



# Kani

A Bit-Precise Rust Verifier

Zyad Hassan

May 19, 2022





# Kani

## A Bit-Precise Rust Verifier

Daniel Schwartz-Narbonne, Celina G. Val  
March 28 2022

© 2022, Amazon Web Services, Inc. or its Affiliates.

0:01 / 53:17



March Session -- Kani

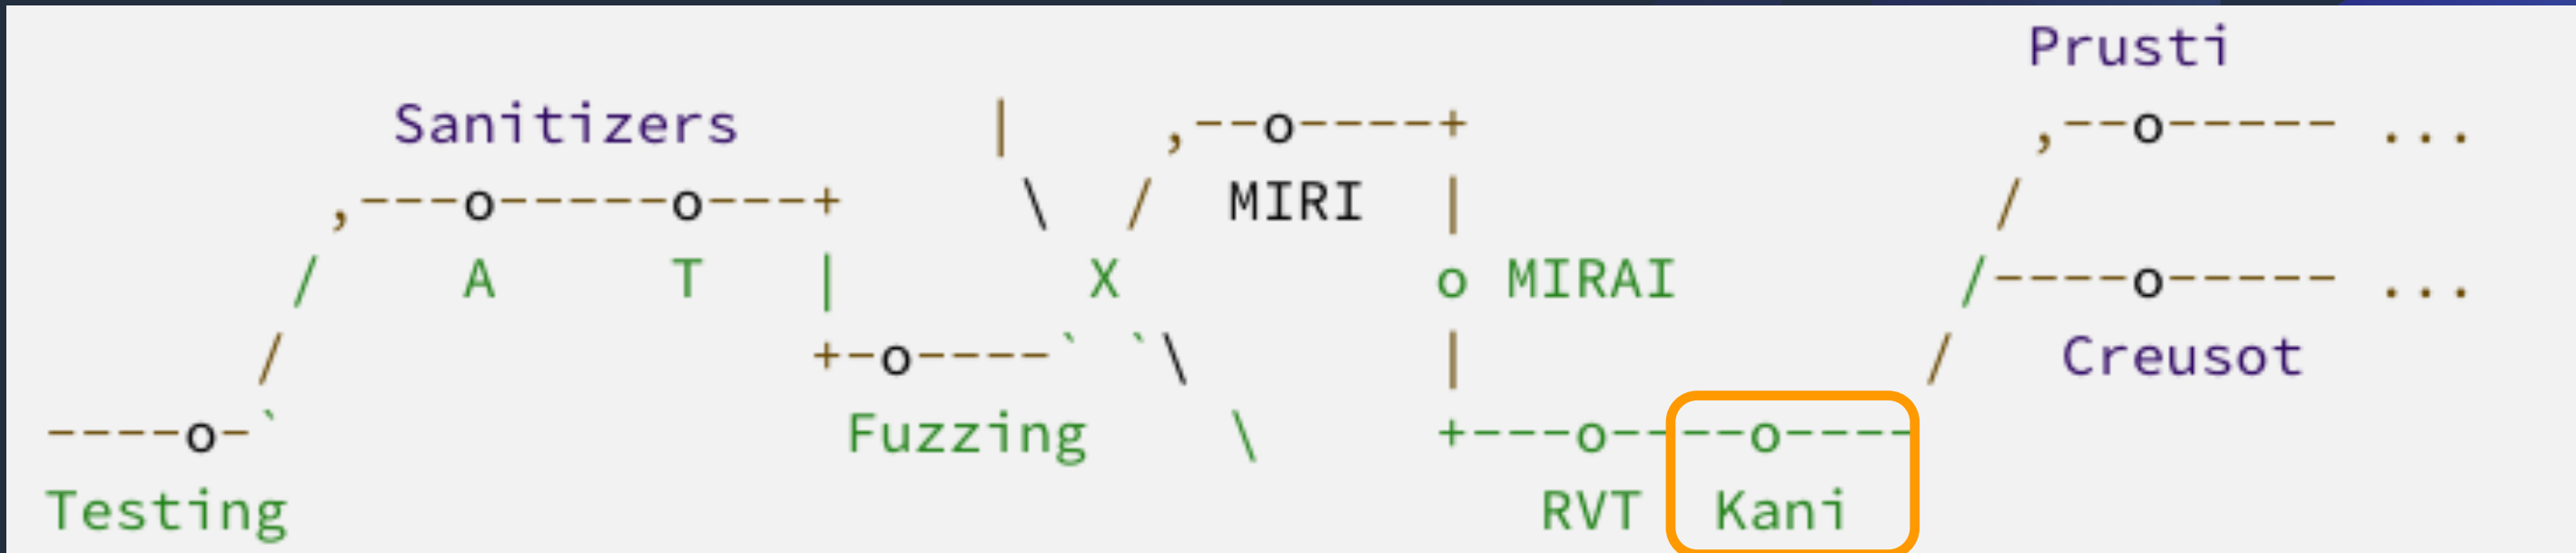
39 views • Mar 28, 2022

👍 0    🗨 DISLIKE    ➦ SHARE    ≡+ SAVE    ...



**Rust Formal Methods IG**  
23 subscribers

SUBSCRIBE



# Demo

# Initial users: Systems Programmers

## Requirements:

1. Bit precision
2. Supporting unsafe code
3. Heap modeling

# Verifying Rust systems code requires bit-precision

```
96  #[cfg_attr(not(test), rustc_diagnostic_item = "OsStr")]
97  #[stable(feature = "rust1", since = "1.0.0")]
98  // FIXME:
99  // `OsStr::from_inner` current implementation relies
100 // on `OsStr` being layout-compatible with `Slice`.
101 // When attribute privacy is implemented, `OsStr` should be annotated as `#[repr(transparent)]`.
102 // Anyway, `OsStr` representation and layout are considered implementation details, are
103 // not documented and must not be relied upon.
104 pub struct OsStr {
105     inner: Slice,
106 }
```

[https://doc.rust-lang.org/stable/src/std/ffi/os\\_str.rs.html#98](https://doc.rust-lang.org/stable/src/std/ffi/os_str.rs.html#98)

# Verifying Rust systems code requires supporting unsafe operations

```
25     // The compiler isn't smart enough to remove all of the bounds checks so we resort to
26     // `get_unchecked`.
27     //
28     // https://godbolt.org/z/45cG1v
29
30     // iterate until we reach one of the ends
31     while from_index < from.len() && to_index < to.len() {
32         let from = unsafe {
33             // Safety: this length is already checked in the while condition
34             debug_assert!(from.len() > from_index);
35             from.get_unchecked(from_index)
36         };
37
38         let to = unsafe {
39             // Safety: this length is already checked in the while condition
40             debug_assert!(to.len() > to_index);
41             to.get_unchecked_mut(to_index)
42         };

```

<https://github.com/aws/s2n-quic/blob/main/quic/s2n-quic-core/src/slice.rs>

# Verifying Rust systems code requires a precise model of the heap

```
51  /// Trait implemented by the underlying `MmapRegion`.
52  pub(crate) trait AsSlice {
53      /// Returns a slice corresponding to the data in the underlying `MmapRegion`.
54      ///
55      /// # Safety
56      ///
57      /// This is unsafe because of possible aliasing.
58      unsafe fn as_slice(&self) -> &[u8];
59
60      /// Returns a mutable slice corresponding to the data in the underlying `MmapRegion`.
61      ///
62      /// # Safety
63      ///
64      /// This is unsafe because of possible aliasing. Accesses done through the resulting slice
65      /// are not visible to dirty bitmap tracking functionality (when present), and have to be
66      /// explicitly accounted for.
67      #[allow(clippy::mut_from_ref)]
68      unsafe fn as_mut_slice(&self) -> &mut [u8];
69  }
```

<https://github.com/rust-vmm/vm-memory/blob/main/src/mmap.rs>



# Properties checked

## Automatic checks

- Buffer overflows
- Pointer safety
- Division by zero
- Pointer arithmetic
- Arithmetic overflows

## User defined properties

- Assertions
- Object invariants

Note that this does not (yet) cover all Rust UB

- Type safety
- Invalid bit patterns
- Aliasing violations

**⚠ Warning:** The following list is not exhaustive. There is no formal model of Rust's semantics for what is and is not allowed in unsafe code, so there may be more behavior considered unsafe. The following list is just what we know for sure is undefined behavior. Please read the [Rustonomicon](#) before writing unsafe code.

# Formal methods in the development workflow



## Code-Level Model Checking in the Software Development Workflow

Nathan Chong  
Amazon

Byron Cook  
Amazon  
UCL

Konstantinos Kallas  
University of Pennsylvania

Kareem Khazem  
Amazon

Felipe R. Monteiro  
Amazon

Daniel Schwartz-Narbonne  
Amazon

Serdar Tasiran  
Amazon

Michael Tautschnig  
Amazon  
Queen Mary University of London

Mark R. Tuttle  
Amazon

### ABSTRACT

This experience report describes a style of applying symbolic model checking developed over the course of four years at Amazon Web Services (AWS). Lessons learned are drawn from proving properties of numerous C-based systems, e.g., custom hypervisors, encryption code, boot loaders, and an IoT operating system. Using our methodology, we find that we can prove the correctness of industrial low-level C-based systems with reasonable effort and predictability. Furthermore, AWS developers are increasingly writing their own formal specifications. All proofs discussed in this paper are publicly available on GitHub.

particular, formal specification of code provides precise, machine-checked documentation for developers and consumers of a code base. They improve code quality by ensuring that the program's *implementation* reflects the developer's *intent*. Unlike testing, which can only validate code against a set of concrete inputs, formal proof can assure that the code is both secure and correct for all possible inputs.

Unfortunately, rapid proof development is difficult in cases where proofs are written by a separate specialized team and *not* the software developers themselves. The developer writing a piece of code has an internal mental model of their code that explains why, and under what conditions, it is correct. However, this model typically remains

<https://ieeexplore.ieee.org/document/9276622>

# Key insight: use artifacts that fit within developers workflow

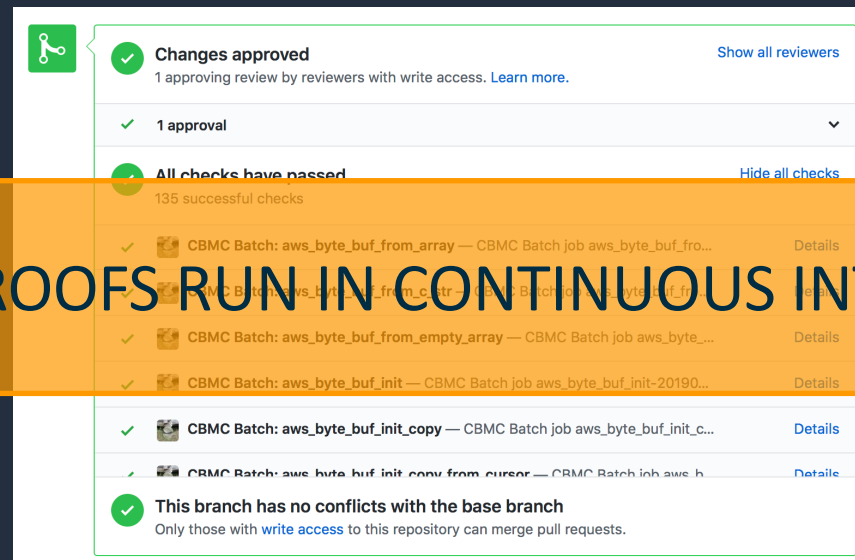
WRITE SPECS IN THEIR LANGUAGE

```
1 /**
2  * Copyright Amazon.com, Inc. or its affiliates. All Rights Reserved.
3  * SPDX-License-Identifier: Apache-2.0.
4  */
5
6 #include <aws/common/array_list.h>
7 #include <proof_helpers/make_common_data_structures.h>
8
9 /**
10 * Runtime: 9s
11 */
12 void aws_array_list_pop_front_harness() {
13     /* data structure */
14     struct aws_array_list list;
15
16     /* assumptions */
17     CPROVER_ASSUME(list.data != NULL);
18     CPROVER_ASSUME(list.alloc == 1);
19     CPROVER_ASSUME(list.current_size == 0);
20     CPROVER_ASSUME(list.item_size == 1);
21     CPROVER_ASSUME(list.length == 0);
22     CPROVER_ASSUME(list.current_index == 0);
23     CPROVER_ASSUME(list.current_index == 0);
24     CPROVER_ASSUME(list.current_index == 0);
25     CPROVER_ASSUME(list.current_index == 0);
26     CPROVER_ASSUME(list.current_index == 0);
27     CPROVER_ASSUME(list.current_index == 0);
28     CPROVER_ASSUME(list.current_index == 0);
29     CPROVER_ASSUME(list.current_index == 0);
30     CPROVER_ASSUME(list.current_index == 0);
31     CPROVER_ASSUME(list.current_index == 0);
32     CPROVER_ASSUME(list.current_index == 0);
33     CPROVER_ASSUME(list.current_index == 0);
34     CPROVER_ASSUME(list.current_index == 0);
35     CPROVER_ASSUME(list.current_index == 0);
36     CPROVER_ASSUME(list.current_index == 0);
37     CPROVER_ASSUME(list.current_index == 0);
38 }
```

EMBED THE SPECS IN THEIR CODE

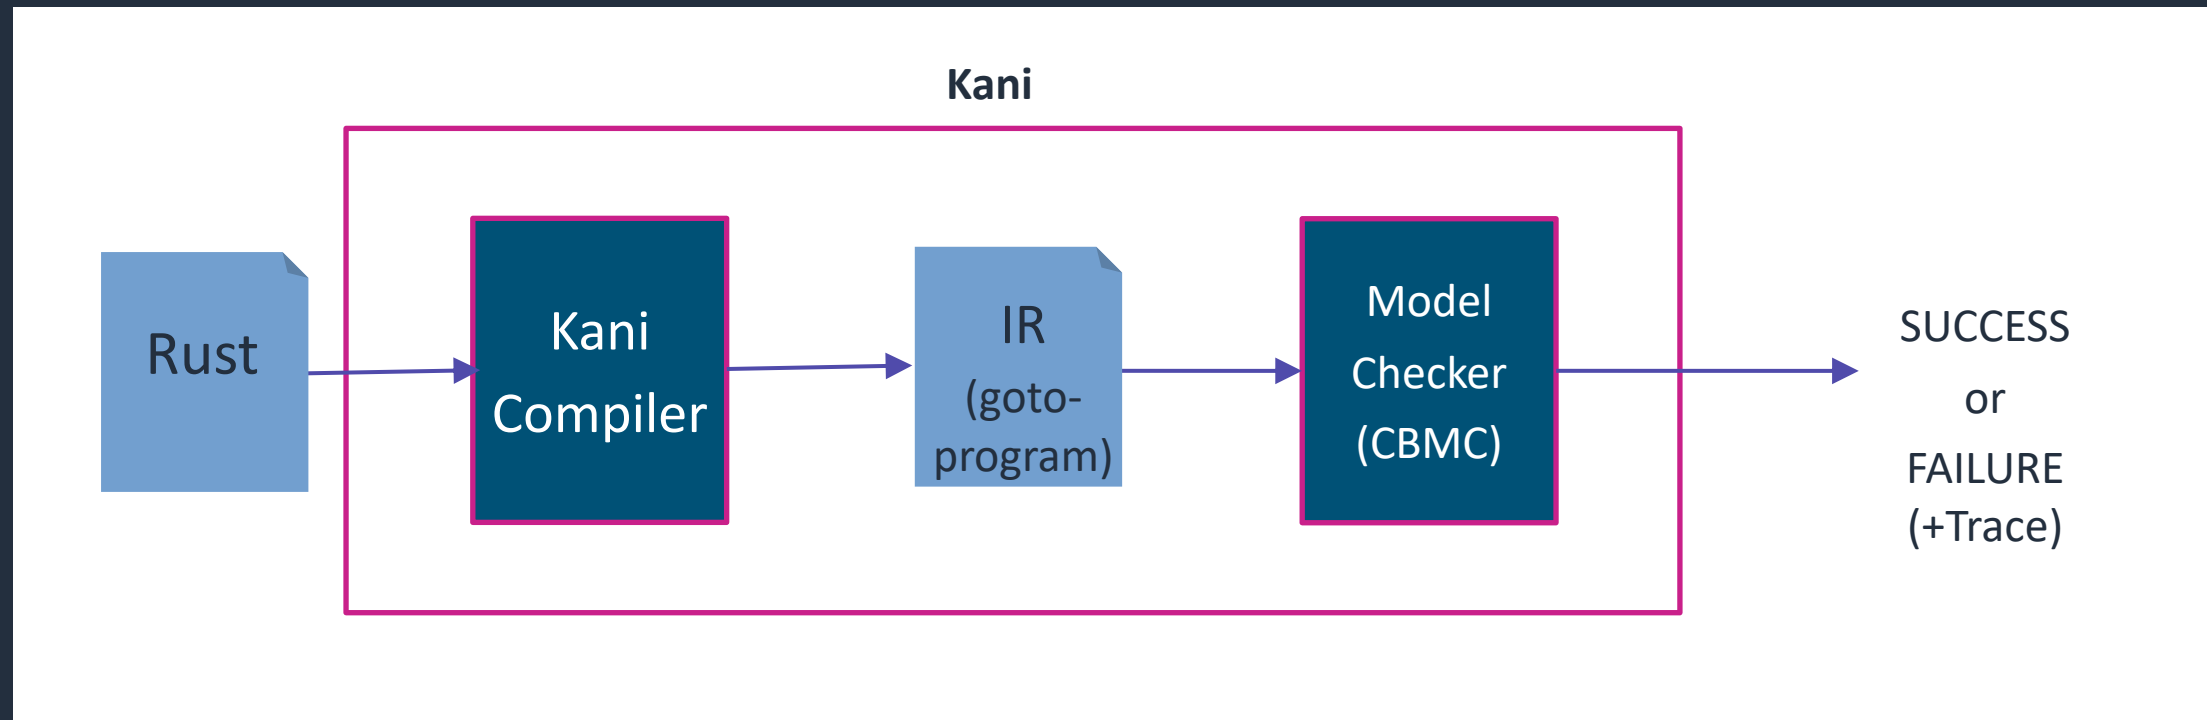
```
164 AWS_STATIC_IMPL
165 int aws_array_list_pop_front(struct aws_array_list *AWS_RESTRICT list) {
166     AWS_PRECONDITION(aws_array_list_is_valid(list));
167     if (aws_array_list_length(list) > 0) {
168         aws_array_list_pop_front_n(list, 1);
169         AWS_POSTCONDITION(aws_array_list_is_valid(list));
170         return AWS_OP_SUCCESS;
171     }
172     AWS_POSTCONDITION(aws_array_list_is_valid(list));
173     return aws_raise_error(AWS_ERROR_LIST_EMPTY);
174 }
175 }
```

PROOFS RUN IN CONTINUOUS INTEGRATION



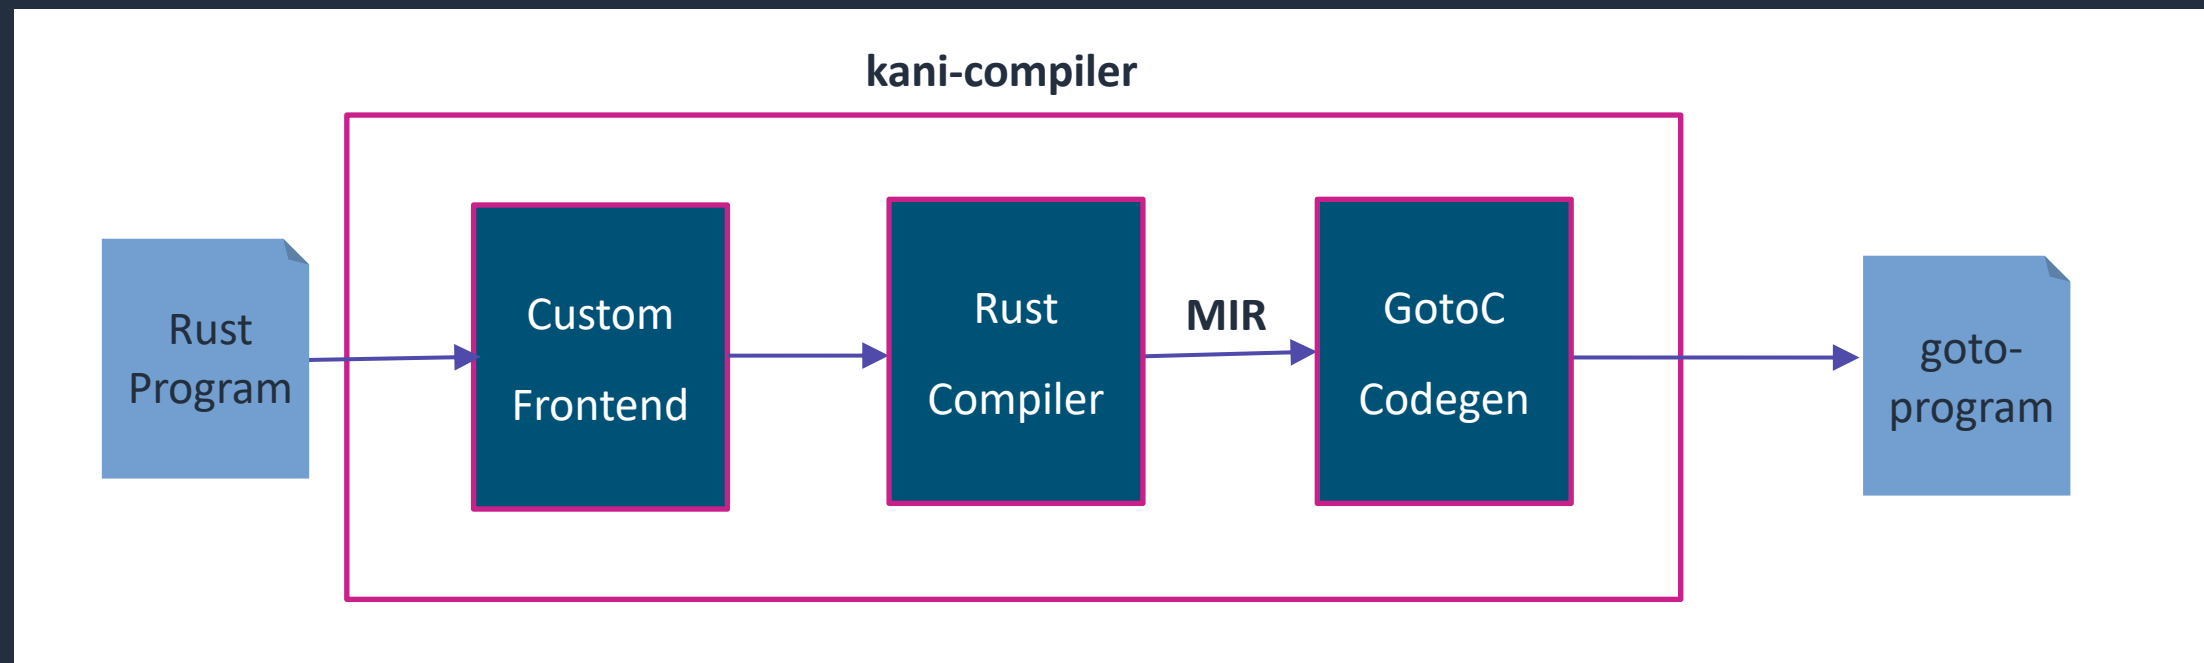
# Under the Hood

# High Level Flow



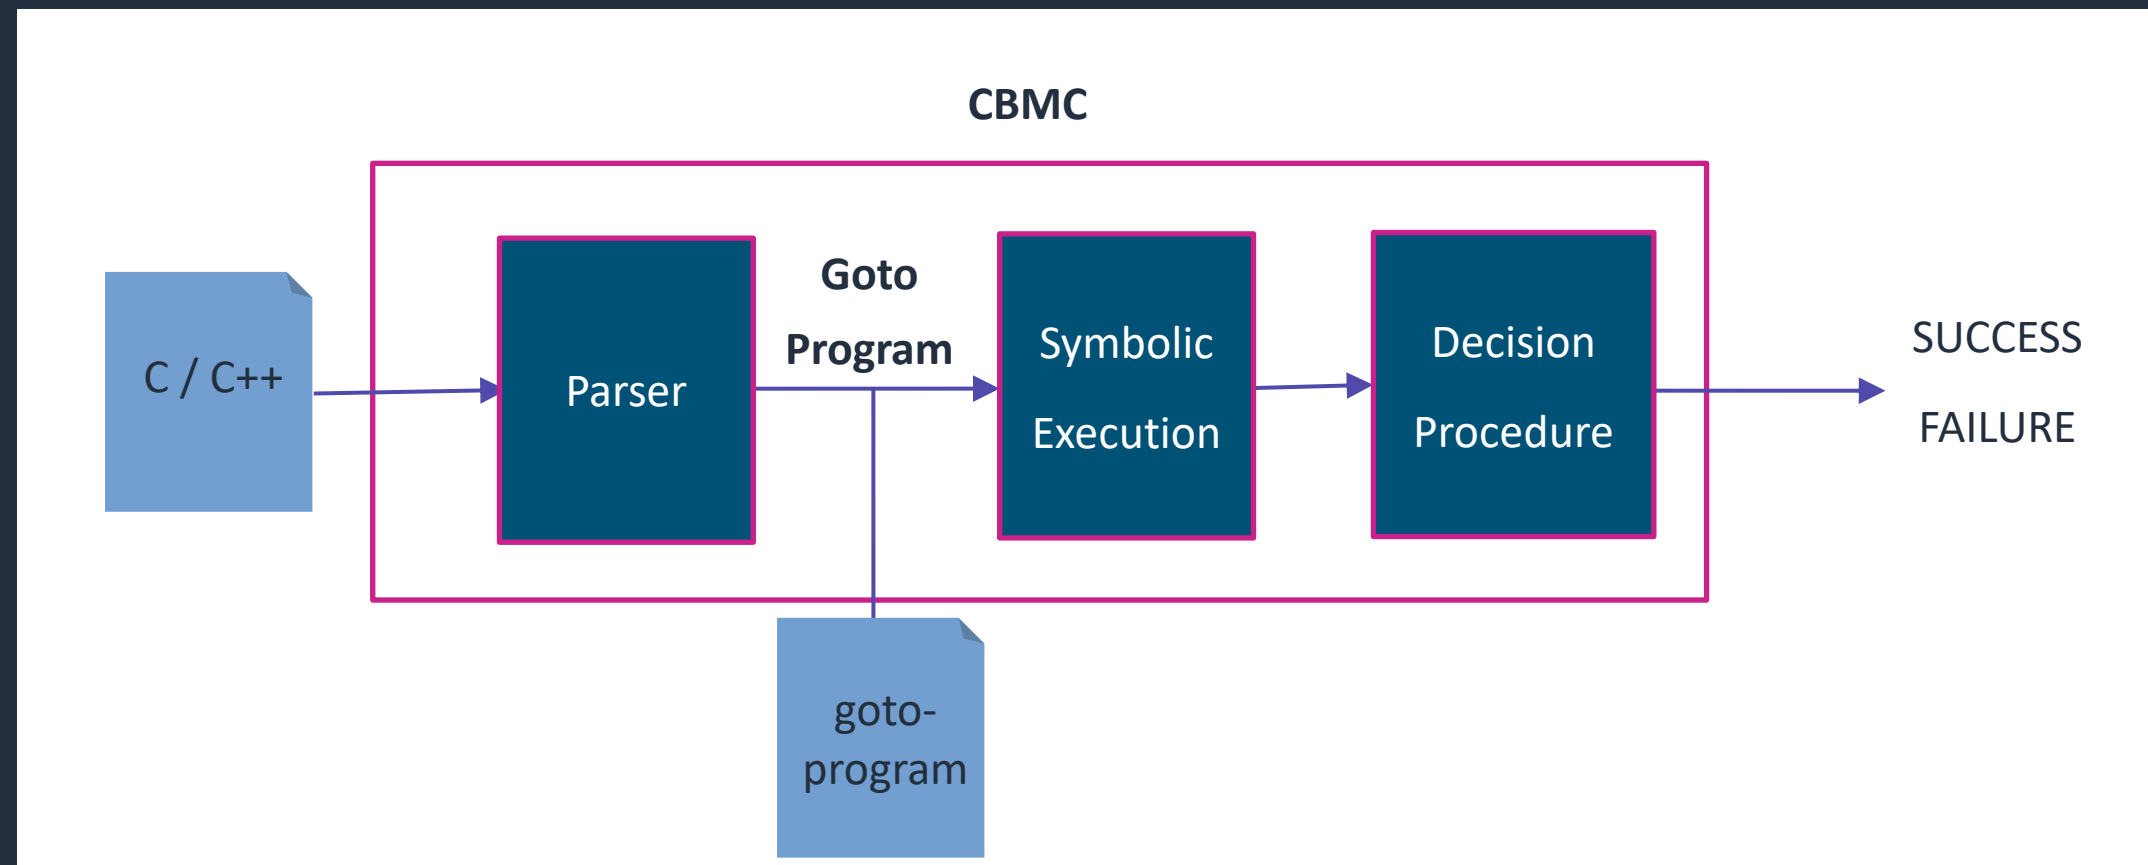
# Kani Compiler

- Adds verification checks
- Includes a custom codegen that translates MIR to goto-program
  - MIR: Mid-Level Intermediate Representation
  - **MIR is the source of truth**



# Model Checker: CBMC

- Bounded Model Checker developed for C and C++ programs.
- Verifies memory safety, undefined behavior, user-specified assertions...
- Uses a MiniSat based solver for bit-vector formulas by default.





# Open Source

The screenshot shows the GitHub repository page for `model-checking/kani`. At the top, there is a search bar and navigation links for Pull requests, Issues, Marketplace, and Explore. The repository name `model-checking / kani` is displayed with a 'Public' badge. Action buttons include 'Edit Pins', 'Unwatch' (9), 'Fork' (27), and 'Starred' (514). Below this is a navigation bar with links for Code, Issues (222), Pull requests (6), Discussions, Actions, Projects (1), Security, Insights, and Settings.

Repository details show the `main` branch selected, with 2 branches and 2 tags. Buttons for 'Go to file', 'Add file', and 'Code' are visible. The 'About' section on the right identifies the project as 'Kani Rust Verifier' and provides the URL `model-checking.github.io/kani`. It also lists tags: `rust`, `verification`, and `model-checking`. Other statistics include 514 stars, 9 watchers, and 27 forks.

The commit history table is as follows:

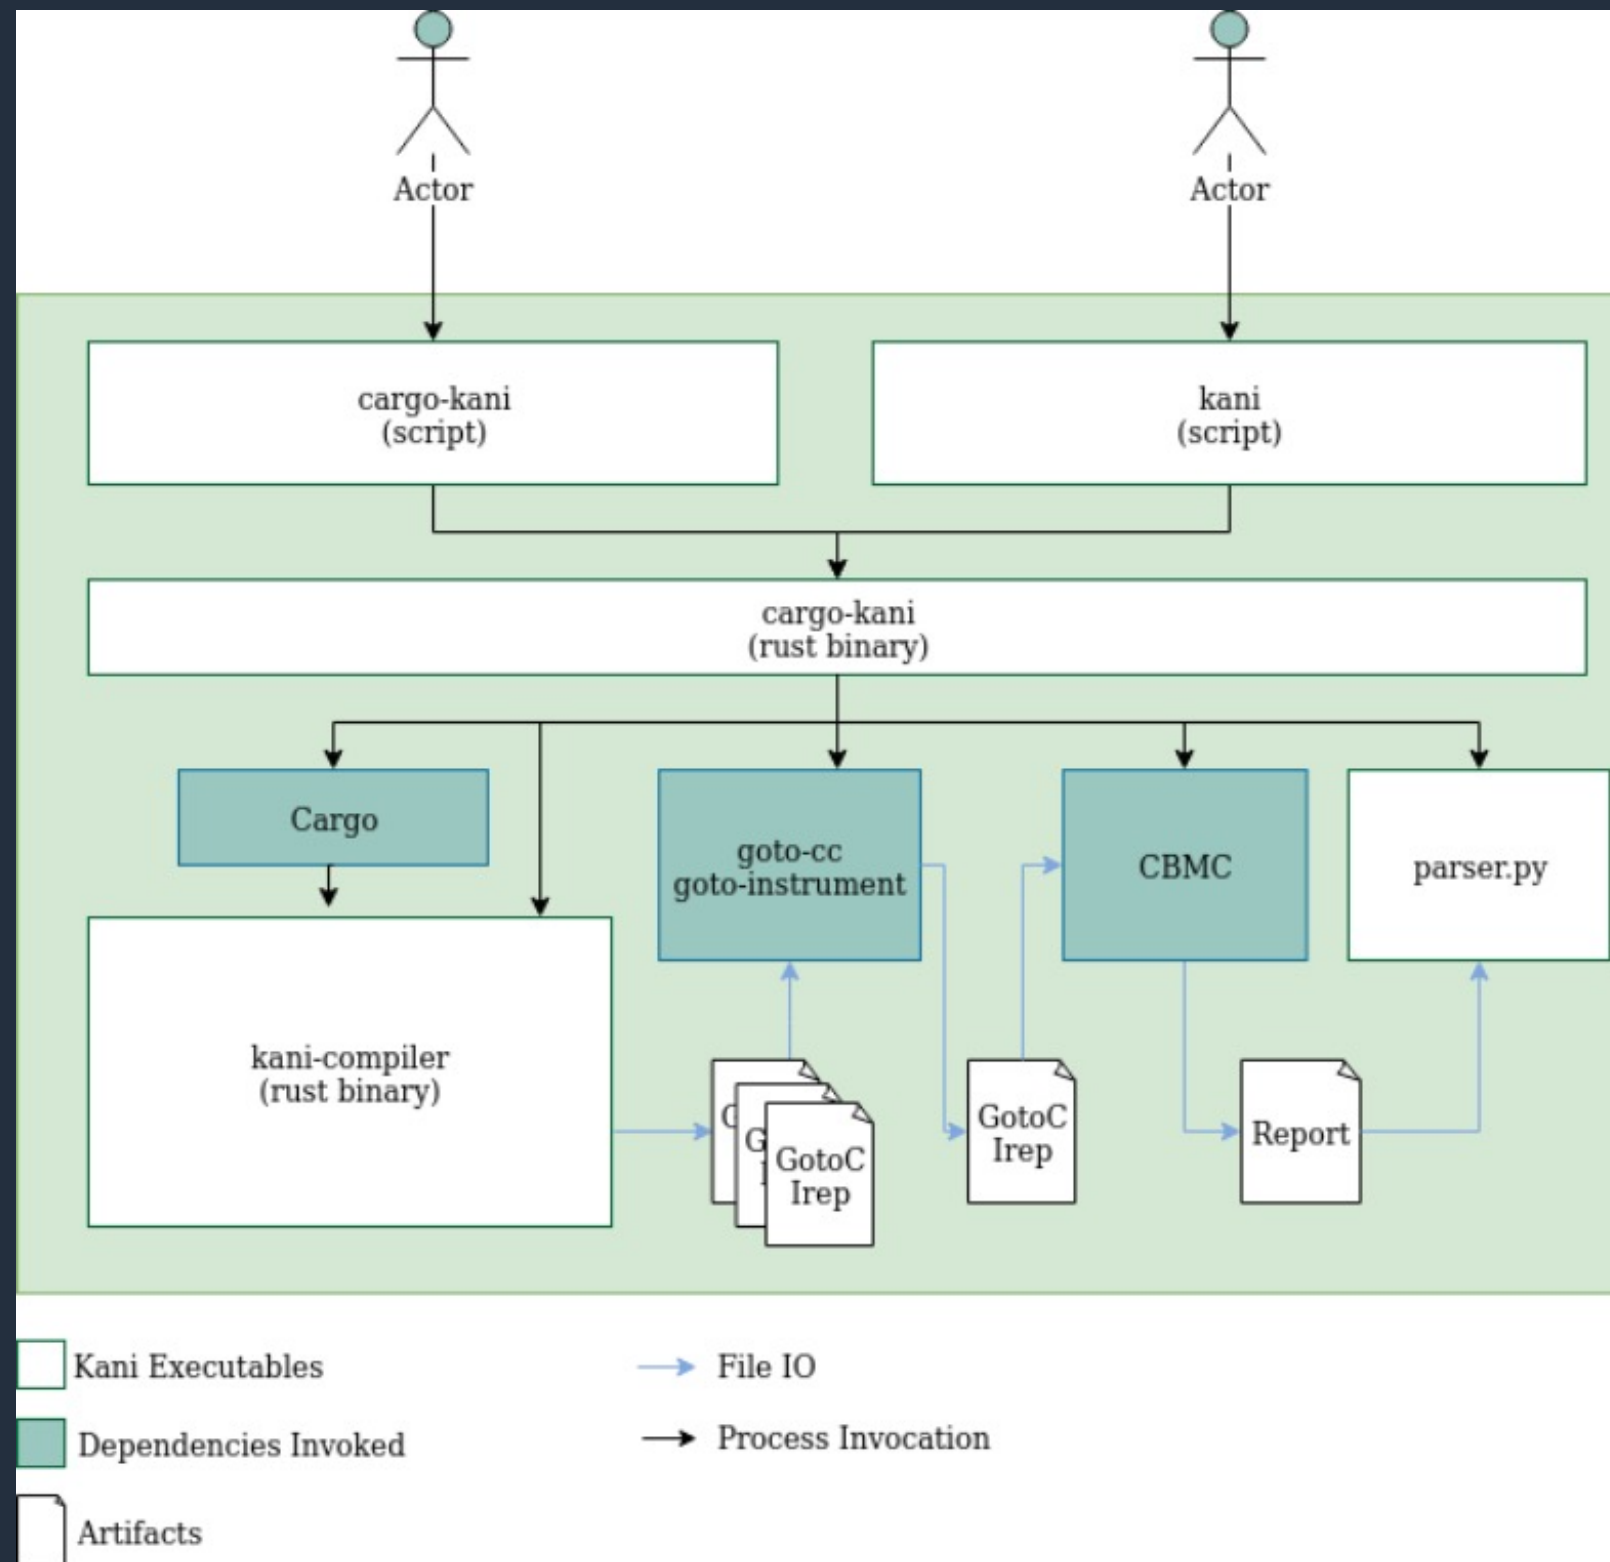
Commit	Description	Time
tedinski Bump version to 0.2.0 (#1204) ✓ c7c0c4f 2 days ago ⌚ 559 commits		
.github	Ensure cargo-kani setup is idempotent (#1193)	7 days ago
cprover_bindings	Update the rust toolchain to <code>nightly-2022-05-03</code> (#1181)	8 days ago
docs	Update python requirement (#1182)	9 days ago
firecracker @ cd36c69	Regression test for codegen'ing Firecracker (#264)	10 months ago
kani-compiler	Bump version to 0.2.0 (#1204)	2 days ago
kani-driver	Bump version to 0.2.0 (#1204)	2 days ago
kani_metadata	Bump version to 0.2.0 (#1204)	2 days ago

# Summary

- Systems programmers need
  - Bit precision
  - Precise modeling of heap
  - Precise modeling of unsafe operations
- Integrate into the developer workflow
  - Unit/Prop test like “proof harnesses”
  - Run using `cargo`
  - Concrete debug traces
- Leverage well established tools
  - Rust Compiler
  - CBMC as bit-precise solver

# Appendix

# Architecture



# Kani Verification: No Error Path

