracing/otel_attributes.py): fi...	2 months ago
<code>ocaml-jsonrpc-tp</code>	style(ocaml-jsonrpc-tp/python): repair format_ci task	2 months ago
<code>ocaml-rocq-simple-api</code>	Merge pull request #320 from SkylabsAI/task/rem...	2 months ago
<code>provenance-toolkit</code>	pyproject+style(provenance-toolkit): small tweaks ...	2 months ago
<code>pyproof-search</code>	[rocq-doc-manager] use variant for vernac data tag.	2 months ago
<code>rocq-agent-toolkit-utils</code>	fix(rocq-agent-toolkit-utils/json): repair handling of...	2 months ago
<code>rocq-agent-tools-pydantic</code>	[rocq-doc-manager] generalize [go_to_first_match].	2 months ago
<code>rocq-doc-manager</code>	Merge pull request #334 from SkylabsAI/janno/ro...	last week
<code>rocq-dune-util</code>	Use the dune environment hack when finding DU...	2 months ago
<code>rocq-fake-repl</code>	Switch to [rocq.theory] everywhere.	3 months ago
<code>rocq-get-env</code>	Switch to [rocq.theory] everywhere.	3 months ago
<code>rocq-goal-to-json</code>	[goal-to-json] fix cram tests.	3 months ago
<code>rocq-ltac-interp</code>	More simplifications to the tactic tracer.	2 months ago
<code>rocq-pipeline</code>	Avoid an error in json_goal.	2 months ago

The repository also has a right-hand sidebar with the following sections:

- About**: Agent infrastructure for Rocq
- Activity**: Custom properties, 9 stars, 0 watching, 0 forks, Audit log, Report repository
- Releases**: 1 tags, [Create a new release](#)
- Packages**: No packages published, [Publish your first package](#)
- Contributors**: 12 contributors
- Languages**: Python 54.8%, TypeScript 24.9%, OCaml 9.8%, Raku 8.8%, Rocq Prover 0.6%, CSS 0.5%, Other 0.6%

AI Oracles

Verification algorithms query agents as oracles for the “hard” problems.

- Avoid asking agents the routine questions.
- Specialize agents and prompts to specific problems:
 - different loop shapes (for, while, for-each),
 - reconciling abstractions, etc.
- General “next tactic prediction”

```
int loop() {
  int i = 0;
  while (i < 10) {
    i++;
  }
  return i;
}
```

```
cpp.spec "loop()" from source with (
  \post[Vint 10] emp
).
```

```
Lemma loop_ok : verify[source] "loop()"
Proof using MOD.
  verify_spec; go.
(* specify the loop invariant *)
```

GPT

```
- : PostCond
- : i_addr |-> intR 1$m 0
-----*
::wps
[region: "i" @ i_addr; return {?: "int"}]
{s: while (($"i" < #10)) {
  $"i"++;
  // end block
}}
```

AI Oracles

Verification algorithms query agents as oracles for the “hard” problems.

- Avoid asking agents the routine questions.
- Specialize agents and prompts to specific problems:
 - different loop shapes (for, while, for-each),
 - reconciling abstractions, etc.
- General “next tactic prediction”

Suitable for more problems, but necessary to engineer context for each agent call.

```
int loop() {
  int i = 0;
  while (i < 10) {
    i++;
  }
  return i;
}
```

```
cpp.spec "loop()" from source with (
  \post[Vint 10] emp
).
```

```
Lemma loop_ok : verify[source] "loop()"
Proof using MOD.
  verify_spec; go.
  (* specify the loop invariant *)
  wp_while (λ ρ ⇒
    ∃ i, _local ρ "i" |-> intR 1$m i **
    [| (i ≤ 10)%Z |])%I.
  go.
Qed.
```

GPT

Limitations AI Oracles

```
int test(int n) {  
  // Sums all integers from 1 up to n  
  int sum = 0;  
  for (int i = 1; i <= n; ++i) {  
    sum += i;  
  }  
  return sum;  
}
```

```
Lemma test_ok : verify[source] "test(int)".  
Proof using MOD.  
verify_spec; go.  
wp_for (λ rho =>  
  \pre{i} _local rho "i" |-> intR 1$m i  
  \post*_local rho "i" |-> anyR "int" 1$m  
  \pre{sum} _local rho "sum" |-> intR 1$m sum  
  \require (2 * sum = (i - 1) * i)%Z  
  \post emp); go.
```

```
H : 2 * sum = (i - 1) * i  
H1 : n < i  
-----  
2 * sum = n * (n + 1)
```

Unprovable!
Need to show $i = n$



Narrow Agentic

Crucial for models to learn from their mistakes.

- Subsequent model calls need **access to failed attempts** to refine their actions.

```
int test(int n) {  
  // Sums all integers from 1 up to n  
  int sum = 0;  
  for (int i = 1; i <= n; ++i) {  
    sum += i;  
  }  
  return sum;  
}
```

```
Lemma test_ok : verify[source] "test (int)".  
Proof using MOD.  
verify_spec; go.
```



```
- : PostCond  
- : n_addr |-> intR 1$m n  
- : sum_addr |-> intR 1$m 0  
- : i_addr |-> intR 1$m 1  
-----  
::wpS  
[region:  
  "i" @ i_addr;  
  return {?: "i"  
{s: for (; ($"i"  
  {?: Eassi  
  // end block  
  }  
}]  
wp_for (λ rho =  
  \pre{i}_local rho "i" |-> intR 1$m i  
  \post*_local rho "i" |-> anyR "int" 1$m  
  \pre{sum}_local rho "sum" |-> intR 1$m sum  
  \require (2 * sum = (i - 1) * i)%Z  
  \post emp); go.
```

The failed attempt...

```
H : 2 * sum = (i - 1) * i  
H1 : n < i  
-----  
2 * sum = n * (n + 1)
```

...and how it failed

Context engineering is crucial!

Limitations AI Oracles

Crucial for models to learn from their mistakes.

- Subsequent model calls need **access to failed attempts** to refine their actions.

```
int test(int n) {  
  // Sums all integers from 1 up to n  
  int sum = 0;  
  for (int i = 1; i <= n; ++i) {  
    sum += i;  
  }  
  return sum;  
}
```

```
Lemma test_ok : verify[source] "test(int)".  
Proof using MOD.  
verify_spec; go.  
wp_for (λ rho =>  
  \pre{i} _local rho "i" |-> intR 1$m i  
  \post*_local rho "i" |-> anyR "int" 1$m  
  \pre{sum} _local rho "sum" |-> intR 1$m sum  
  \require (2 * sum = (i - 1) * i)%Z  
  \require 1 ≤ i ≤ n+1  
  \post emp); go.
```

GPT

The refined invariant

Context engineering is crucial!

Narrow Agentic

Put agents in control **for a single proof**

- Constrained and engineered context
- **Longer-lived agent** interacts with Rocq through tools, including backtracking.
 - Tool selection is the agent's responsibility.
- **Enabled by more powerful models**
- Easy to prevent cheating during eval

More tokens, but more flexibility.

Built on rocq-agent-tools-pydantic.

```
int test(int n) {  
  // Sums all integers from 1 up to n  
  int sum = 0;  
  for (int i = 1; i <= n; ++i) {  
    sum += i;  
  }  
  return sum;  
}
```

```
lemma test_ok : verify_spec test test (int) = true :=  
  Proof using [wp_for, local_rho] {  
    wp_for (fun i => test i) (local_rho "i" |-> intR 1 $m i)  
    \pre{sum} sum  
    \require (2 * sum = (i - 1) * i)%Z  
    \post eq := go
```

run: "verify_spec; go"

run: "wp_for (...)"

run: "go"

backtrack:2

run: "wp_for (...)"

run: "go"

qed

GPT



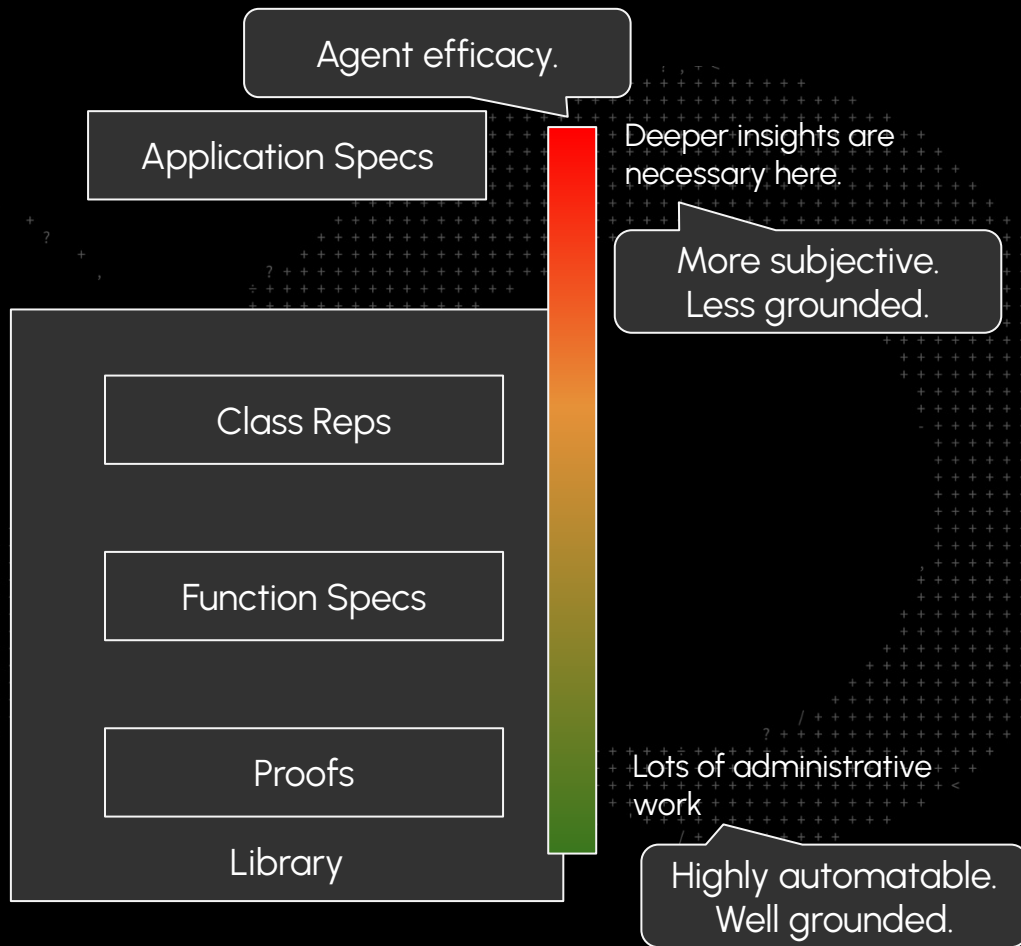
The Future

Spec-driven verification.

- Humans focus on specs & high-level abstractions.
- Agents fill in the proofs.
- **Good specifications signal good abstractions!**

Farther ahead (but not too far?)

- Expanding agent capabilities:
 - Writing specs and representation predicates.



Questions

Authors

Lennart Beringer
Jasper Haag
Rodolphe Lepigre
Gregory Malecha
Rayan Soban
Ehtesham Zahoor

Additional Contributors

Janno Kaiser
Simon Hudon
Paolo G. Giarrusso
Hai Dang (BlueRock Security)
Abhishek Anand (Category Labs)

- Separation Logic as a foundation for verification
 - Language-agnostic
 - Full system (multi-program)
 - Concurrency
 - Refinement
- Tools for humans & agents
 - Rocq: general purpose proof assistant
 - Embedded logic **w/ automation**
- Algorithmic \leftrightarrow Agentic architectures