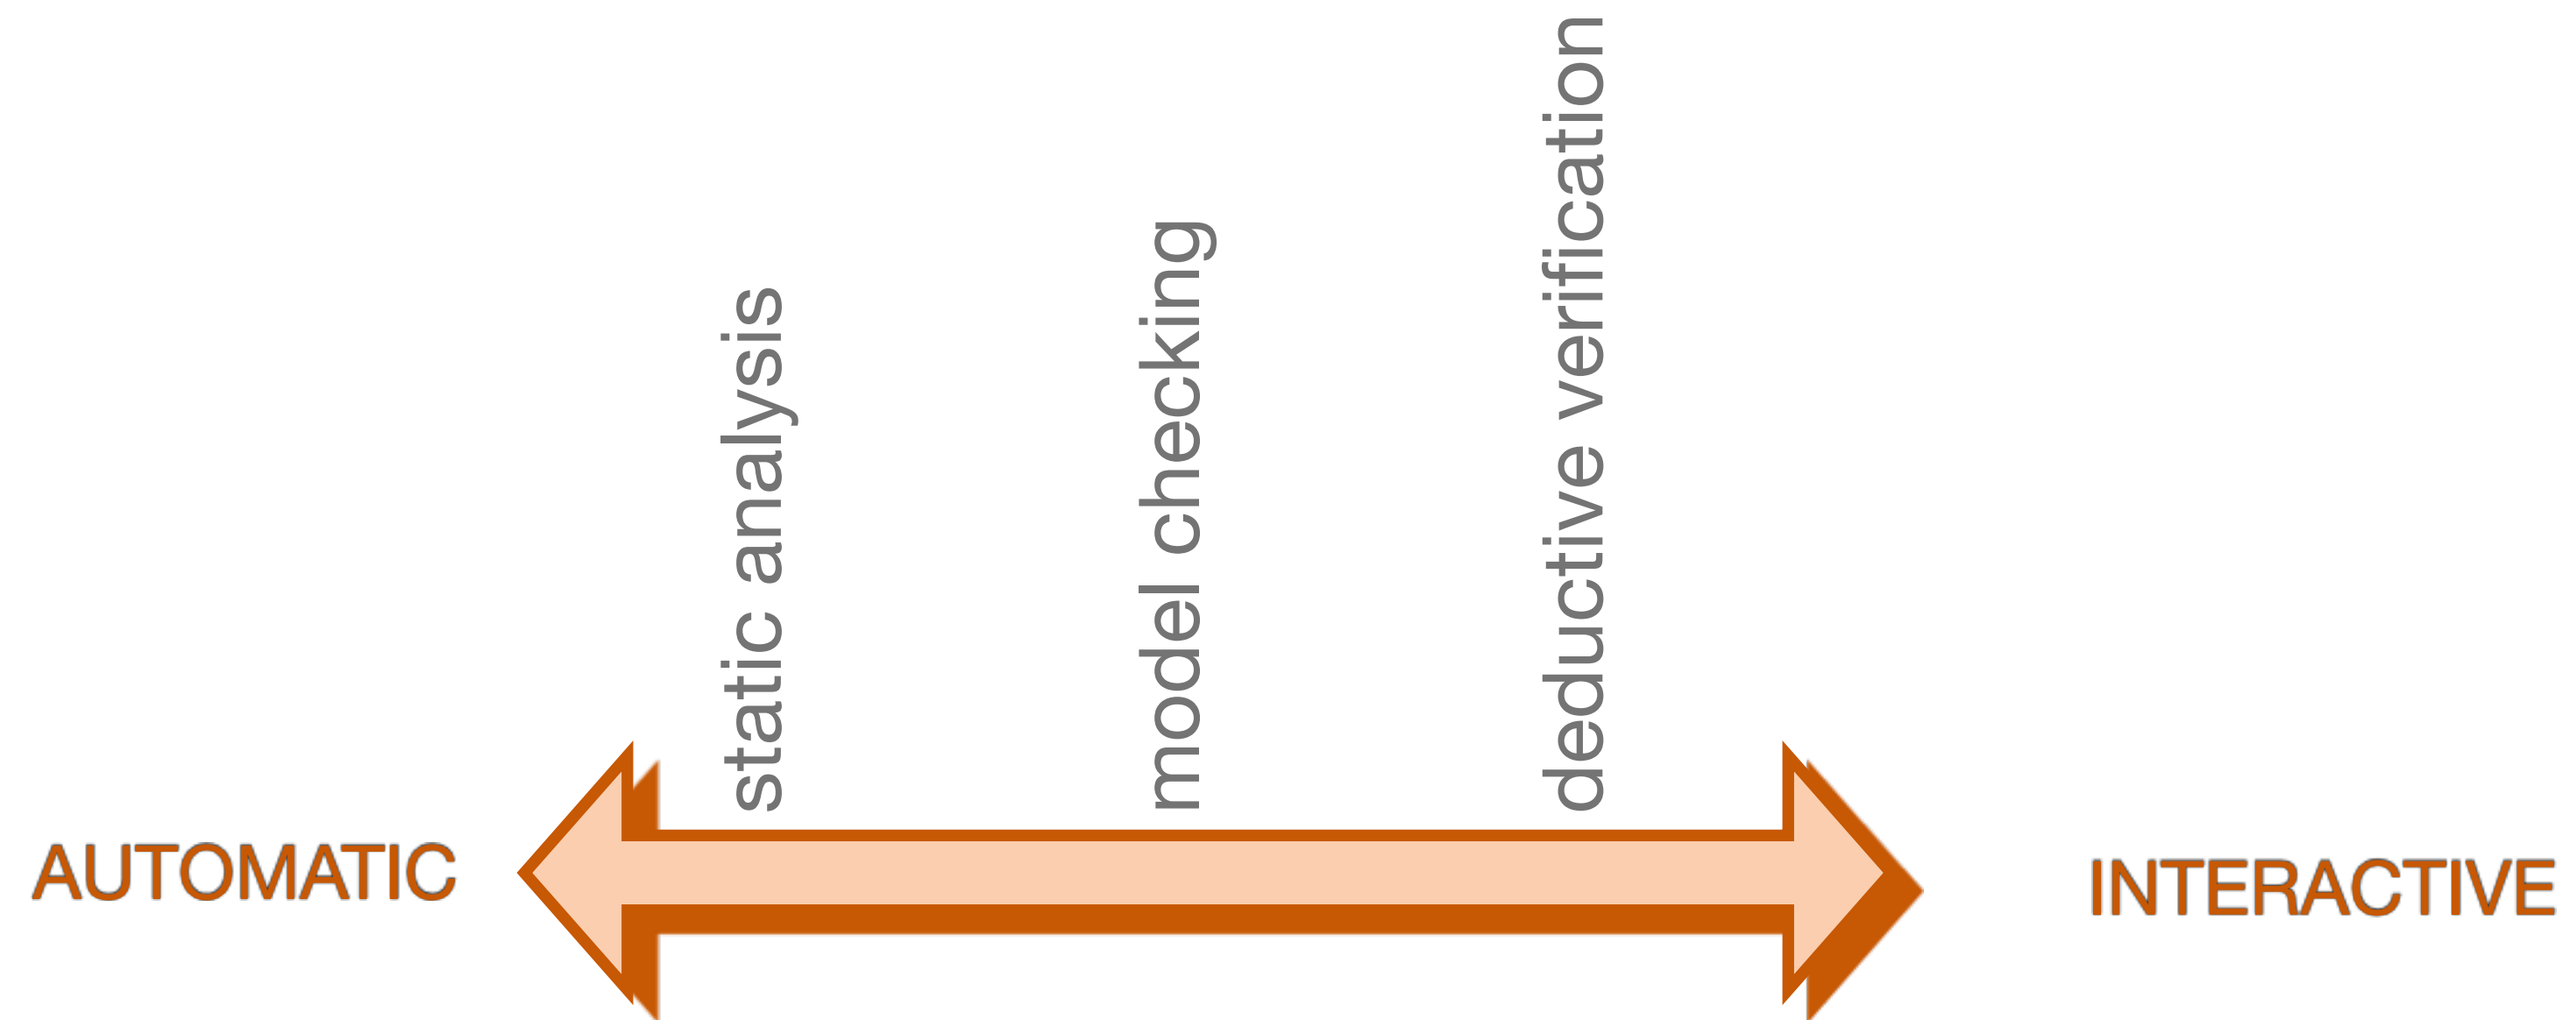


Verified compilation: towards zero-defect software


Sandrine Blazy



Formal verification of software: tool-assisted techniques

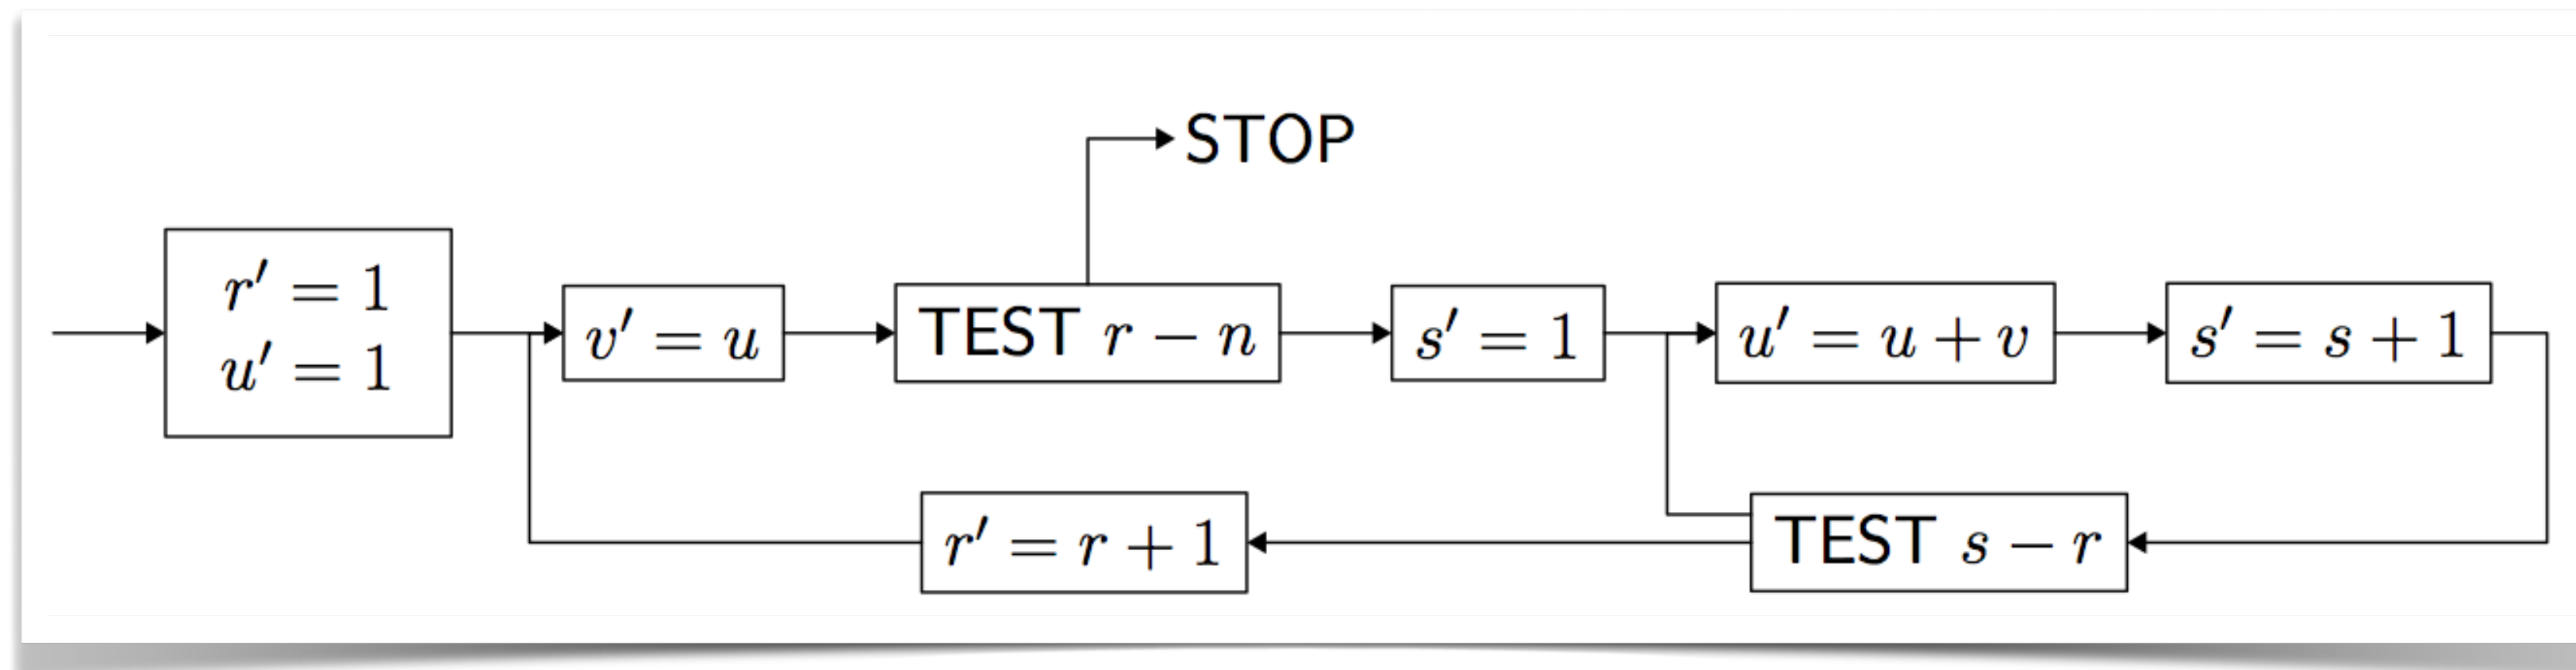
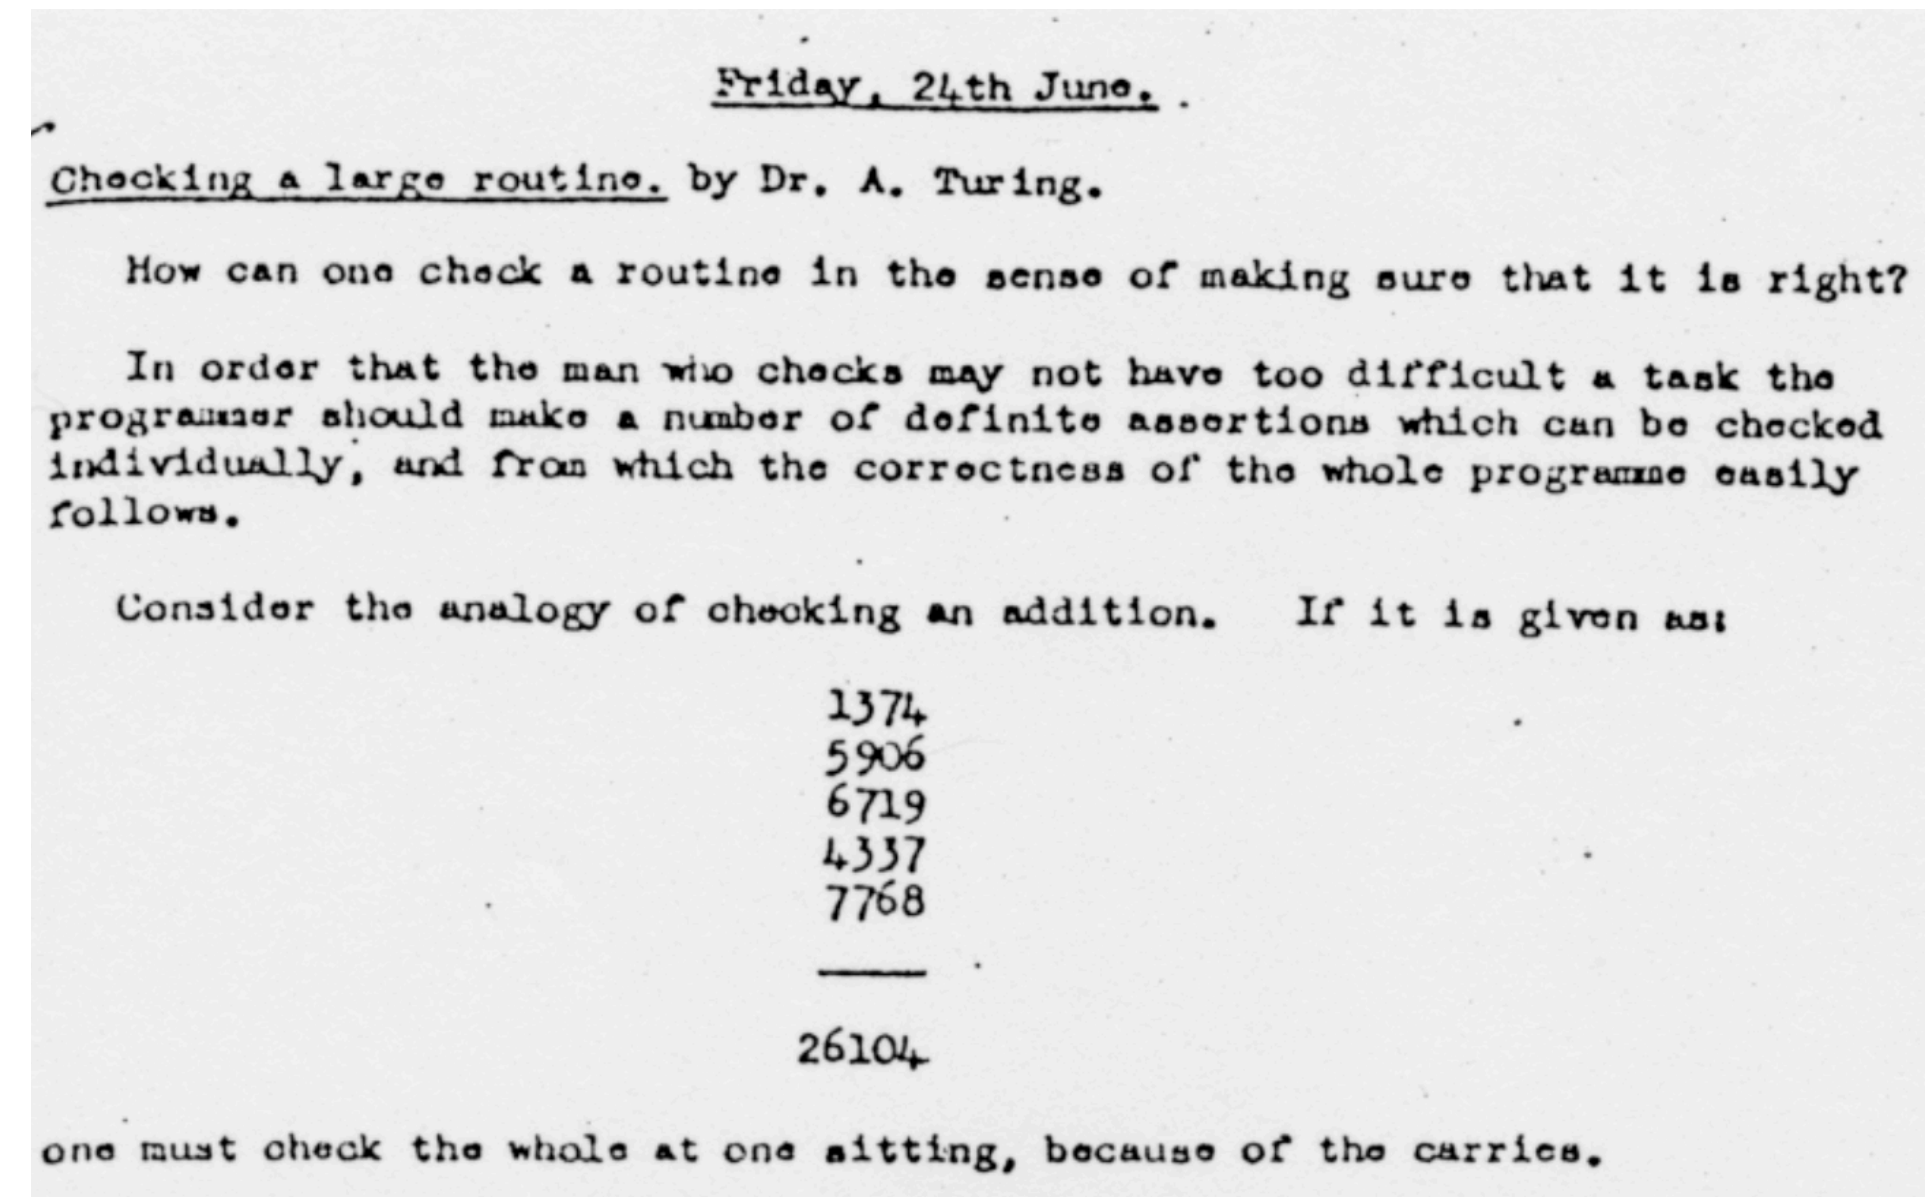


Deductive verification



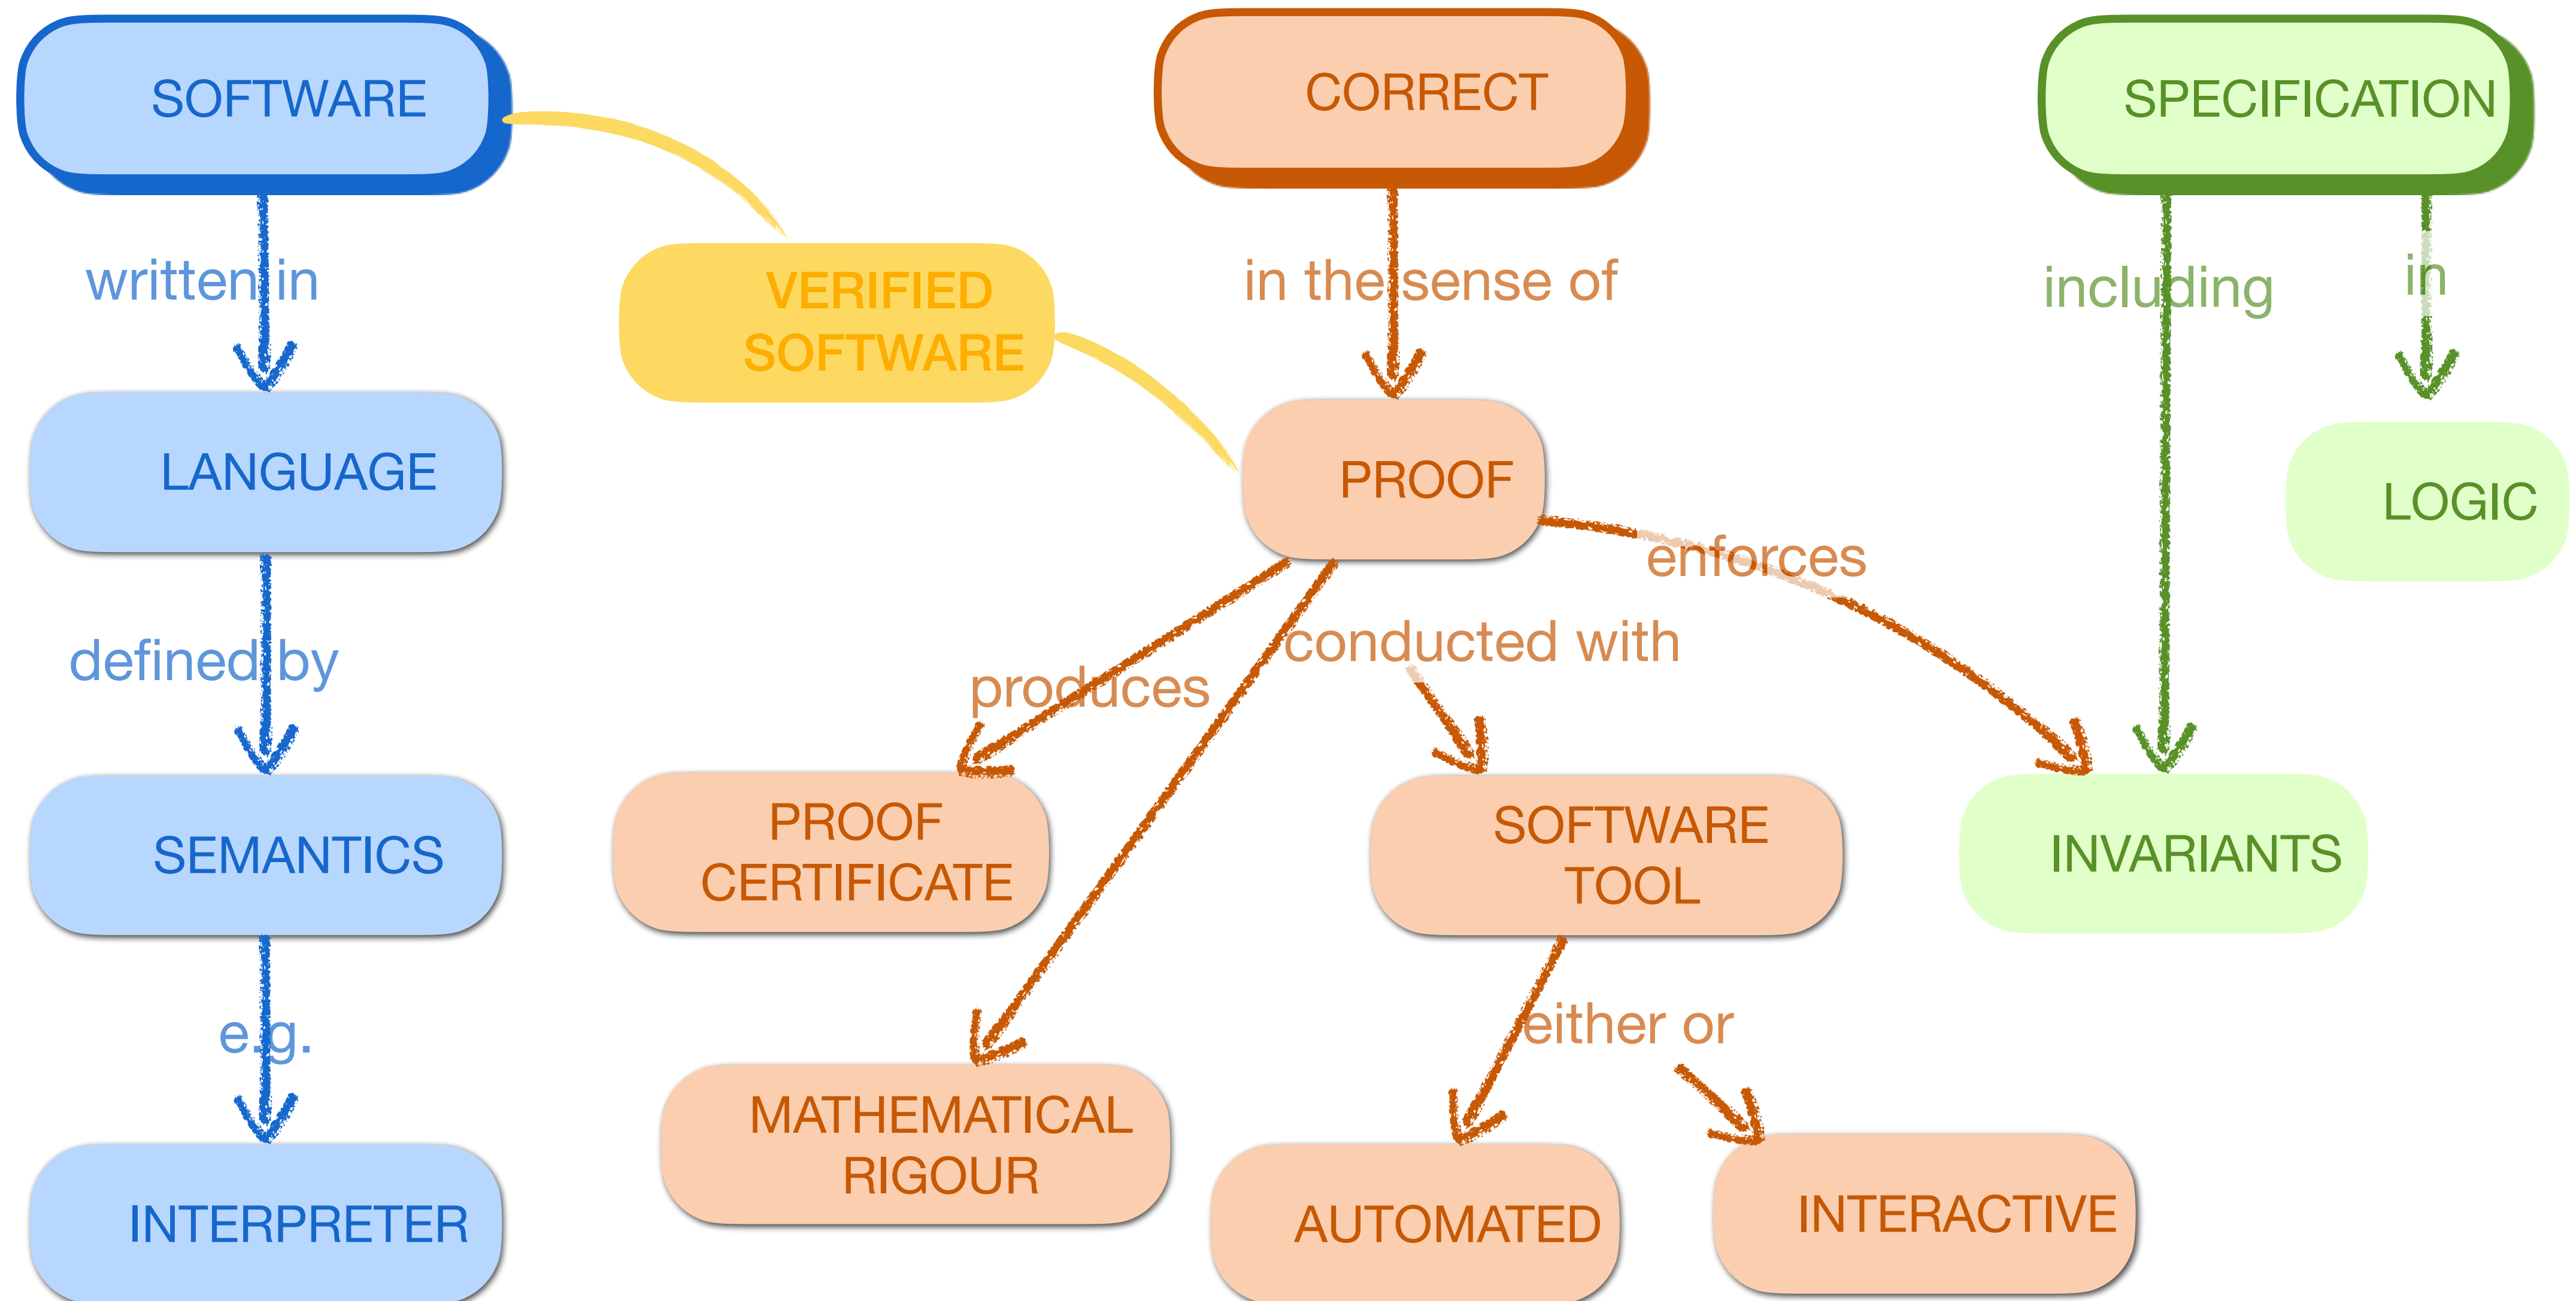
From early intuitions ...

A. M. Turing.
Checking a large routine. 1949.



... to deductive-verification and automated tools

Floyd 1967, Hoare 1969



Another historical example

Boyer-Moore's majority. 1980

Given N votes, determine the majority if any

A	A	A	C	C	B	B	C	C	C	B	C	C
---	---	---	---	---	---	---	---	---	---	---	---	---



majority = A

delta = 3

MJRTY—A Fast Majority Vote Algorithm¹

Robert S. Boyer and J Strother Moore

Computer Sciences Department
University of Texas at Austin

and

Computational Logic, Inc.
1717 West Sixth Street, Suite 290
Austin, Texas

Abstract

A new algorithm is presented for determining which, if any, of an arbitrary number of candidates has received a majority of the votes cast in an election. The number of comparisons required is at most twice the number of votes. Furthermore, the algorithm uses storage in a way that permits an efficient use of magnetic tape. A Fortran version of the algorithm is exhibited. The Fortran code has been proved correct by a mechanical verification system for Fortran. The system and the proof are discussed.

Another historical example

Boyer-Moore's majority. 1980

Given N votes, determine the majority if any

A	A	A	C	C	B	B	C	C	C	B	C	C
---	---	---	---	---	---	---	---	---	---	---	---	---



majority = A
delta = 3

A	A	A	C	C	B	B	C	C	C	B	C	C
---	--------------	--------------	--------------	--------------	---	---	---	---	---	---	---	---



majority = A
delta = 1

MJRTY—A Fast Majority Vote Algorithm¹

Robert S. Boyer and J Strother Moore

Computer Sciences Department
University of Texas at Austin

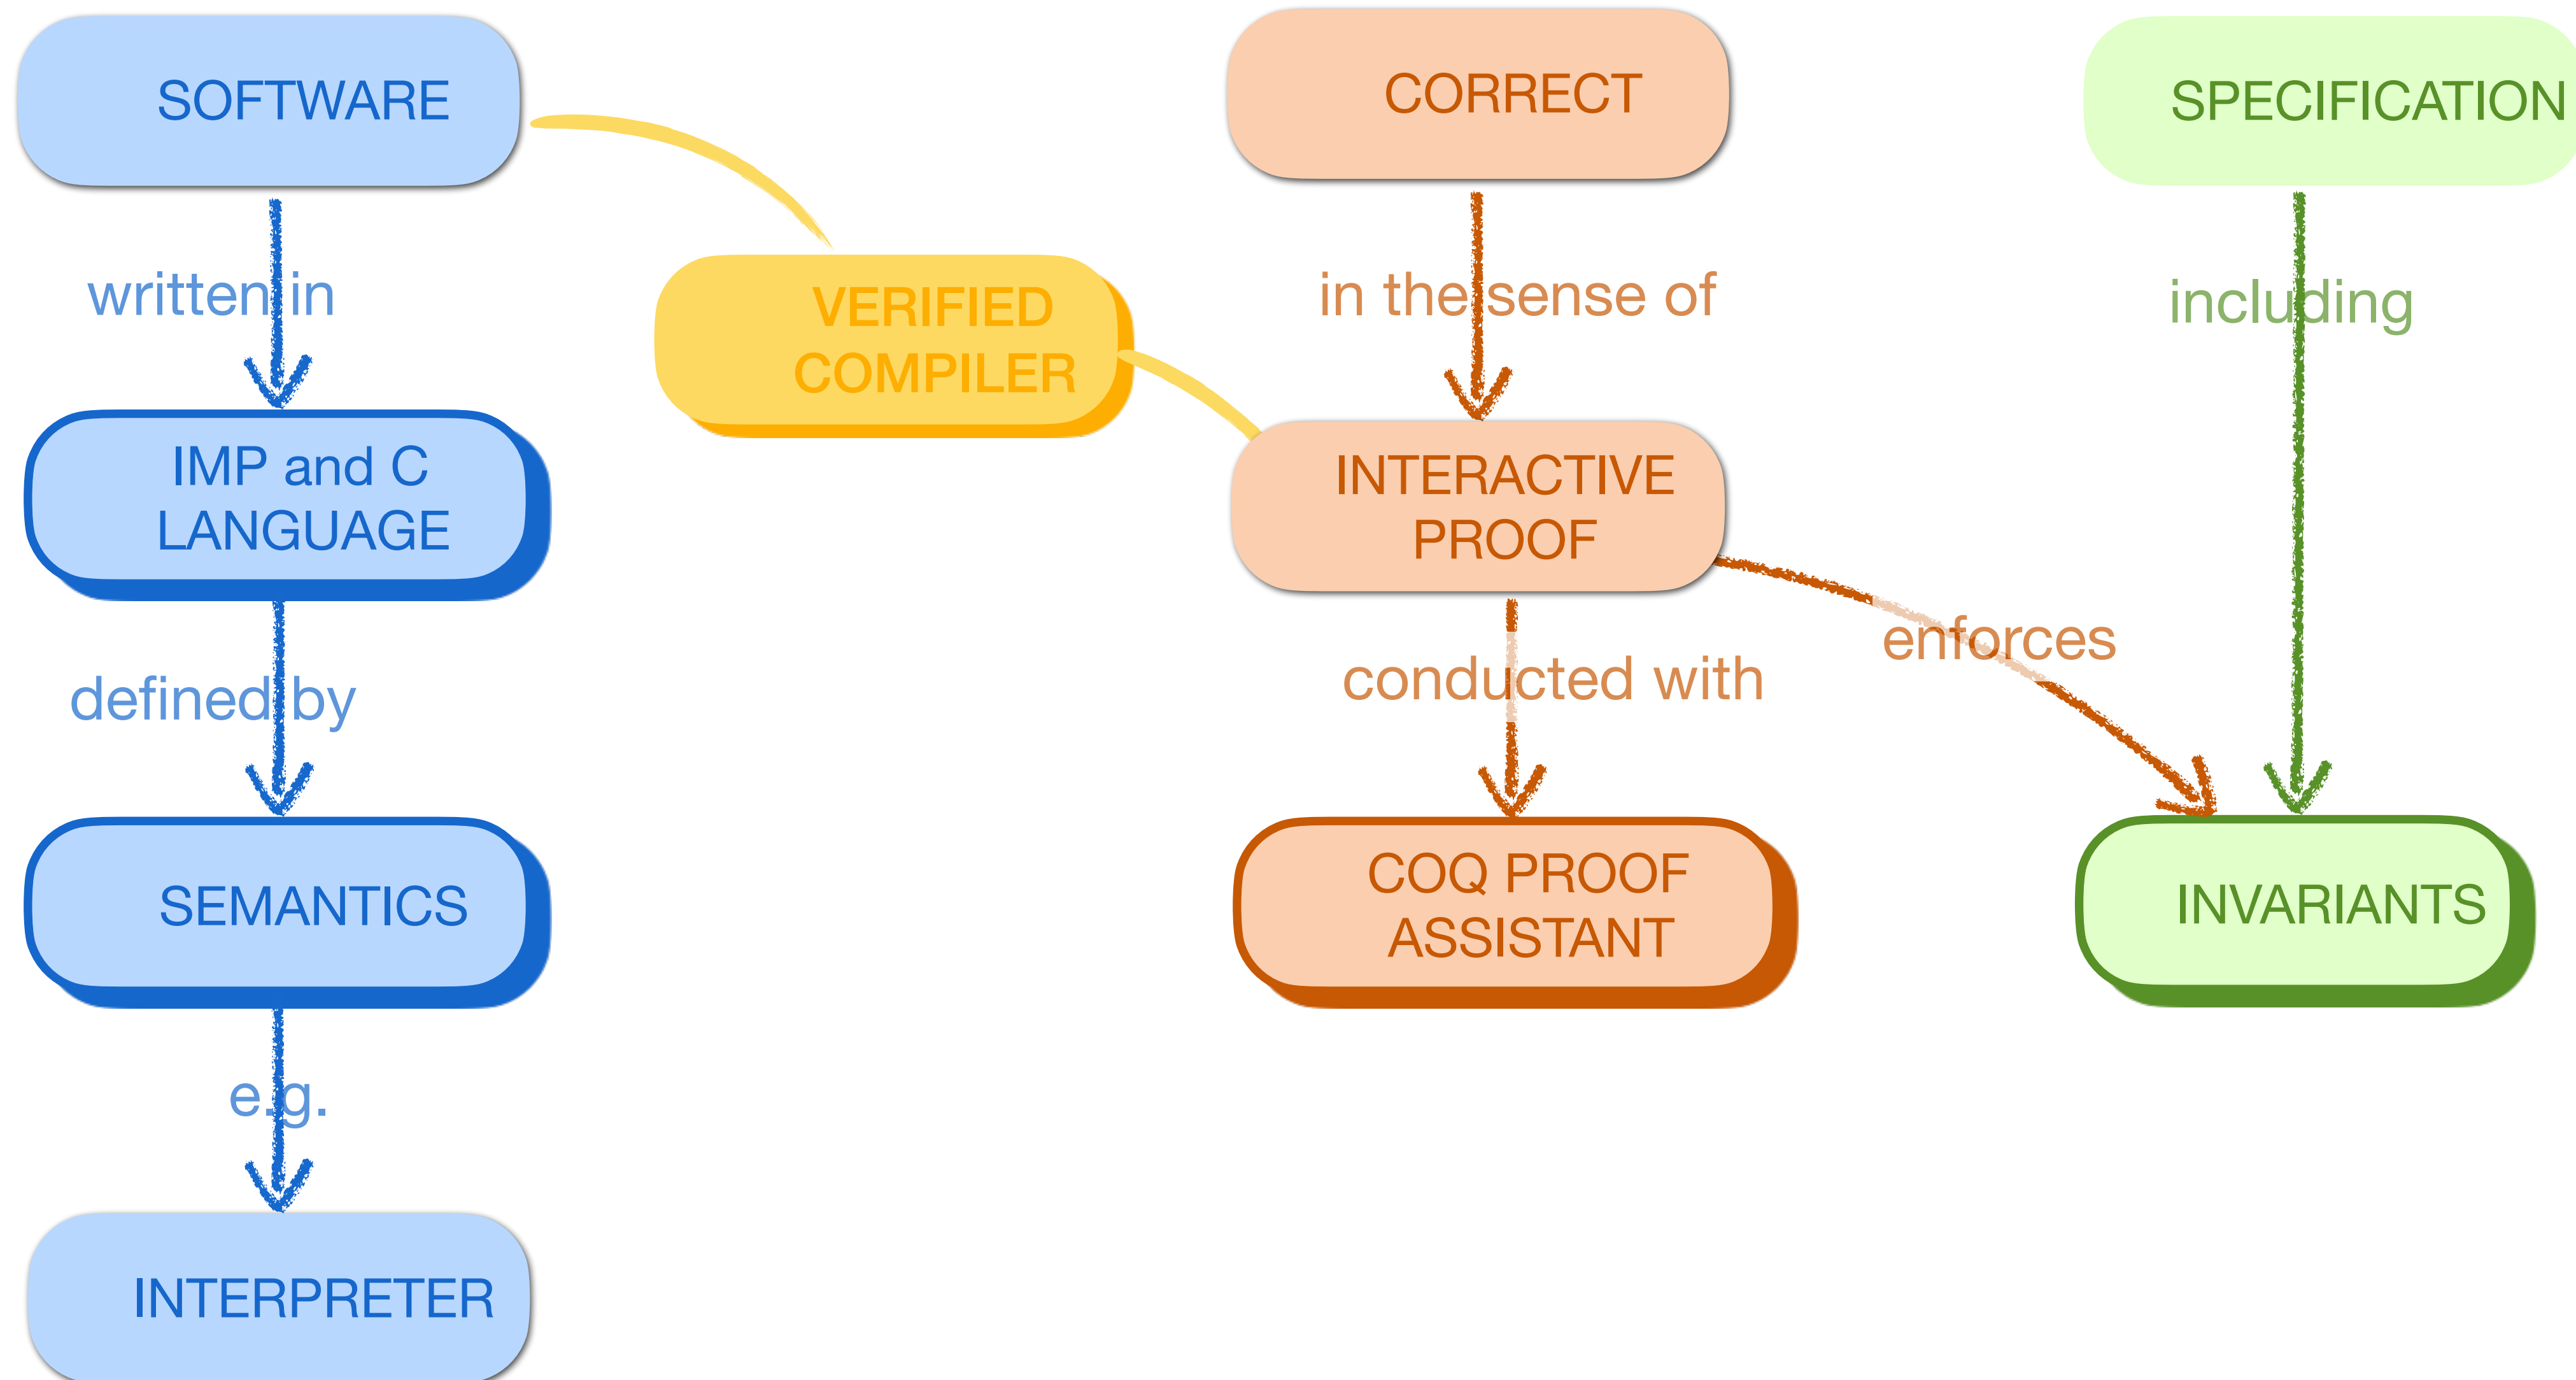
and

Computational Logic, Inc.
1717 West Sixth Street, Suite 290
Austin, Texas

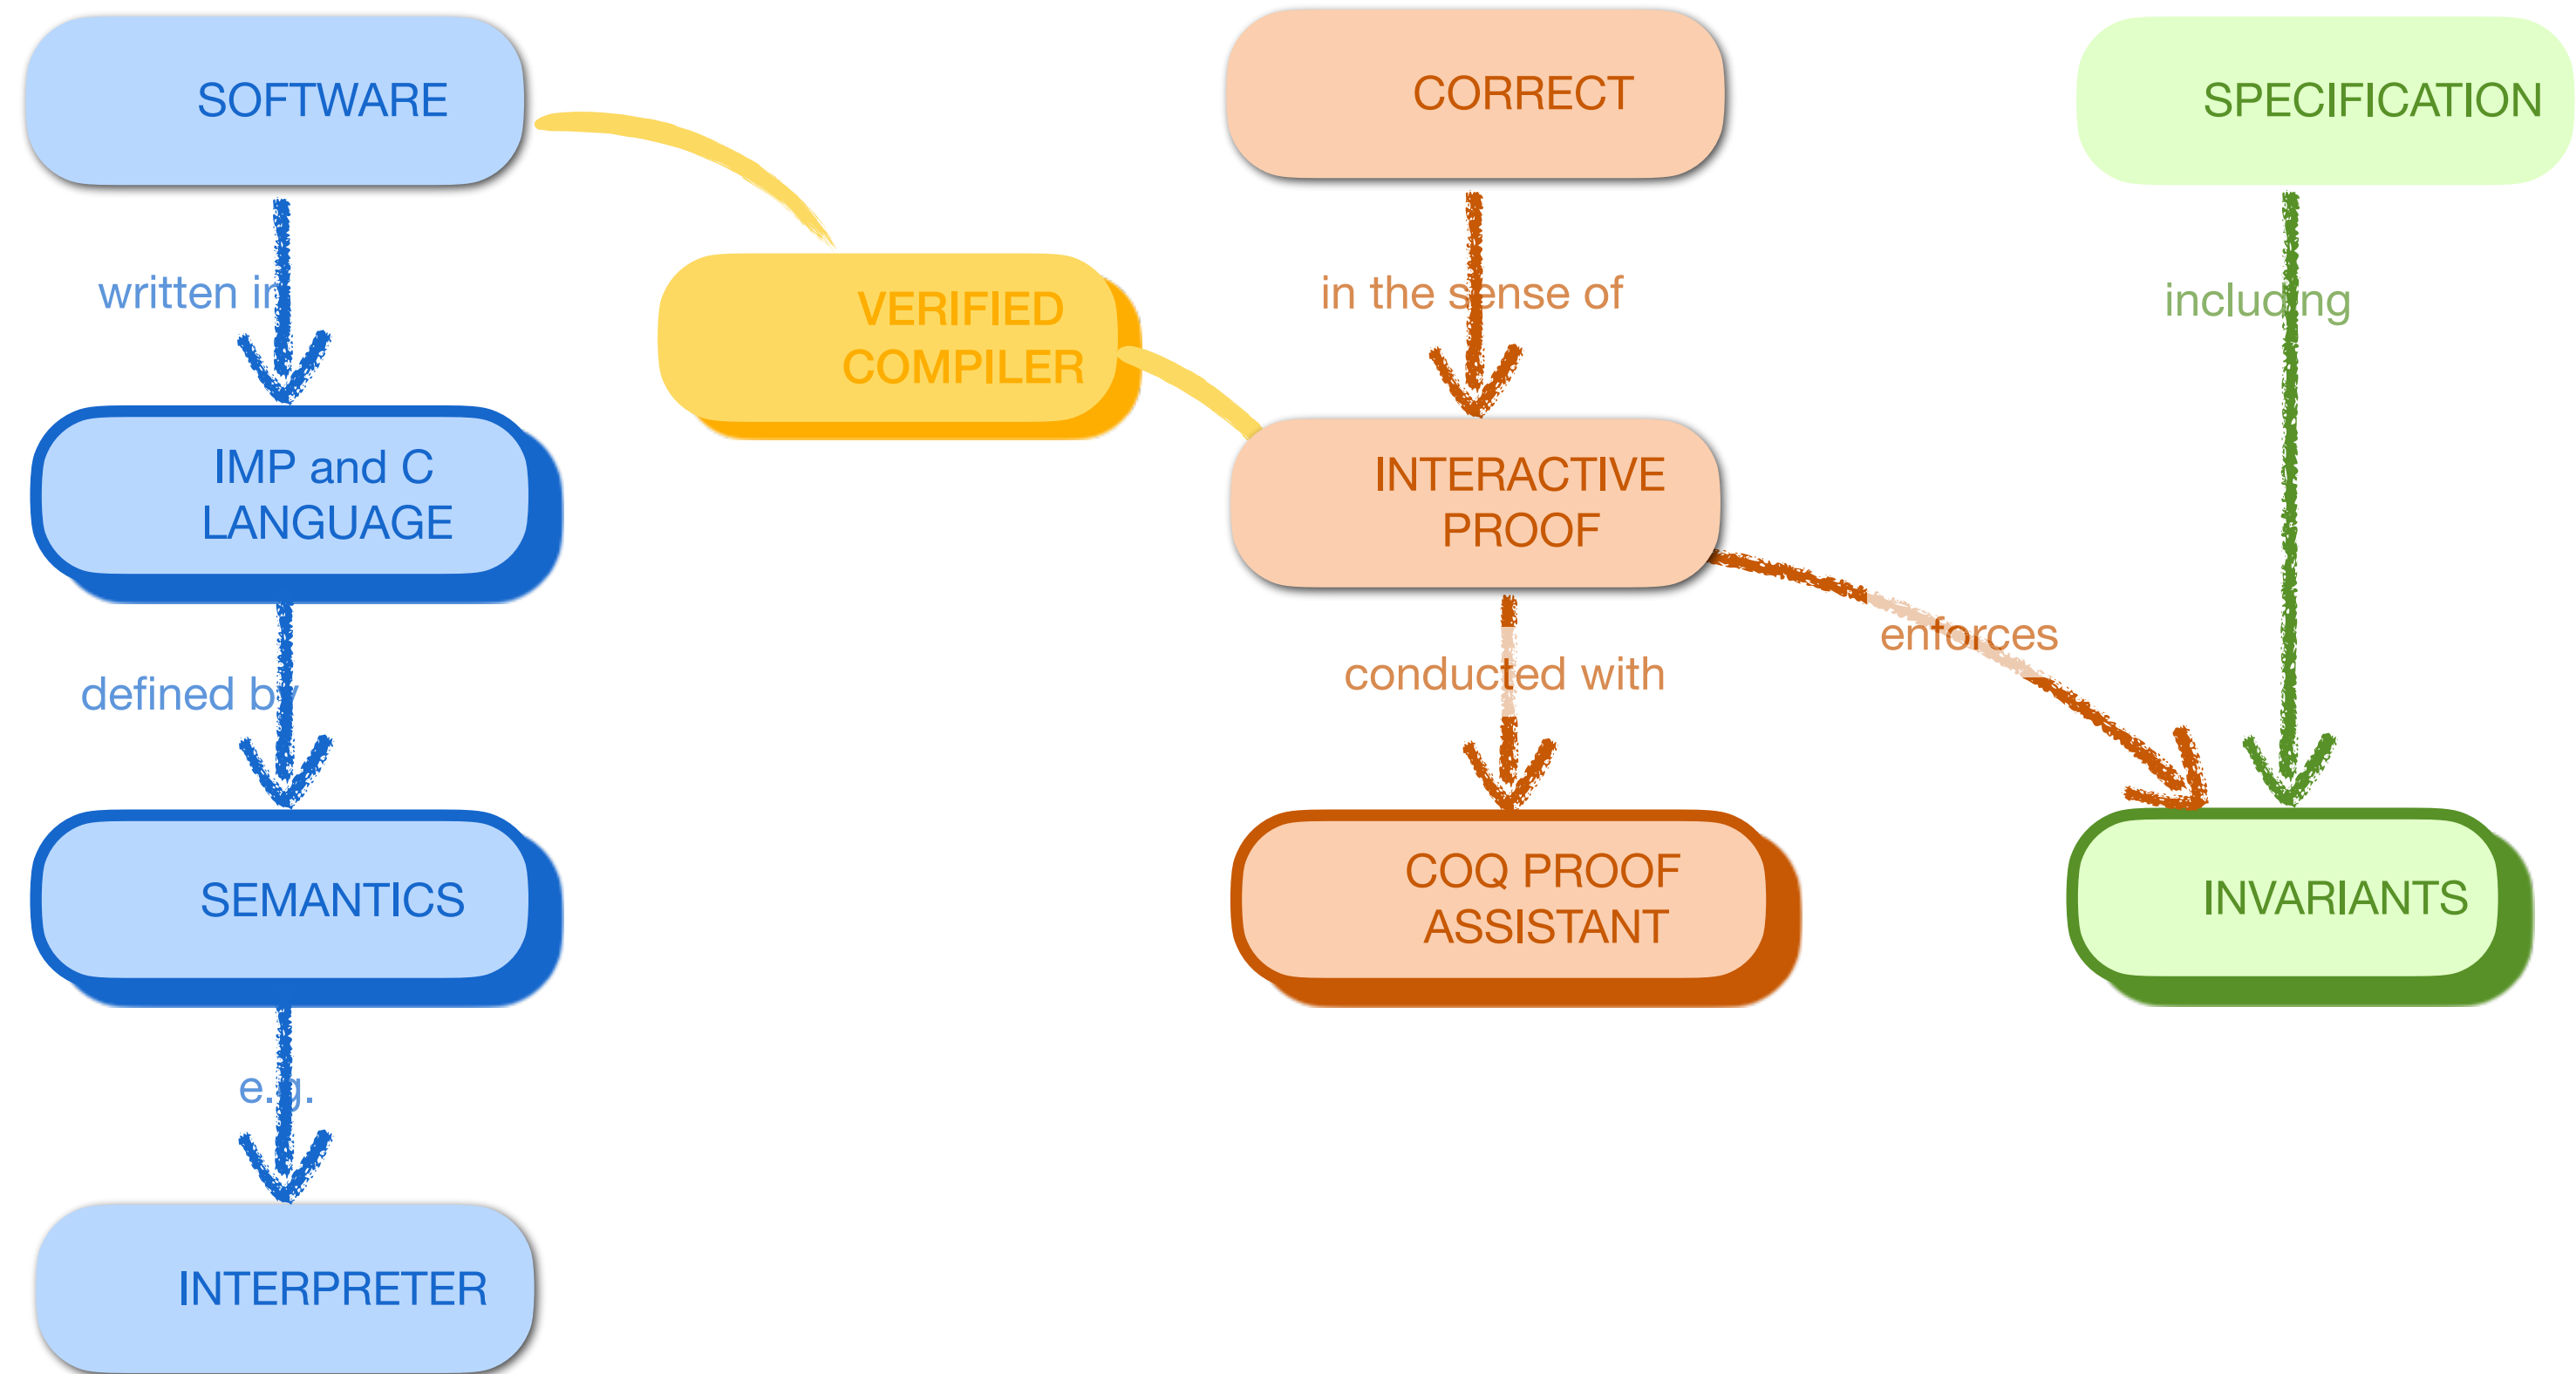
Abstract

A new algorithm is presented for determining which, if any, of an arbitrary number of candidates has received a majority of the votes cast in an election. The number of comparisons required is at most twice the number of votes. Furthermore, the algorithm uses storage in a way that permits an efficient use of magnetic tape. A Fortran version of the algorithm is exhibited. The Fortran code has been proved correct by a mechanical verification system for Fortran. The system and the proof are discussed.

Part 1: summary



Part 2 Early intuitions



Verified compilation

Compilers are complicated programs, but have a rather simple end-to-end specification:

The generated code must behave as prescribed by the semantics of the source program.

This specification becomes mathematically precise as soon as we have formal semantics for the source language and the machine language.

Then, a formal verification of a compiler can be considered.

An old idea ...

John McCarthy
James Painter¹

CORRECTNESS OF A COMPILER FOR ARITHMETIC EXPRESSIONS²

1. Introduction. This paper contains a proof of the correctness of a simple compiling algorithm for compiling arithmetic expressions into machine language.

The definition of correctness, the formalism used to express the description of source language, object language and compiler, and the methods of proof are all intended to serve as prototypes for the more complicated task of proving the correctness of usable compilers. The ultimate goal, as outlined in references [1], [2], [3] and [4] is to make it possible to use a computer to check proofs that compilers are correct.

Mathematical Aspects of Computer Science, 1967

3

Proving Compiler Correctness in a Mechanized Logic

R. Milner and R. Weyhrauch

Computer Science Department
Stanford University

Abstract

We discuss the task of machine-checking the proof of a simple compiling algorithm. The proof-checking program is LCF, an implementation of a logic for computable functions due to Dana Scott, in which the abstract syntax and extensional semantics of programming languages can be naturally expressed. The source language in our example is a simple ALGOL-like language with assignments, conditionals, whiles and compound statements. The target language is an assembly language for a machine with a pushdown store. Algebraic methods are used to give structure to the proof, which is presented only in outline. However, we present in full the expression-compiling part of the algorithm. More than half of the complete proof has been machine checked, and we anticipate no difficulty with the remainder. We discuss our experience in conducting the proof, which indicates that a large part of it may be automated to reduce the human contribution.

Machine Intelligence (7), 1972

Now taught as an exercise to Masters students

(Mechanized semantics: when machines reason about their languages, X.Leroy)

(Software foundations, B.Pierce et al.)

```
type exp = Nb int | Id string | Plus exp exp
```

```
type state = string → int
```

```
let rec eval (s:state)(a:exp): int =  
  match a with  
  | Nb n → n  
  | Id x → s x  
  | Plus (a1,a2) → (eval s a1)+(eval s a2)
```

```
let rec compile (a:exp): instr list = match a with  
  | Nb n → [ Push n ]  
  | Id x → [ Load x ]  
  | Plus (a1,a2) → (compile a1)@ (compile a2)@ [ IPlus ]
```

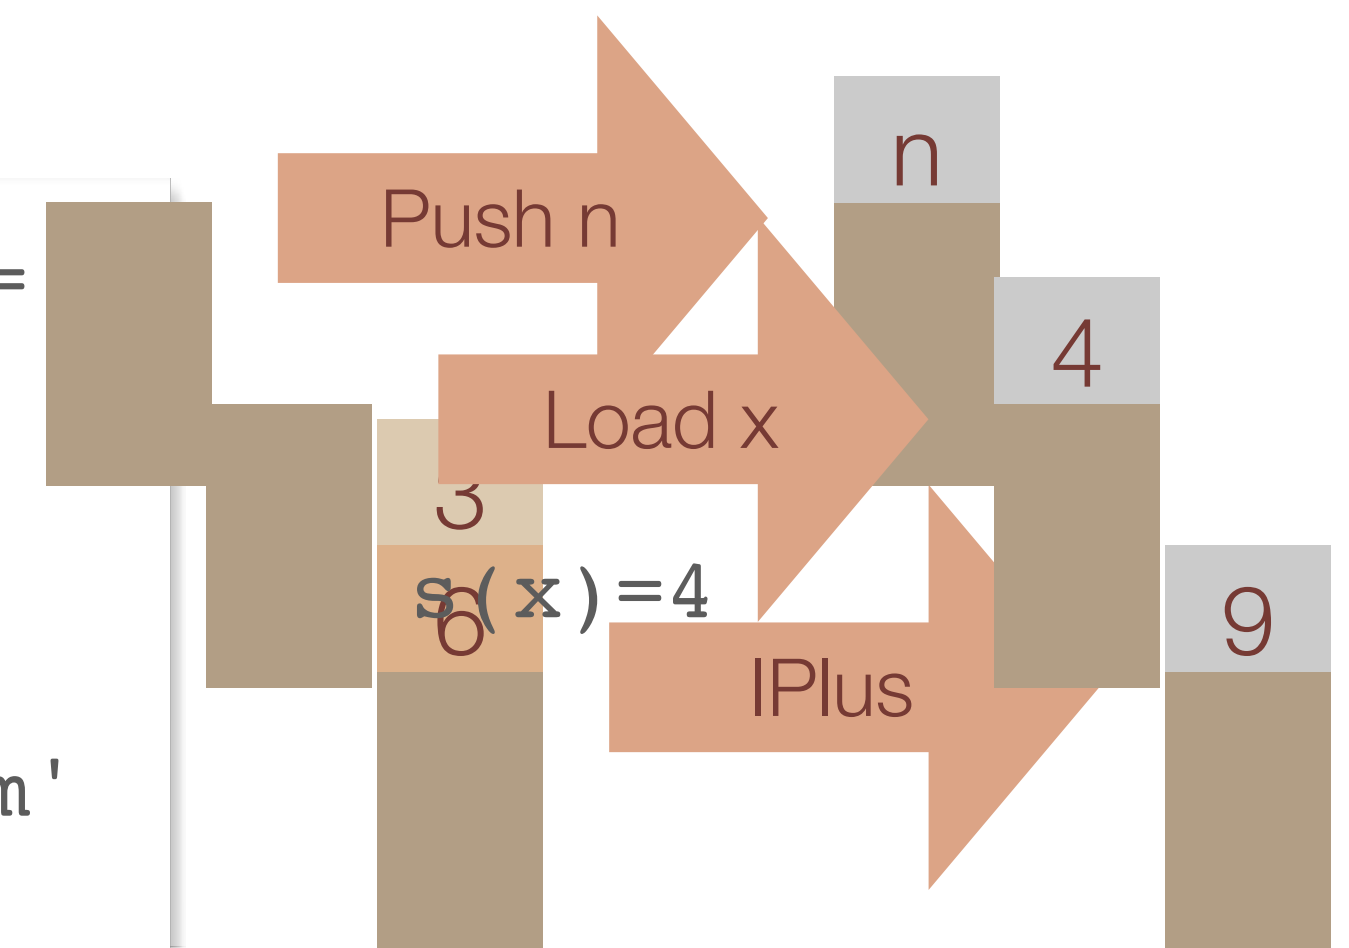
```
type instr = Push int | Load string | IPlus
```

```
let rec exec(s:state)(stack: int list)(pgm: instr list): int list =  
  match (pgm, stack) with  
  | ([], _) → stack  
  | (Push n :: pgm', _) → exec s (n :: stack) pgm'  
  | (Load x :: pgm', _) → exec s (s x :: stack) pgm'  
  | (IPlus :: pgm', n:: m :: stack') → exec s ((m+n) :: stack') pgm'  
  | (_ :: pgm', _) → exec s stack pgm'
```

semantics
(eval, exec)

compiler
(compile)

compilation



Proving a property with the Coq software

ACM SIGPLAN Programming Languages Software award 2013

ACM Software System award 2013

coq.inria.fr

```
Theorem toy-compiler-correct:  
  forall s a,  
  exec s [] (compile a) = [eval s a].
```

semantics
(eval, exec)

compiler
(compile)

Proving a property with the Coq software

ACM SIGPLAN Programming Languages Software award 2013

ACM Software System award 2013

coq.inria.fr

```
Theorem toy-compiler-correct:  
  forall s a,  
    exec s [] (compile a) = [eval s a].  
Proof.  
  intros;  
  ... (* not shown here *)  
Qed.
```

```
Extraction compile.
```

semantics
(eval, exec)

compiler
(compile)

proof
guided by Coq


extraction

compiler.ml



Part 3

How to turn CompCert
from a prototype in a lab
into a real-world compiler?



A selection of formally verified compilers

CompCert C compiler (Coq) [Leroy, POPL'06]

CakeML ML bootstrapped compiler (HOL)
[Kumar, Myreen, Norrish, Owens, POPL'14]

CertiCoq Gallina compiler (Coq) [Appel et al., CoqPL'17]

Jasmin language and compiler for cryptographic implementations (Coq)
[Almeida et.al, CCS'17]

The CompCert formally verified compiler

(X.Leroy, S.Blazy et al.)

<https://compcert.org>

A moderately optimizing C compiler

Targets several architectures (PowerPC, ARM, RISC-V and x86)

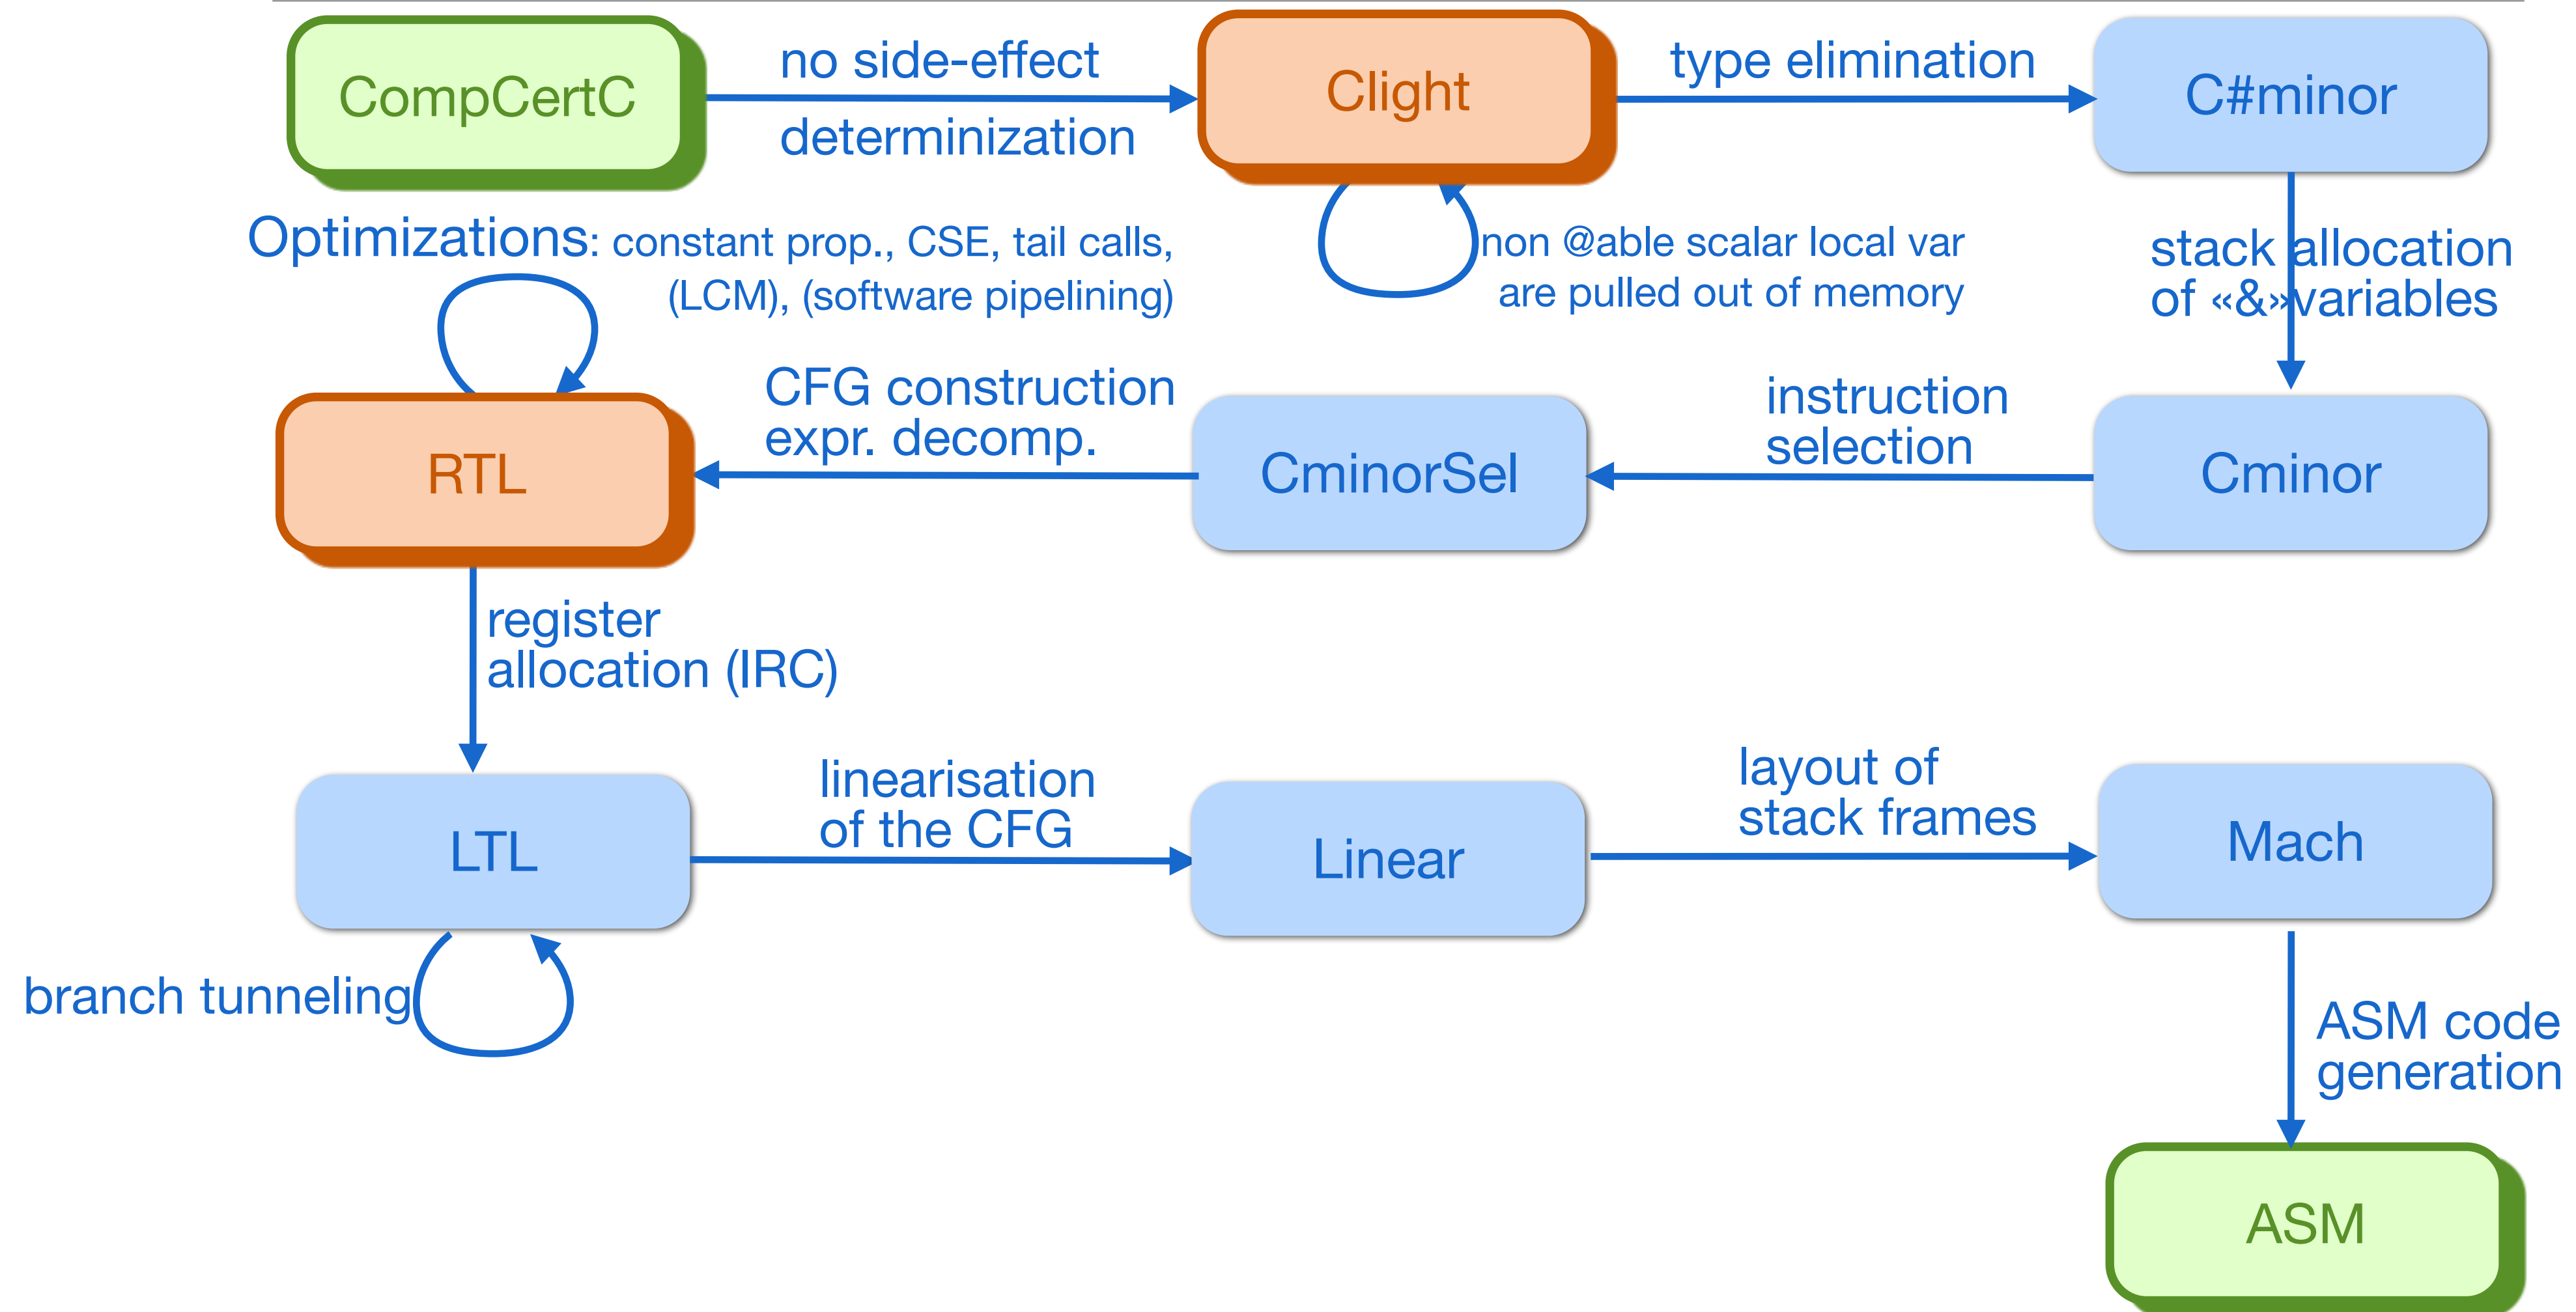
Used in commercial settings (for emergency power generators and flight control navigation algorithms) and for software certification - AbsInt company

Improved performances of the generated code while providing proven traceability information

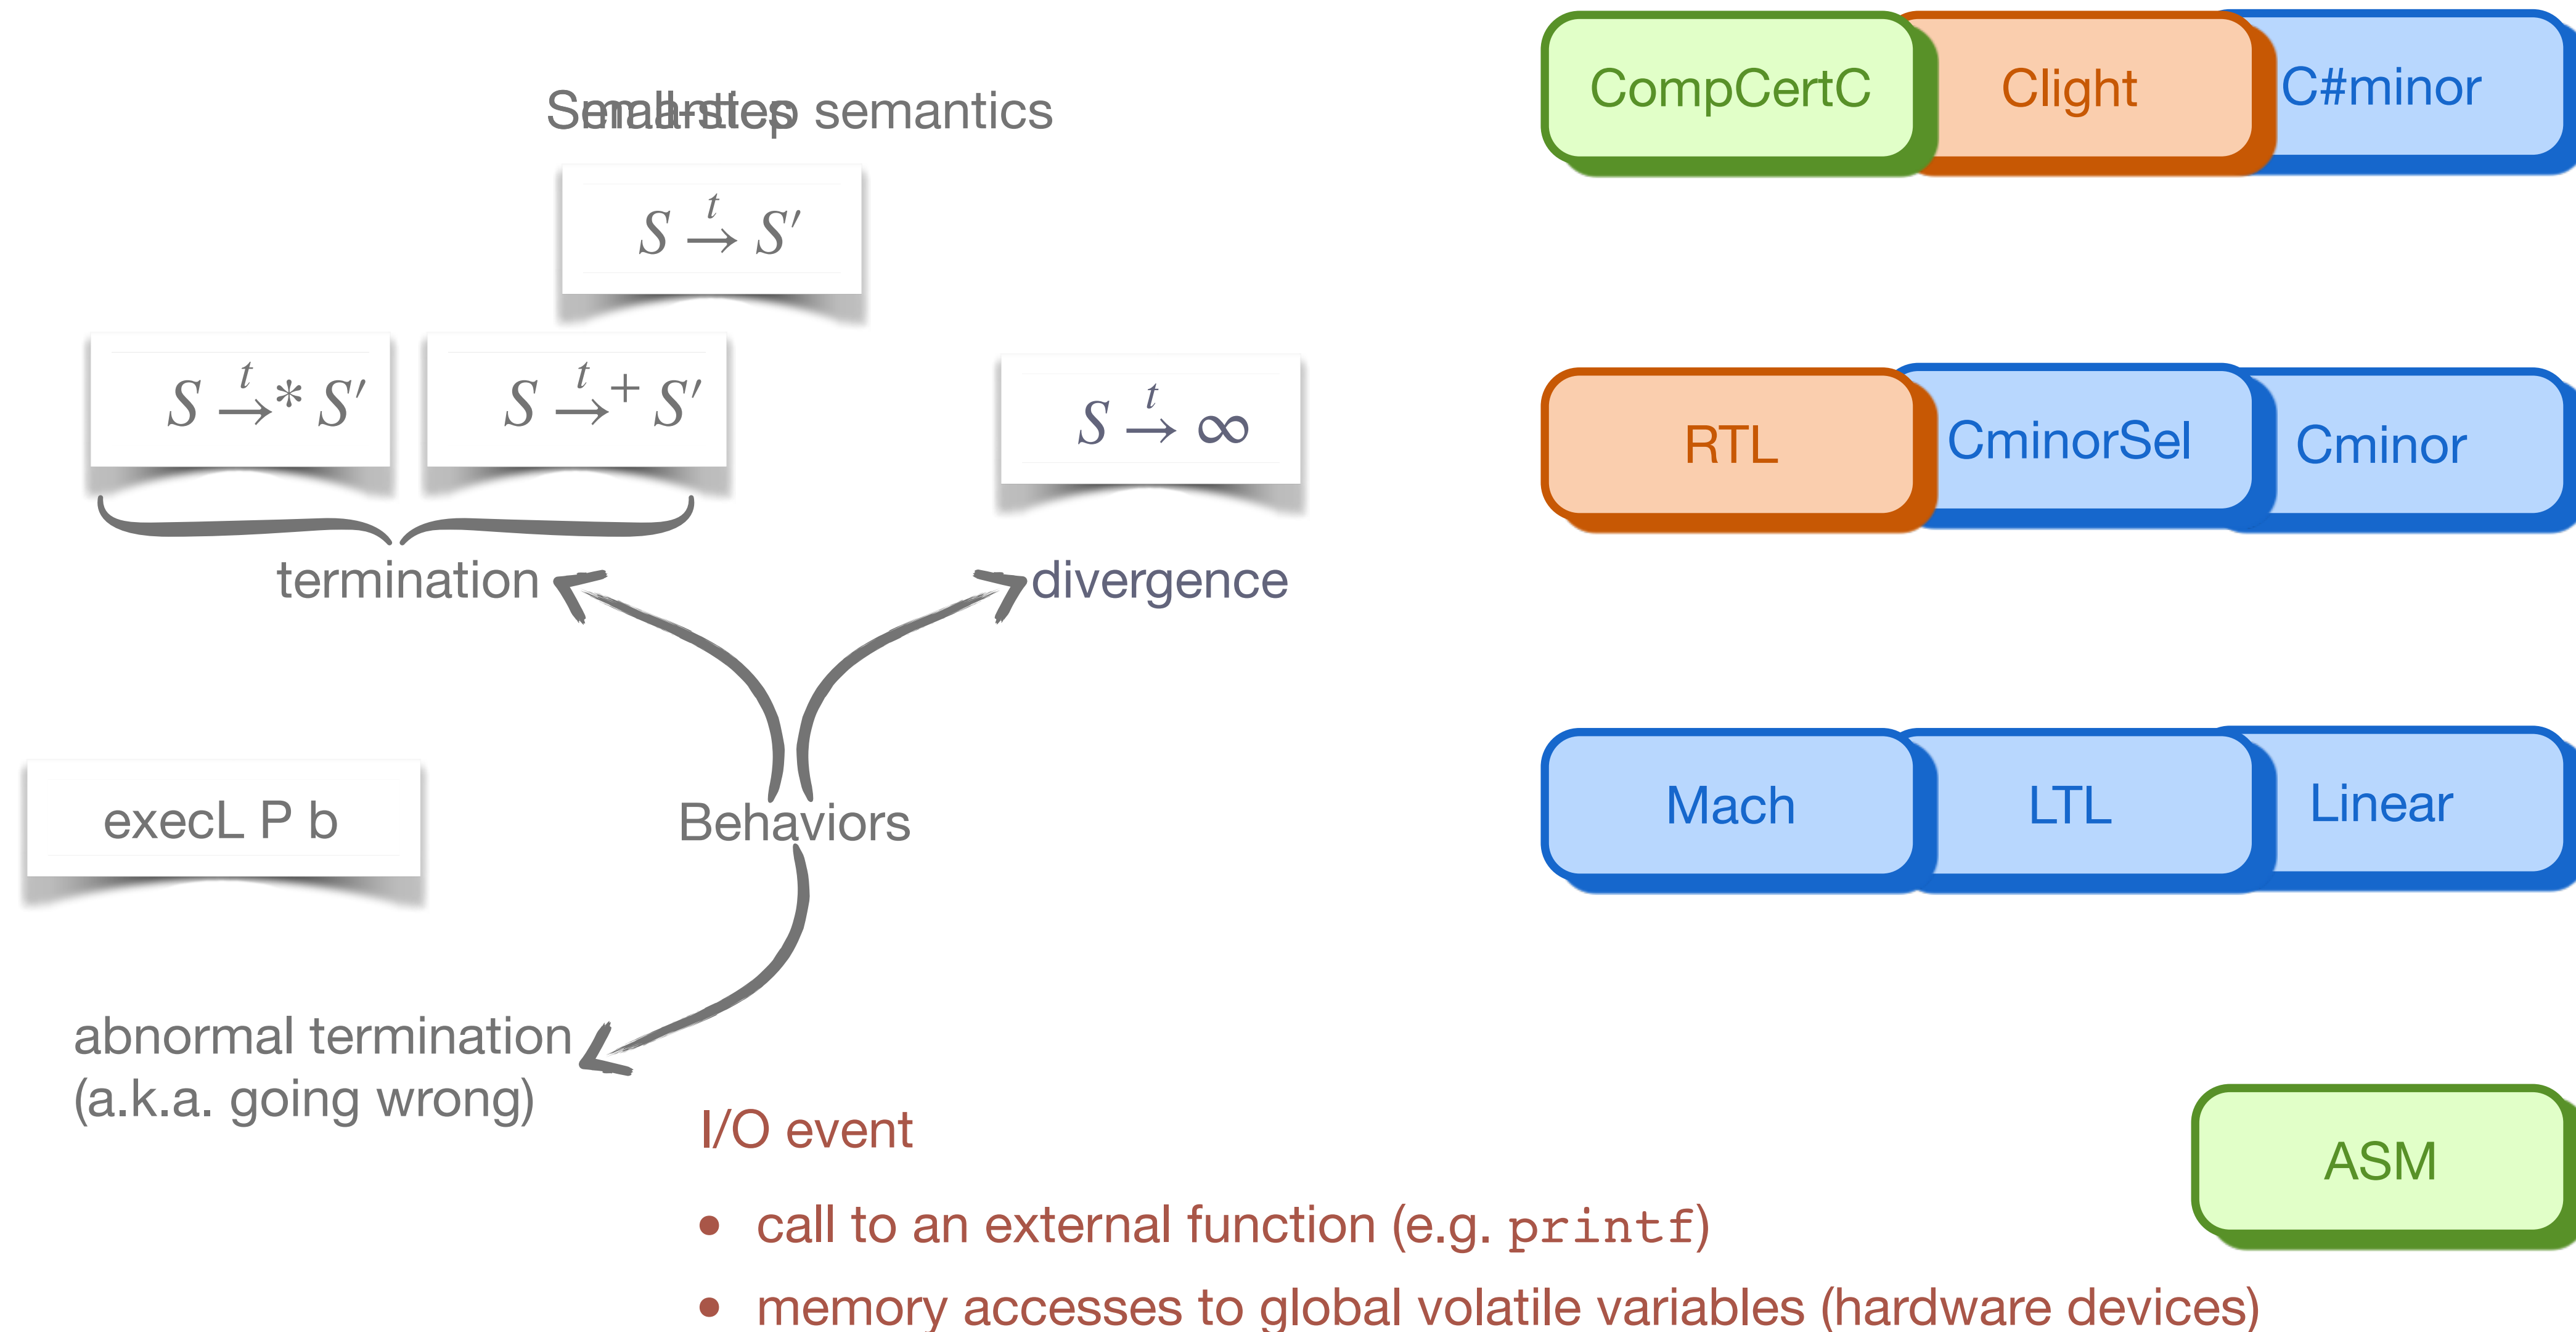
ACM Software System award 2021

ACM SIGPLAN Programming Languages Software award 2022

CompCert compiler: 10 languages, 18 passes



CompCert compiler: 10 languages, 18 passes



Proving semantics preservation: the simulation approach

semantics
(`execSource`, `execTarget`)


compiler

Preserved **behaviors** = termination and divergence

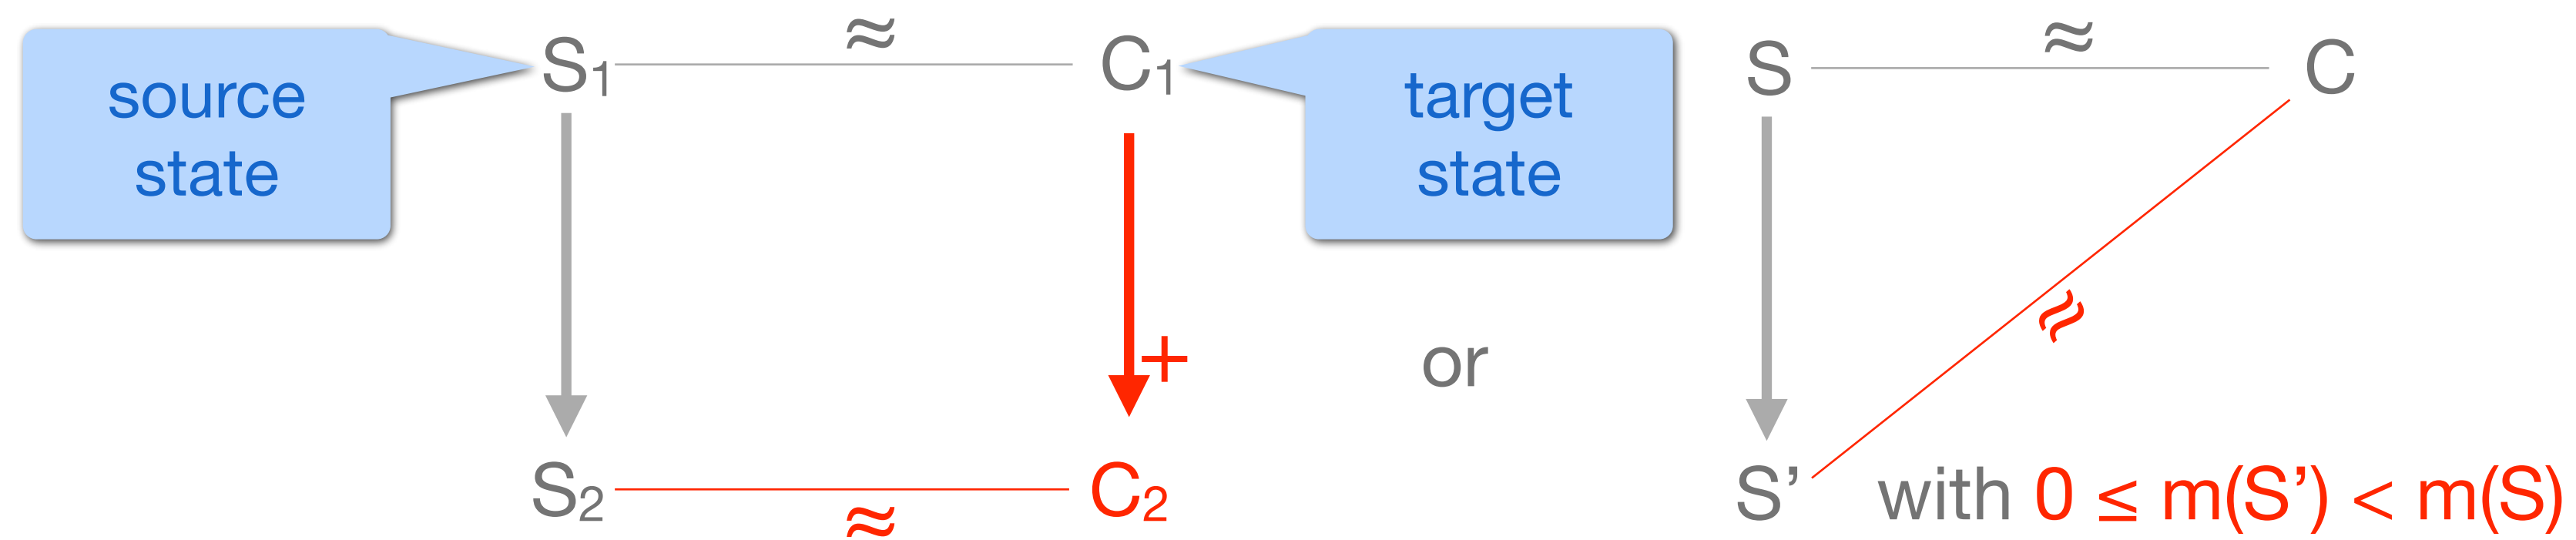
Theorem compiler-correct:
 $\forall s c b,$
`compiler` $s = \text{OK } c \rightarrow$
`execSource` $s b \rightarrow$
`execTarget` $c b.$

« The generated code must
behave as prescribed by the
semantics of the source
program. »

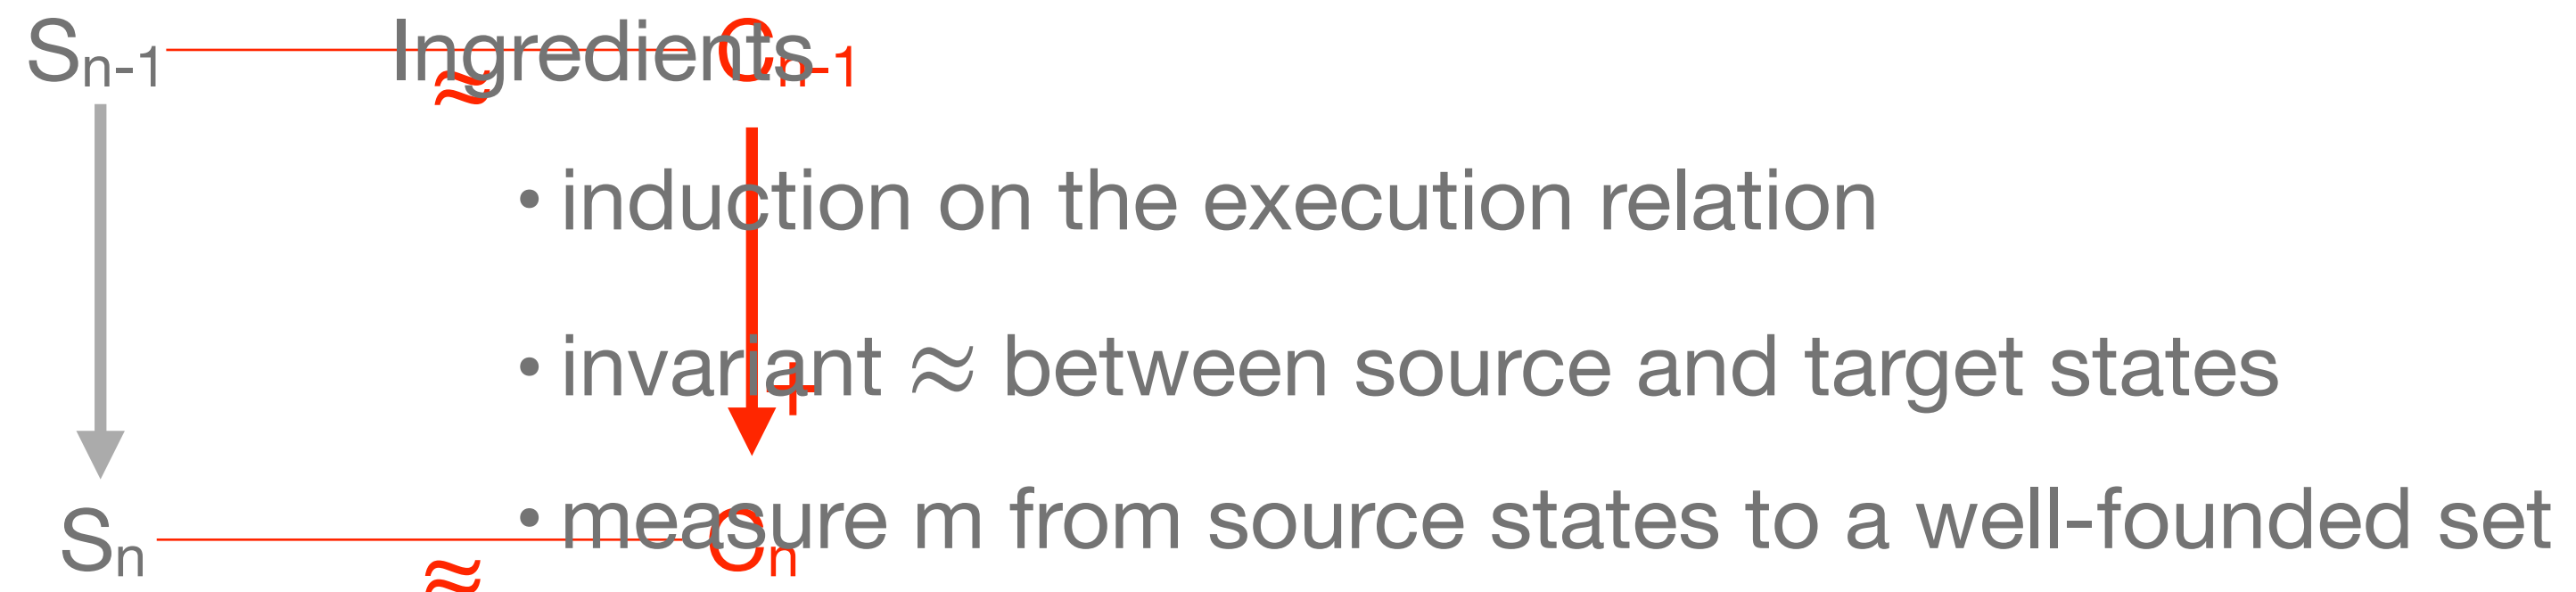
Proof technique: simulation diagram



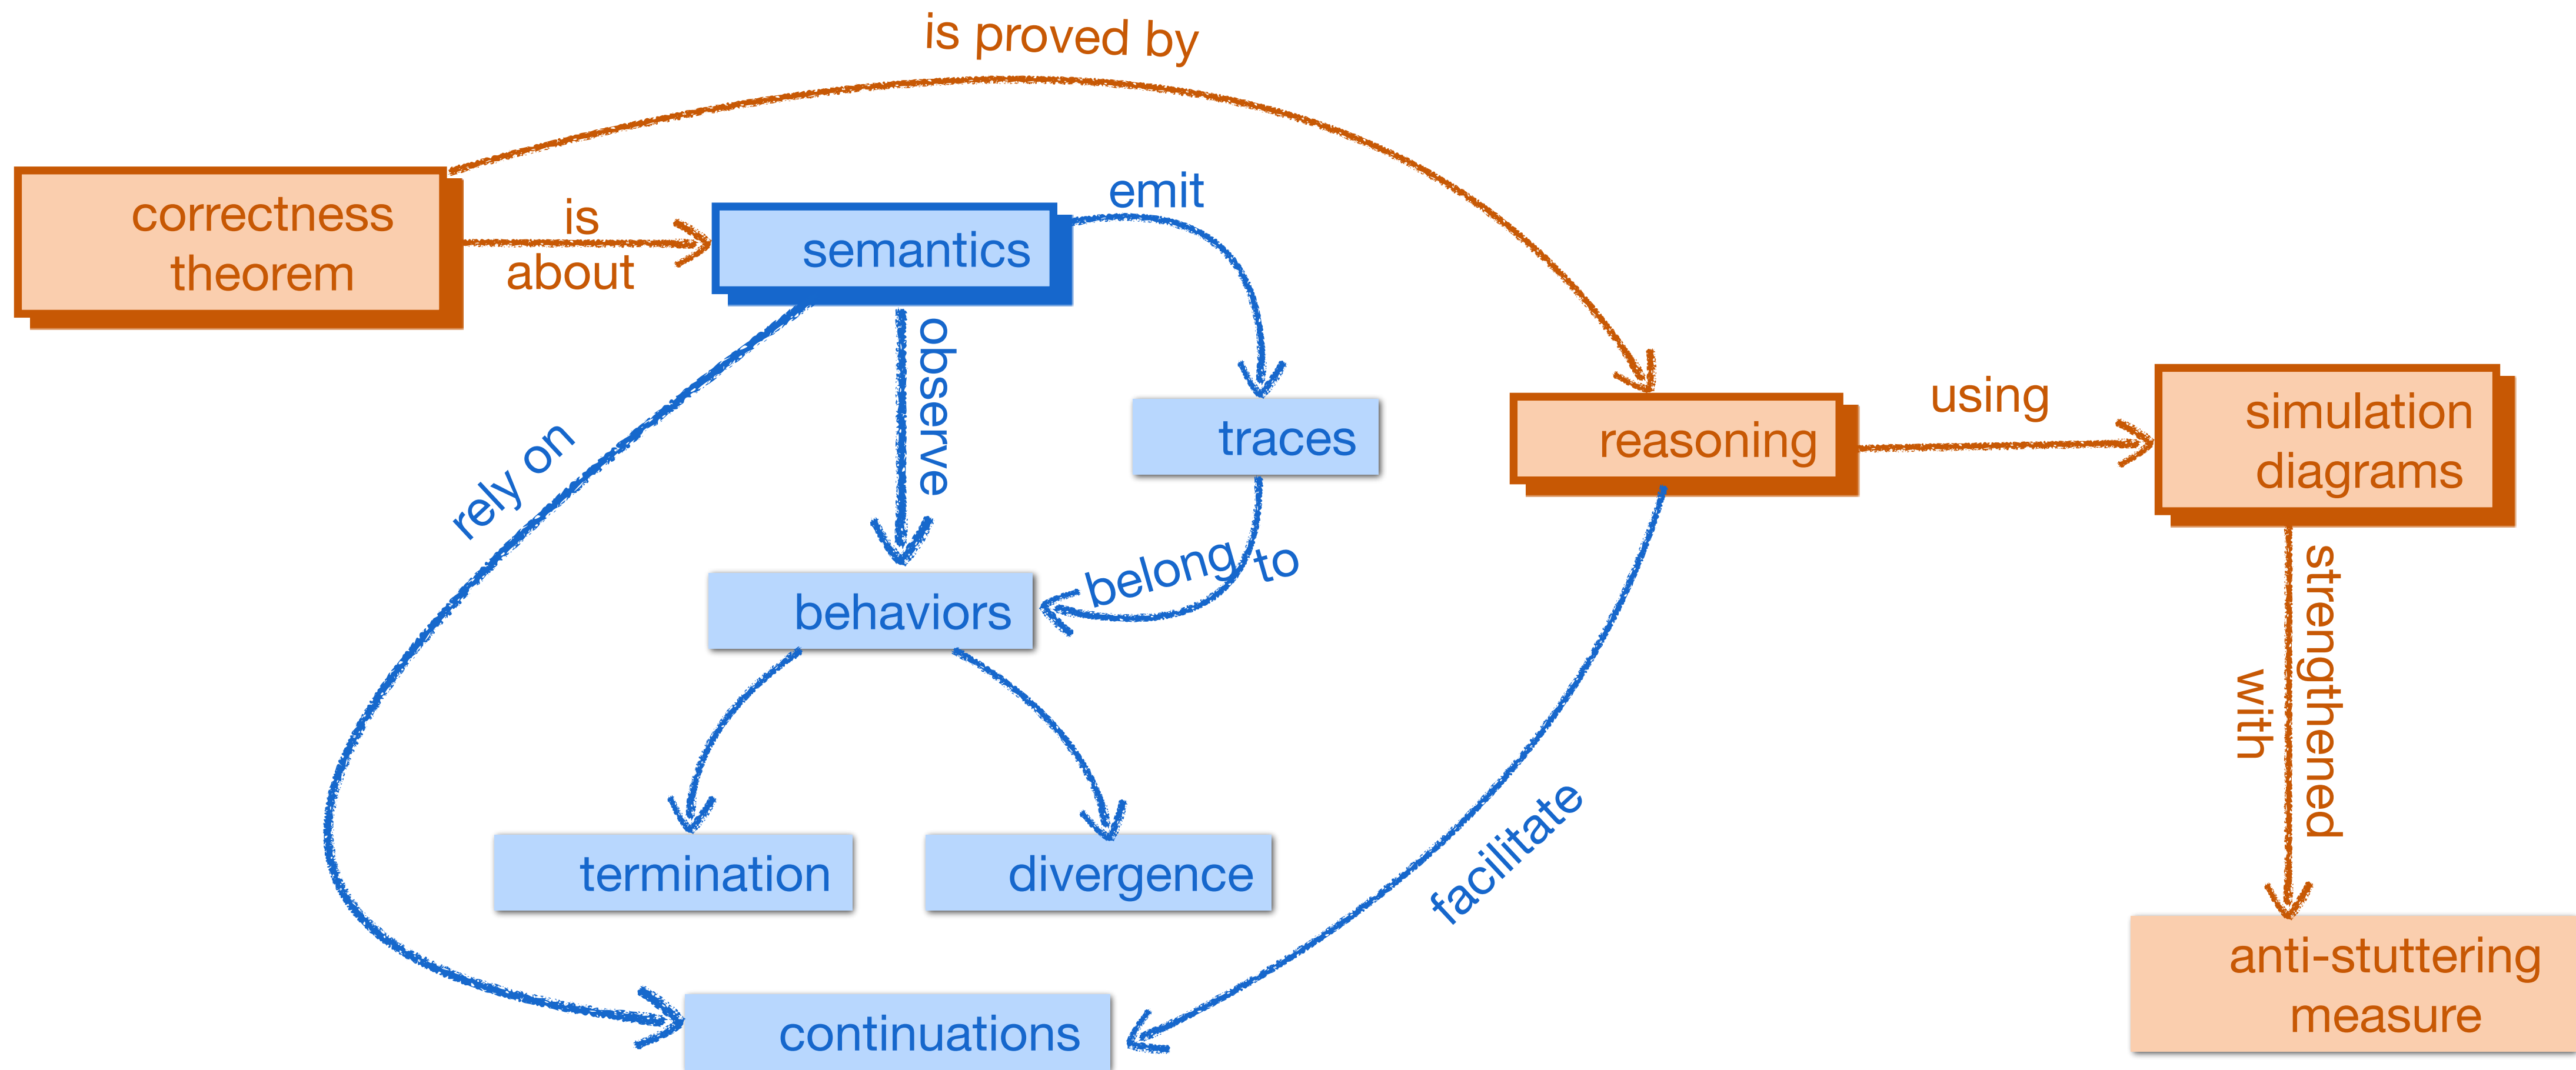
Proving semantics preservation: the simulation approach



If the source program diverges, it must perform infinitely many non-stuttering steps, so the compiled code executes infinitely many steps.



Semantic reasoning for compiler correctness: summary



CompCert verified compiler: main ingredients

Compilation	Formal semantics	Deductive verification
Source and target languages	Observable behaviors	Proof assistant
Intermediate language	Traces of ext. I/O events	Semantic preservation theorem
Optimizations	Small-step style	Simulation diagram
Data-flow analysis	Continuations	Anti-stuttering measure
Register allocation	Memory model	A posteriori validation
Other passes		

Part 4

Beyond CompCert:

- secure compilation
- just-in-time compilation



Turning CompCert into a secure compiler

CT-CompCert [Barthe, Blazy, Grégoire, Hutin, Laporte, Pichardie, Trieu, POPL'20]



Cryptographic constant-time (CCT) programming discipline

```
unsigned nok-function (unsigned x, unsigned y, bool secret)
{ if (secret) return y; else return x; }
```

```
✓ unsigned ok-function (unsigned x, unsigned y, bool secret)
{ return x ^ ((y ^ x) & (-(unsigned)secret)); }
```

How to turn CompCert into a formally-verified secure compiler?

Theorem compiler-correct:

```
∀ S C b,
  compiler S = OK C →
  execCompCertC S b →
  execASM C b.
```

Theorem compiler-preserves-CCT:

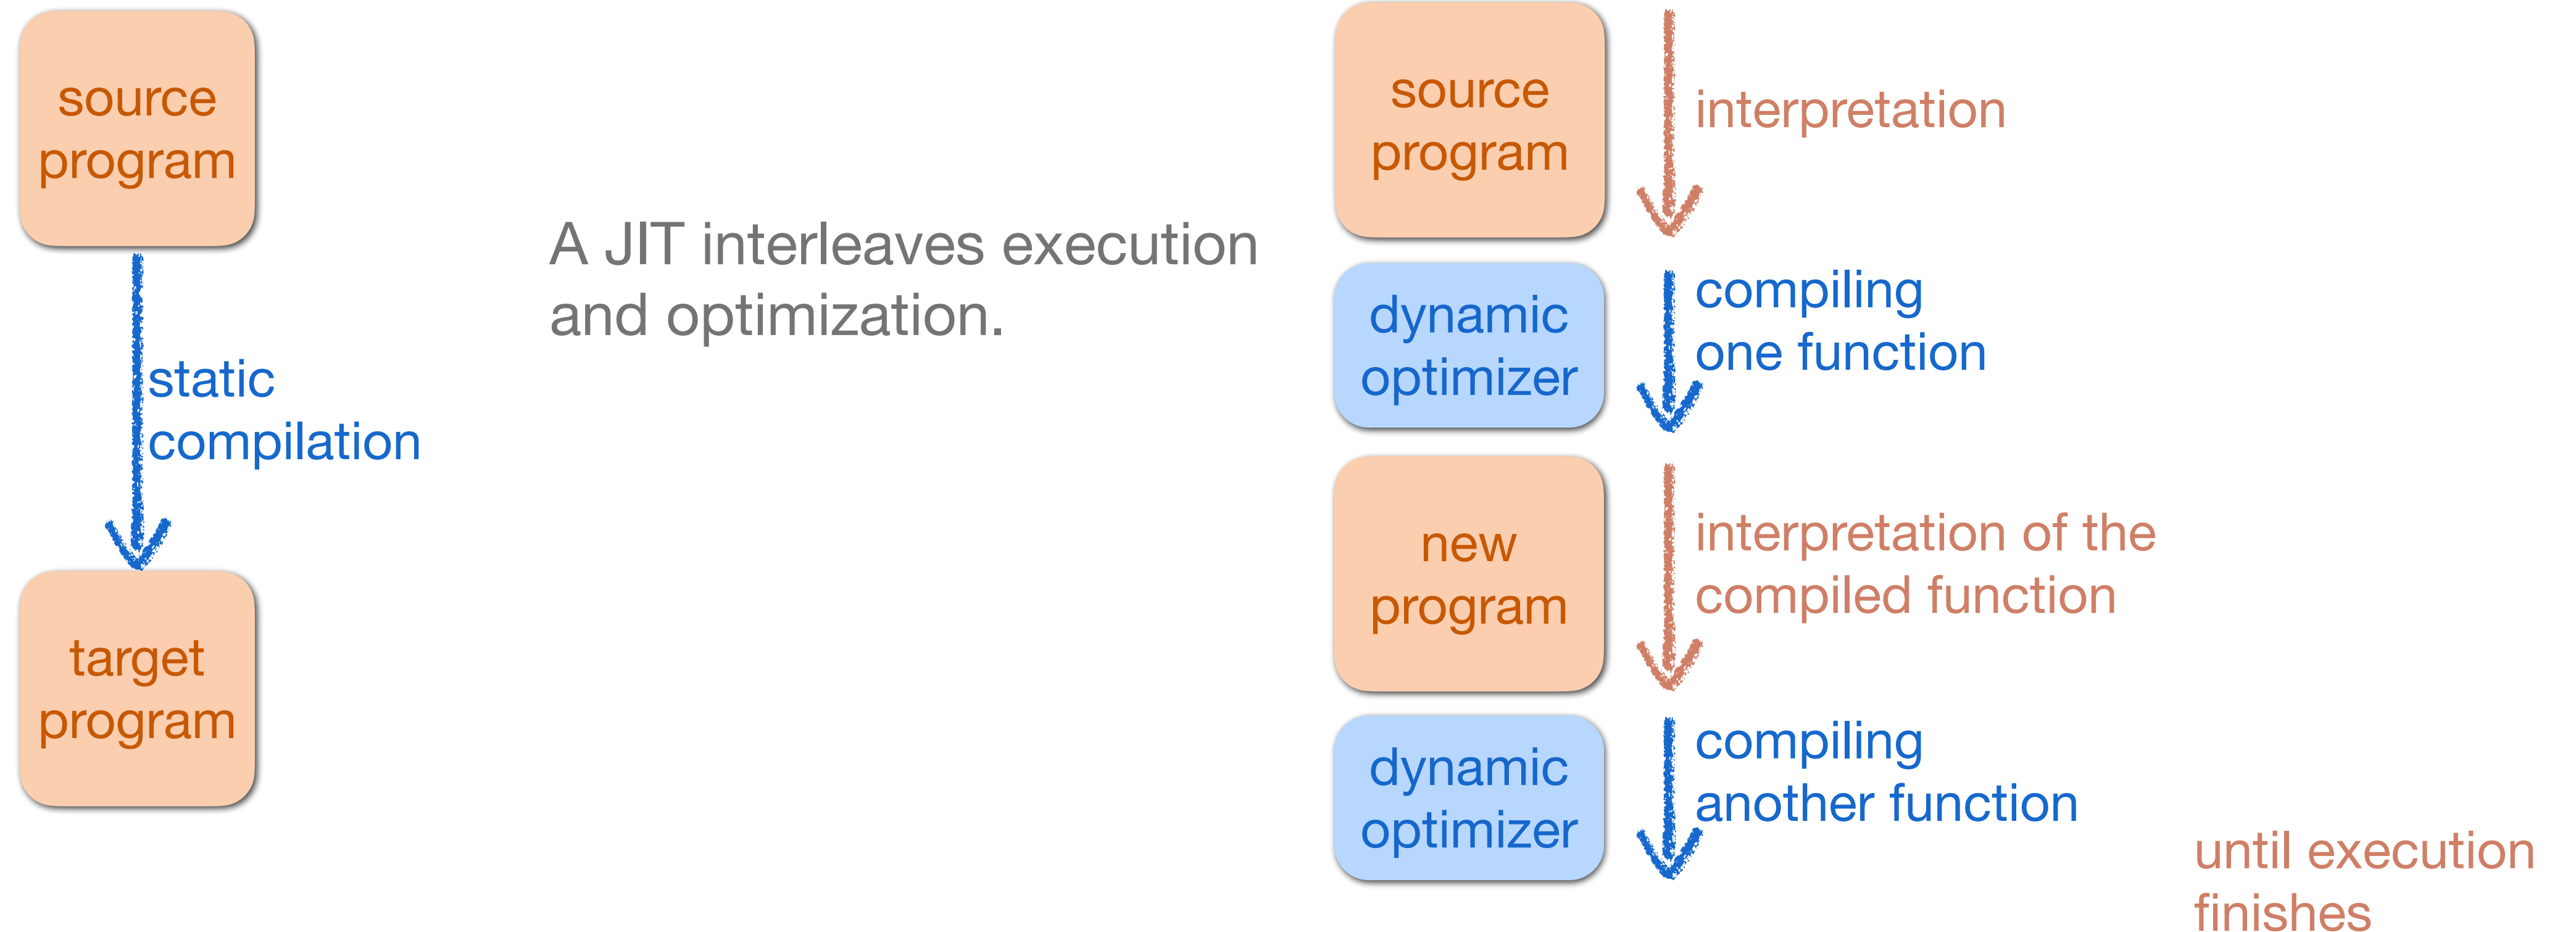
```
∀ S C,
  compiler S = OK C →
  isCCT S →
  isCCT C.
```

observe
program leakages (boolean guards
and memory accesses)

2 executions of S from 2
indistinguishable states (only share
public values)



Just-in-time (JIT) compilation vs. static compilation



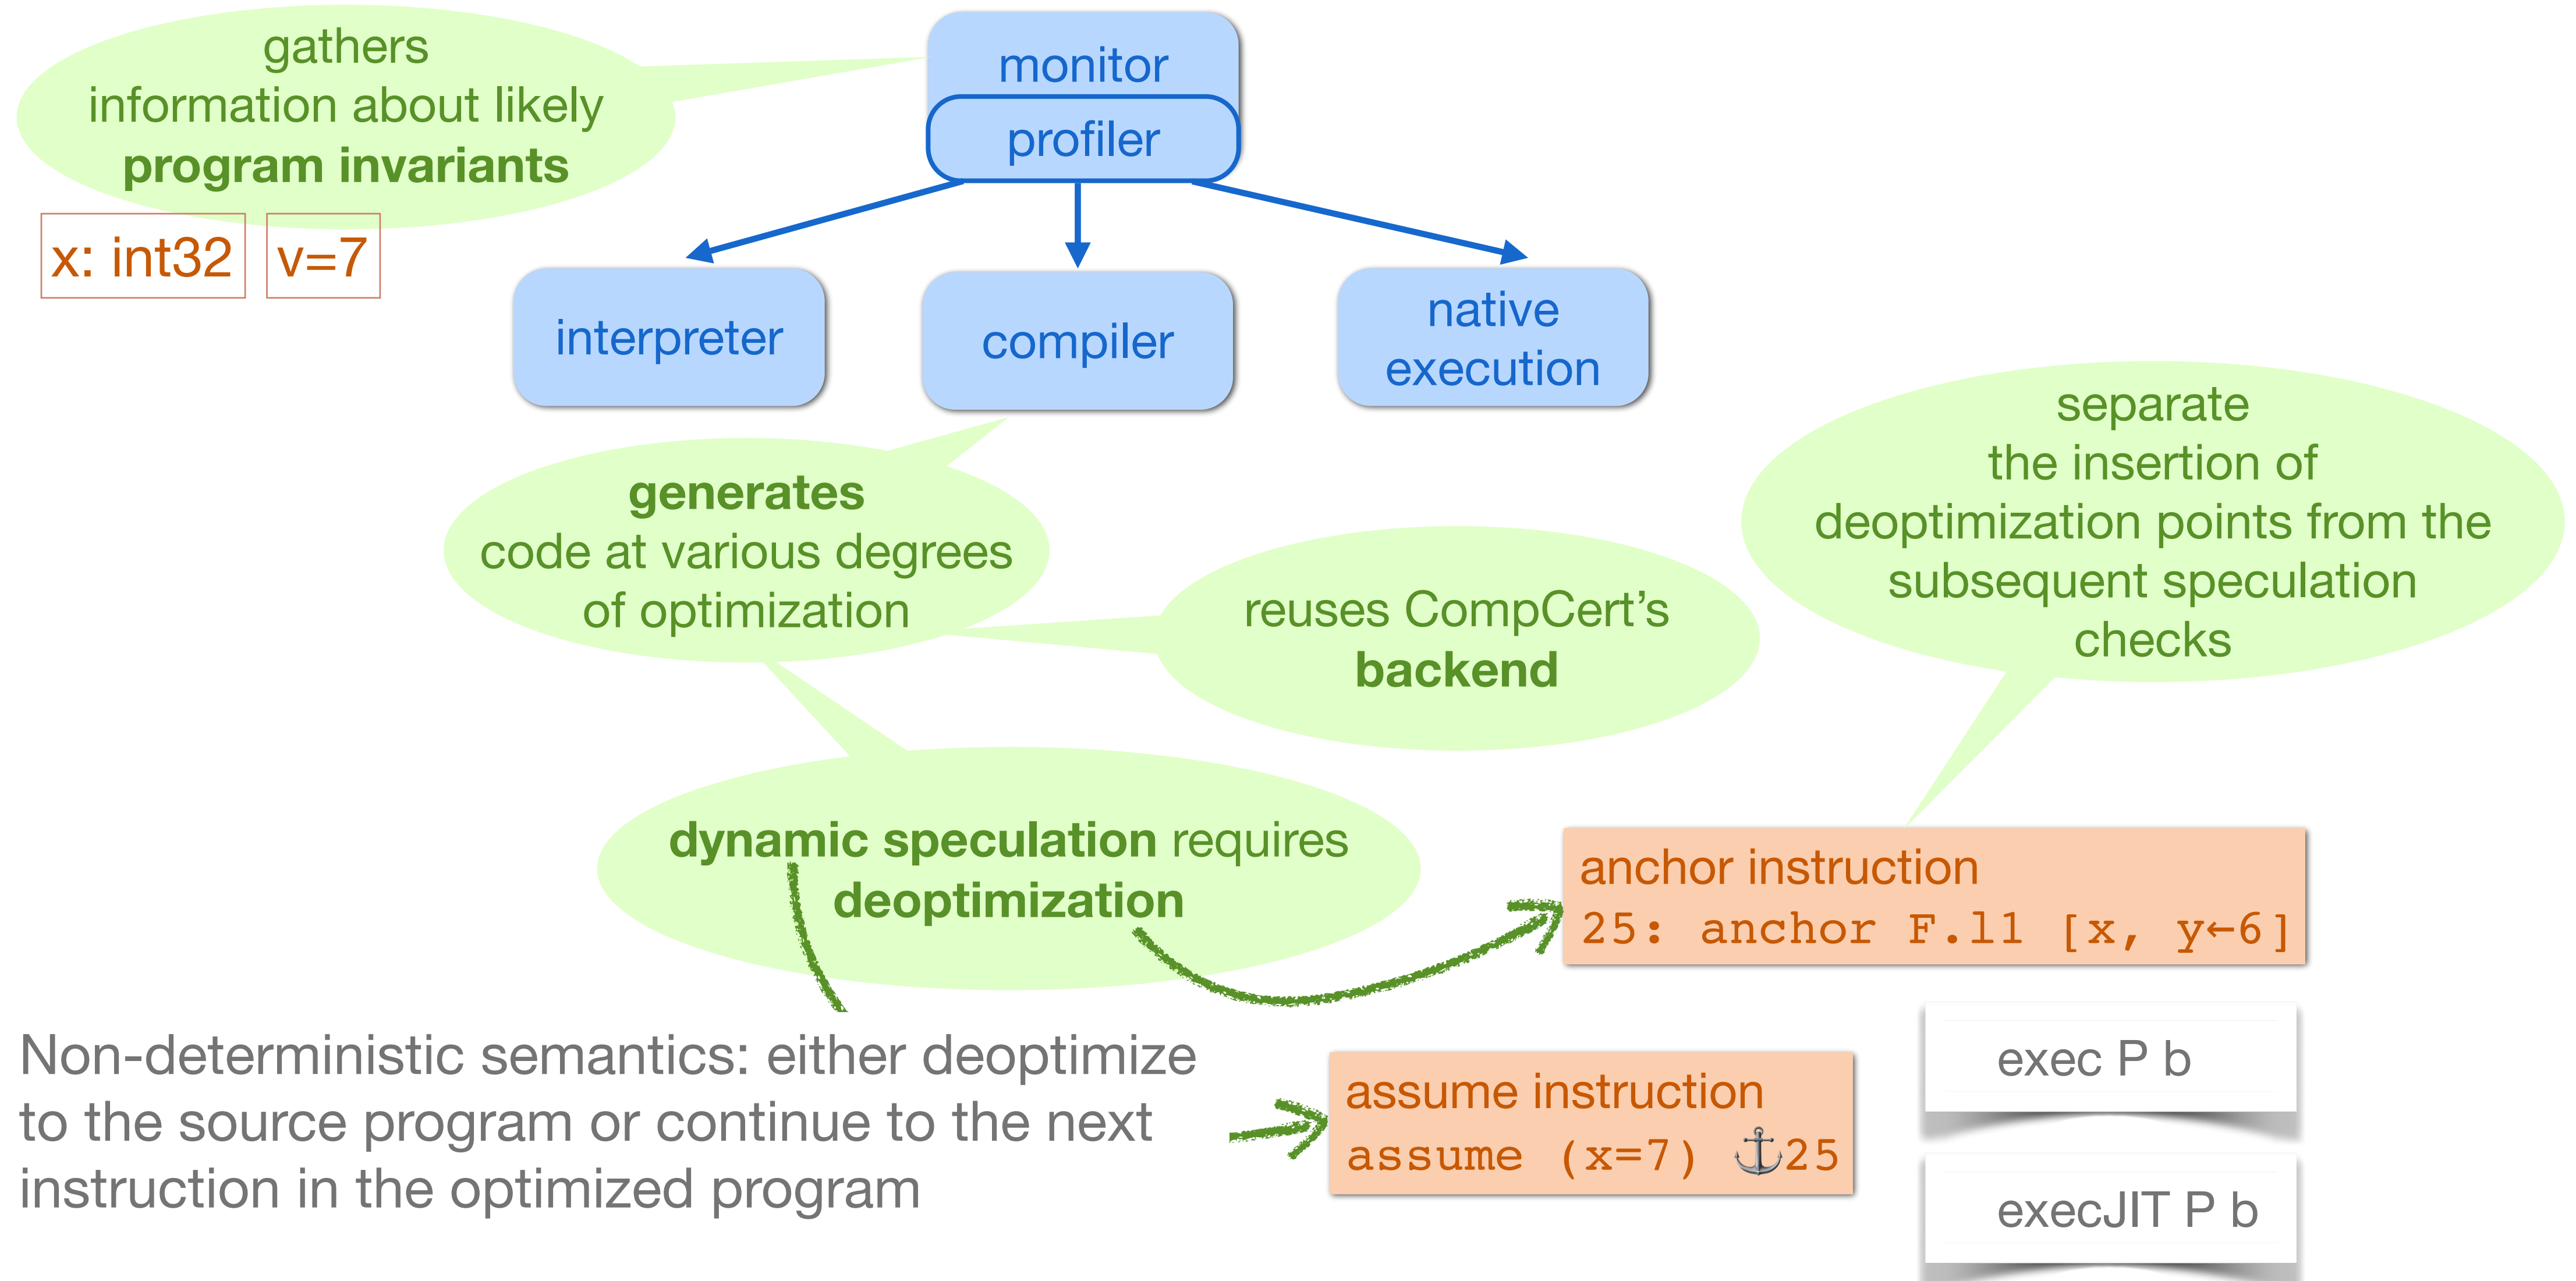
Dynamic speculation generates specialized functions

Deoptimization requires the JIT to synthesize interpreter stackframes in the middle of a function

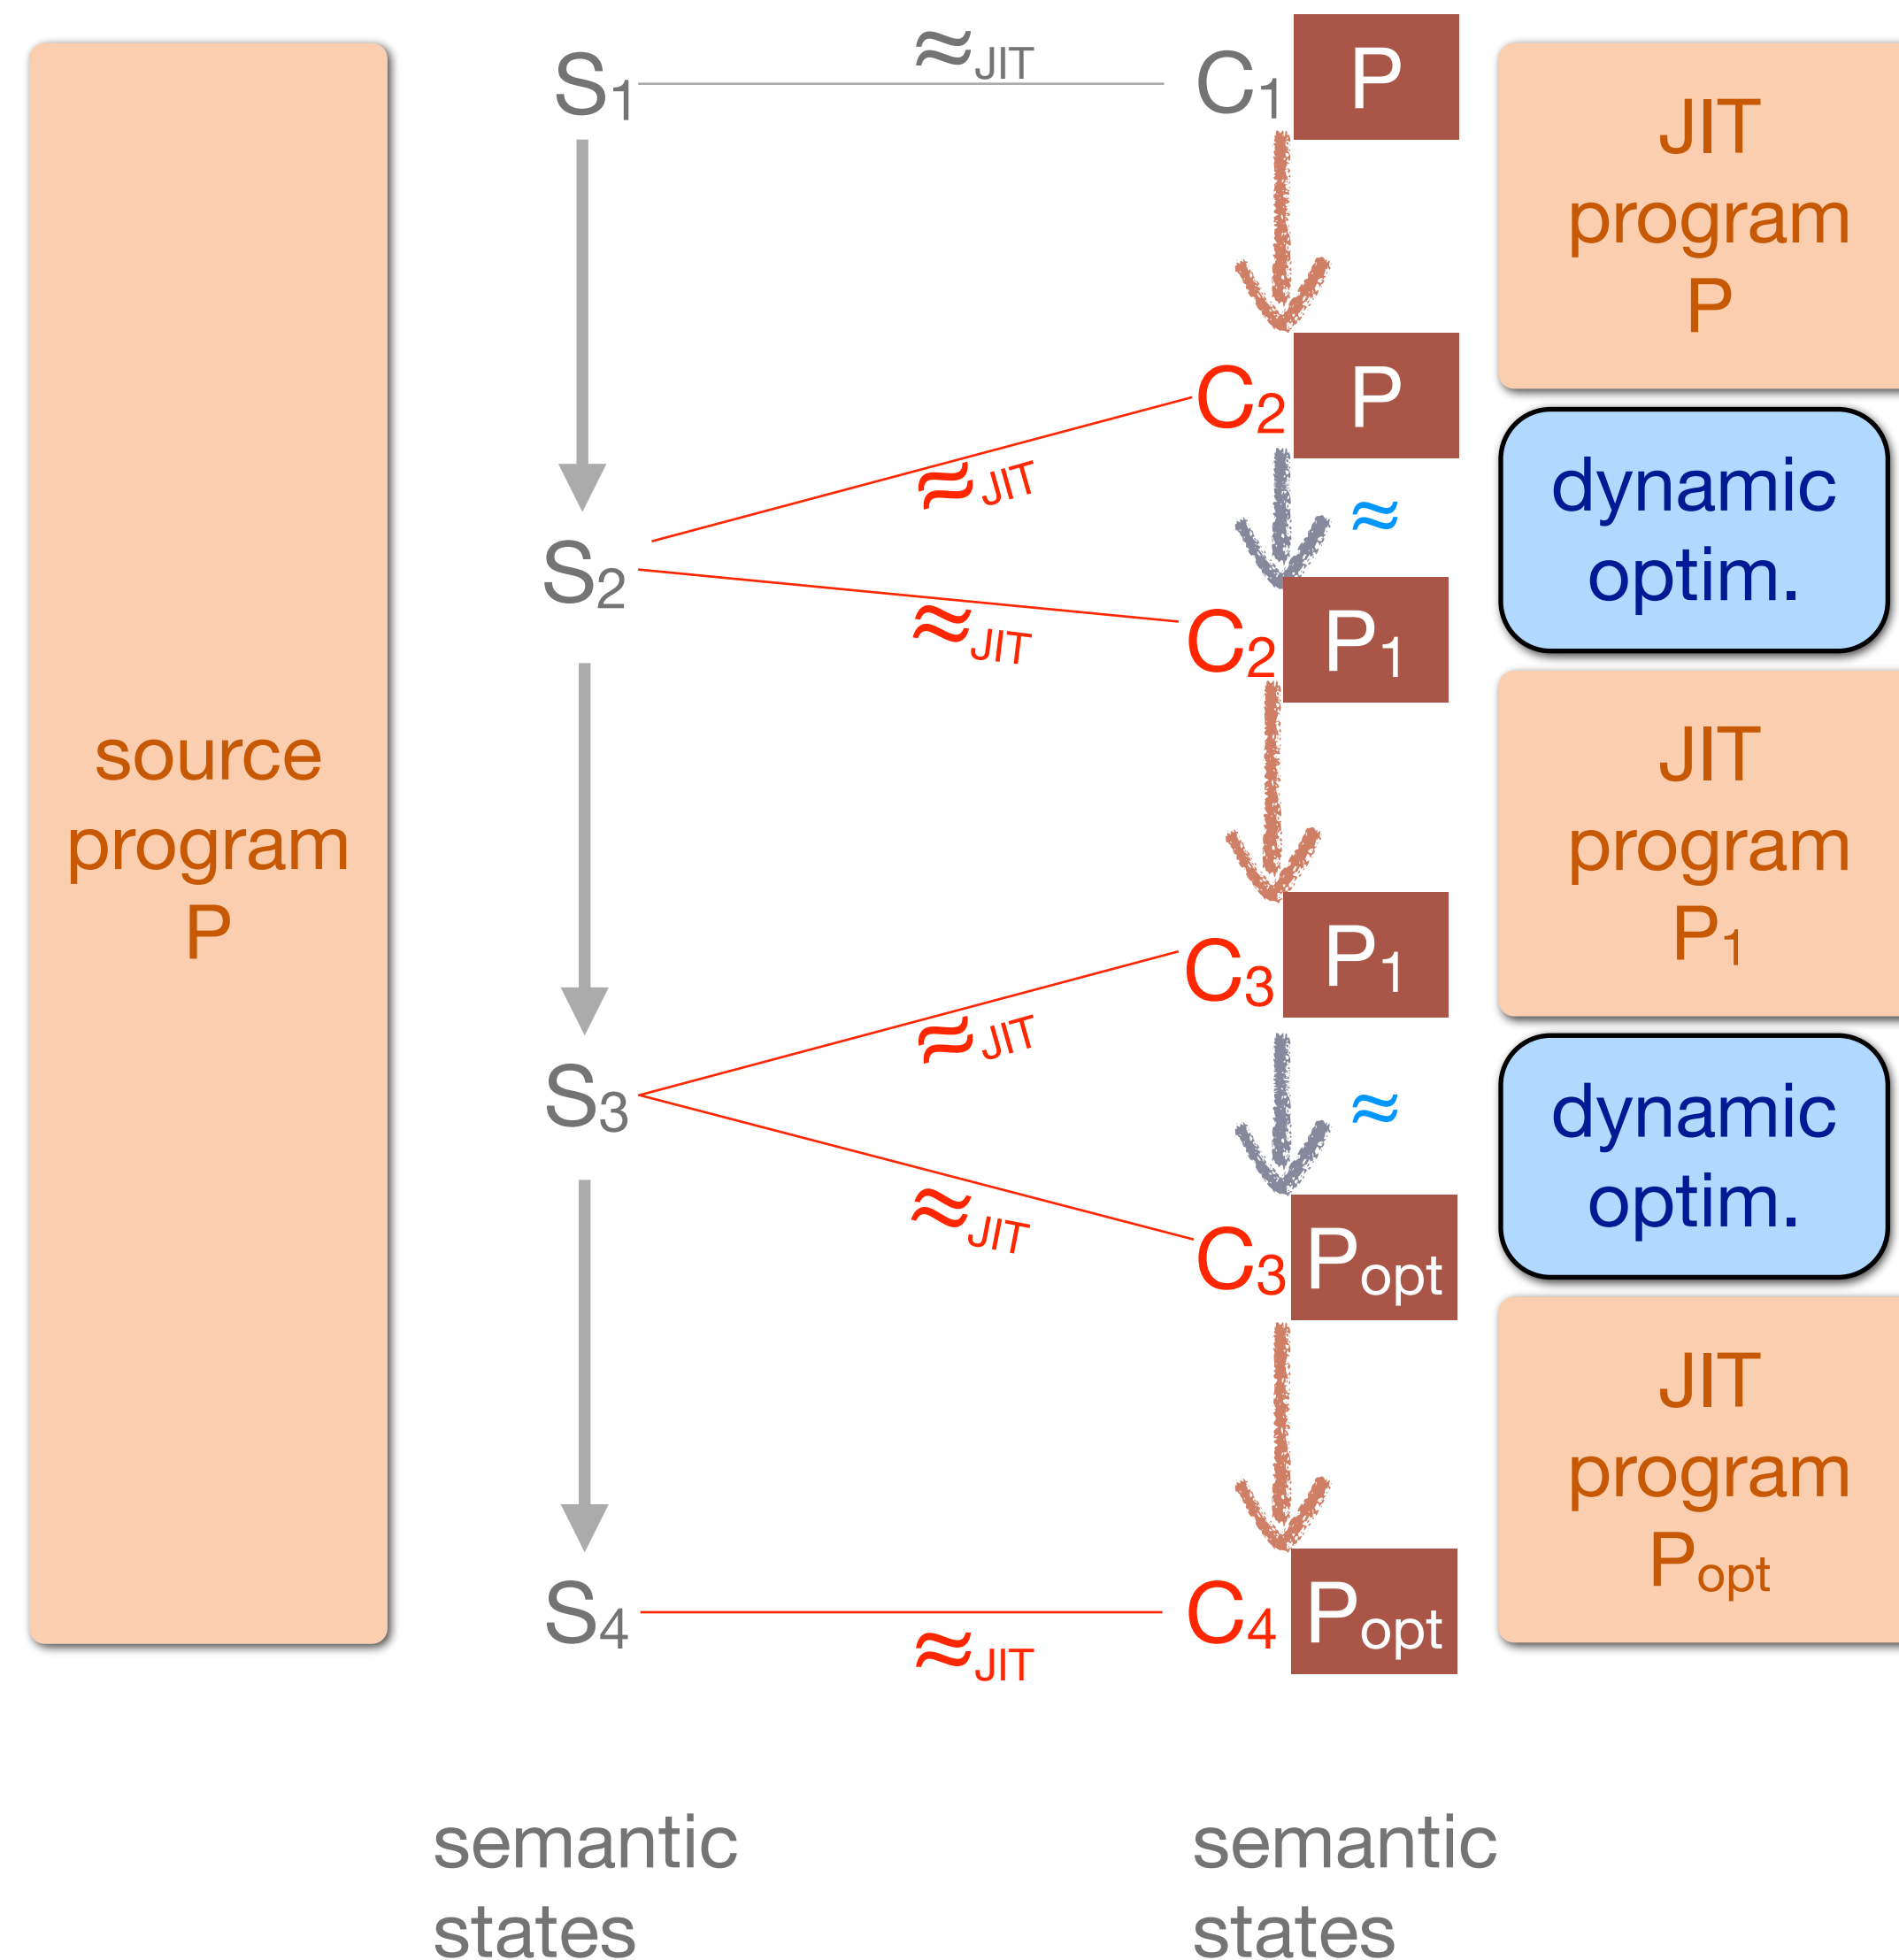
Verifying just-in-time (JIT) compilation: FM JIT

[Aurèle Barrière's PhD 12/2022]

[Barrière, Blazy, Flückiger, Pichardie, Vitek, POPL'21] [Barrière, Blazy, Pichardie, POPL'23]



Nested simulations for JIT verification



Theorem JITcompiler-correct:

$$\forall P P_{\text{opt}} b,$$

$$\text{JITcompiler } P = P_{\text{opt}} \rightarrow$$

$$\text{exec } P b \rightarrow$$


$$\text{execJIT } P_{\text{opt}} b.$$

Invariant \approx_{JIT} : at any point i during JIT execution

- C_i correspond to S_i
- P_i is **equivalent** to P

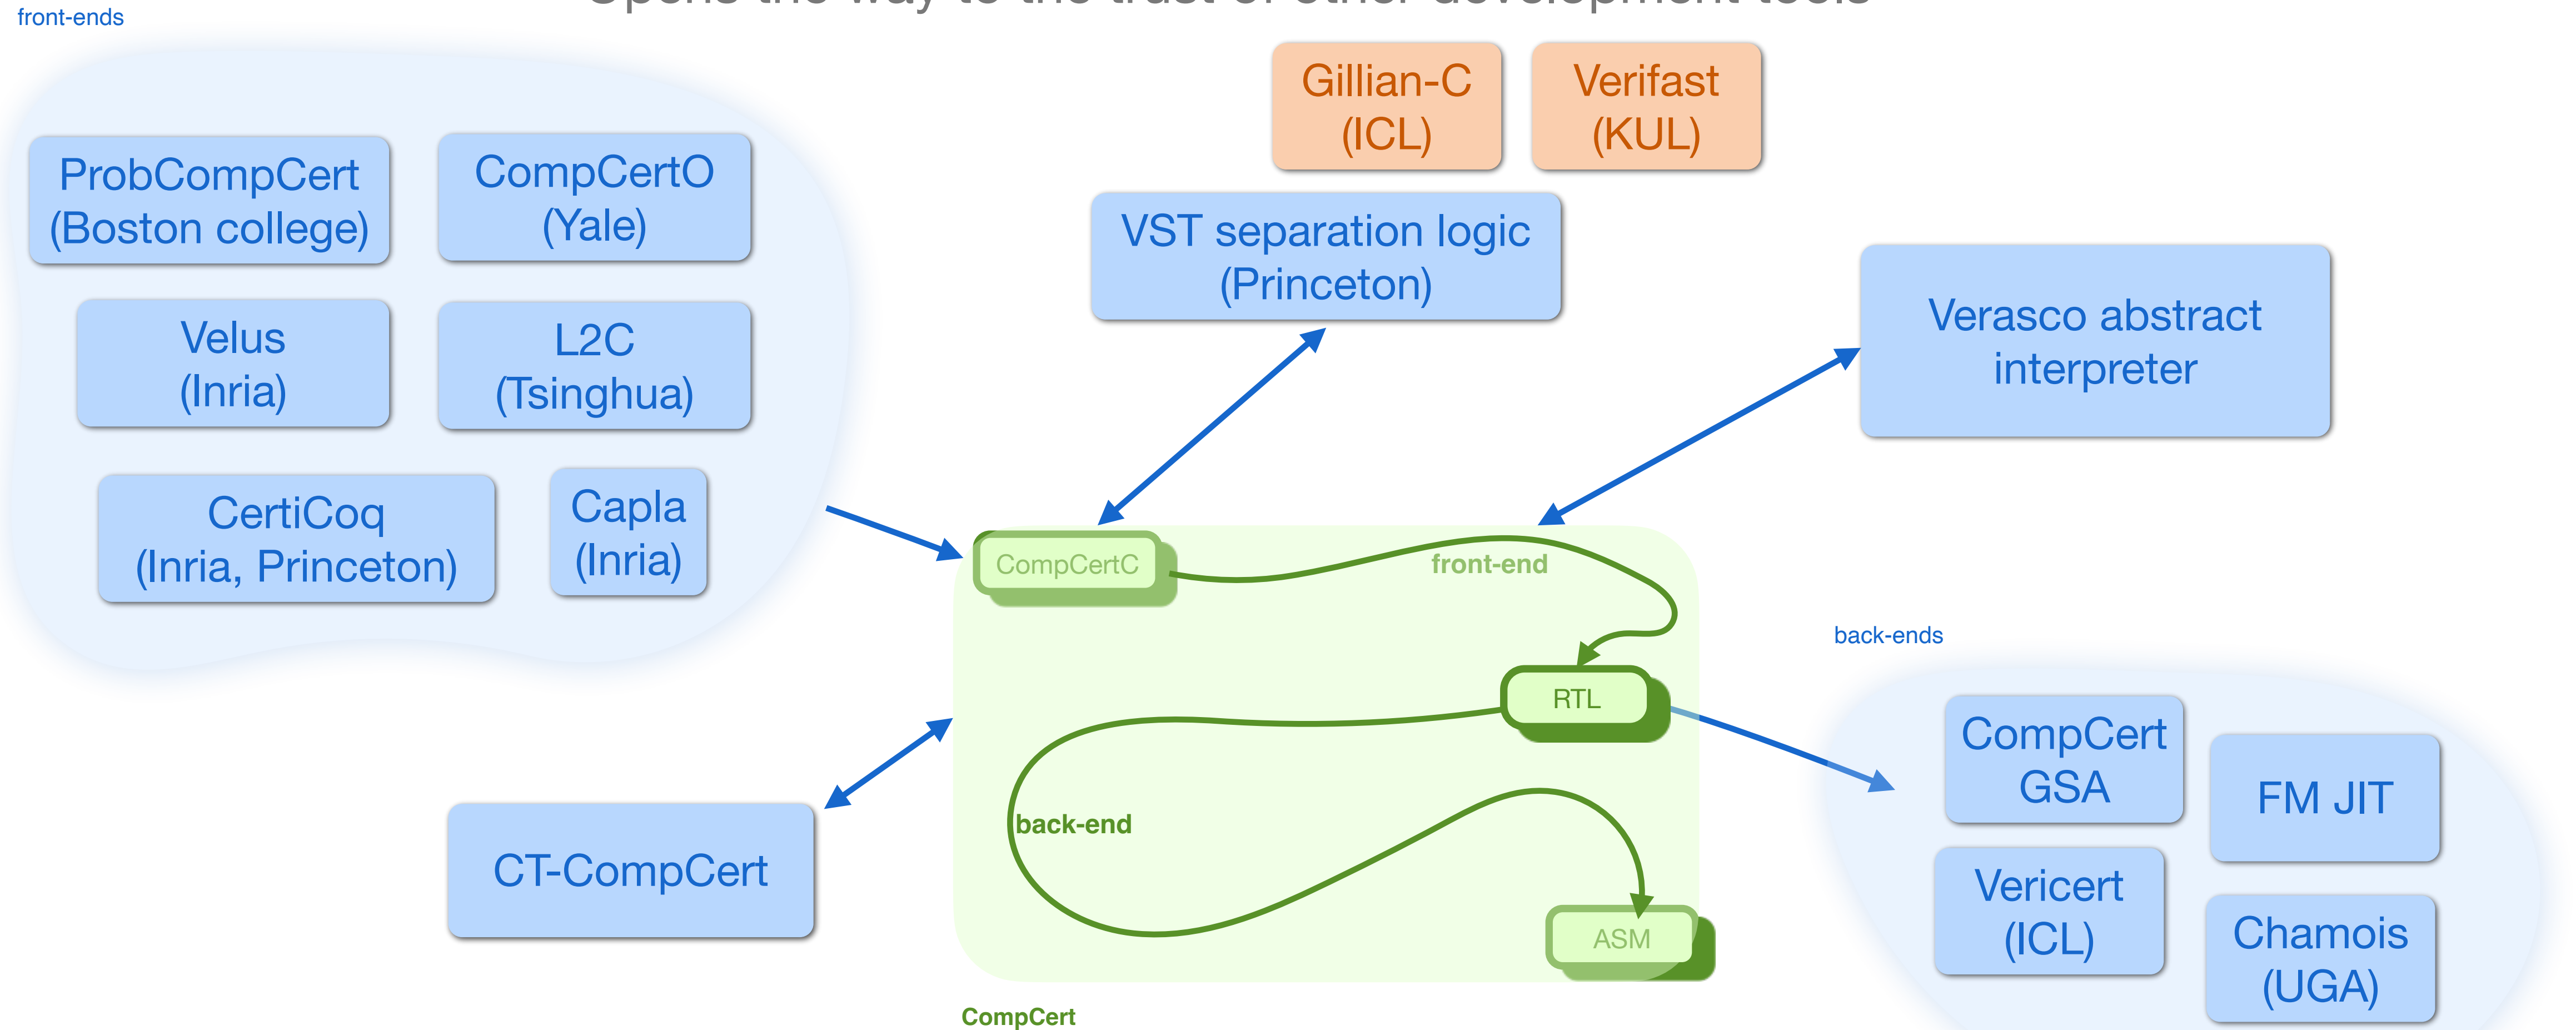
Nested simulation: this **equivalence** is expressed with another (internal) simulation \approx between compiled programs

Conclusion



CompCert, an open infrastructure for research

Opens the way to the trust of other development tools



Mechanized semantics are the shared basis for verified compilers, sound program logics, and sound static analyzers

Thank you!

Questions?

