


Formal Verification of a Constant-Time Preserving C Compiler

Gilles Barthe, Sandrine Blazy, Benjamin Grégoire, Rémi Hutin,
Vincent Laporte, David Pichardie and Alix Trieu



Cache timing attacks against cryptographic implementations

- Common side-channel: cache timing attacks
- Exploit the latency between cache hits and misses
- Attackers can recover cryptographic keys
 - Tromer et al (2010), Gullasch et al (2011) show efficient attacks on AES implementations
- Based on the use of look-up tables
 - Access to memory addresses that depend on the key



Constant-time programming

A programming discipline for crypto programmers

- Constant-time programs should not
 - branch on secrets
 - perform memory accesses that depend on secrets
- This is a strictly stronger property than « time execution does not depend on secrets » !
- There are constant-time implementations of many cryptographic algorithms: AES, DES, RSA, etc

```
if (secret)  
then do1()  
else do2()
```



not constant-time

```
a[secret] = ...
```



not constant-time


Cryptographic constant-time verification

- Several verification tools have been built and used for checking that popular libraries are constant-time [Almeida16, Rodrigues16]
- But checking low-level implementations is not ideal
 - it makes the analysis work harder (e.g. alias analysis)
 - it makes the results of the analysis difficult to understand for programmers

Our Research Program

Our Research Program

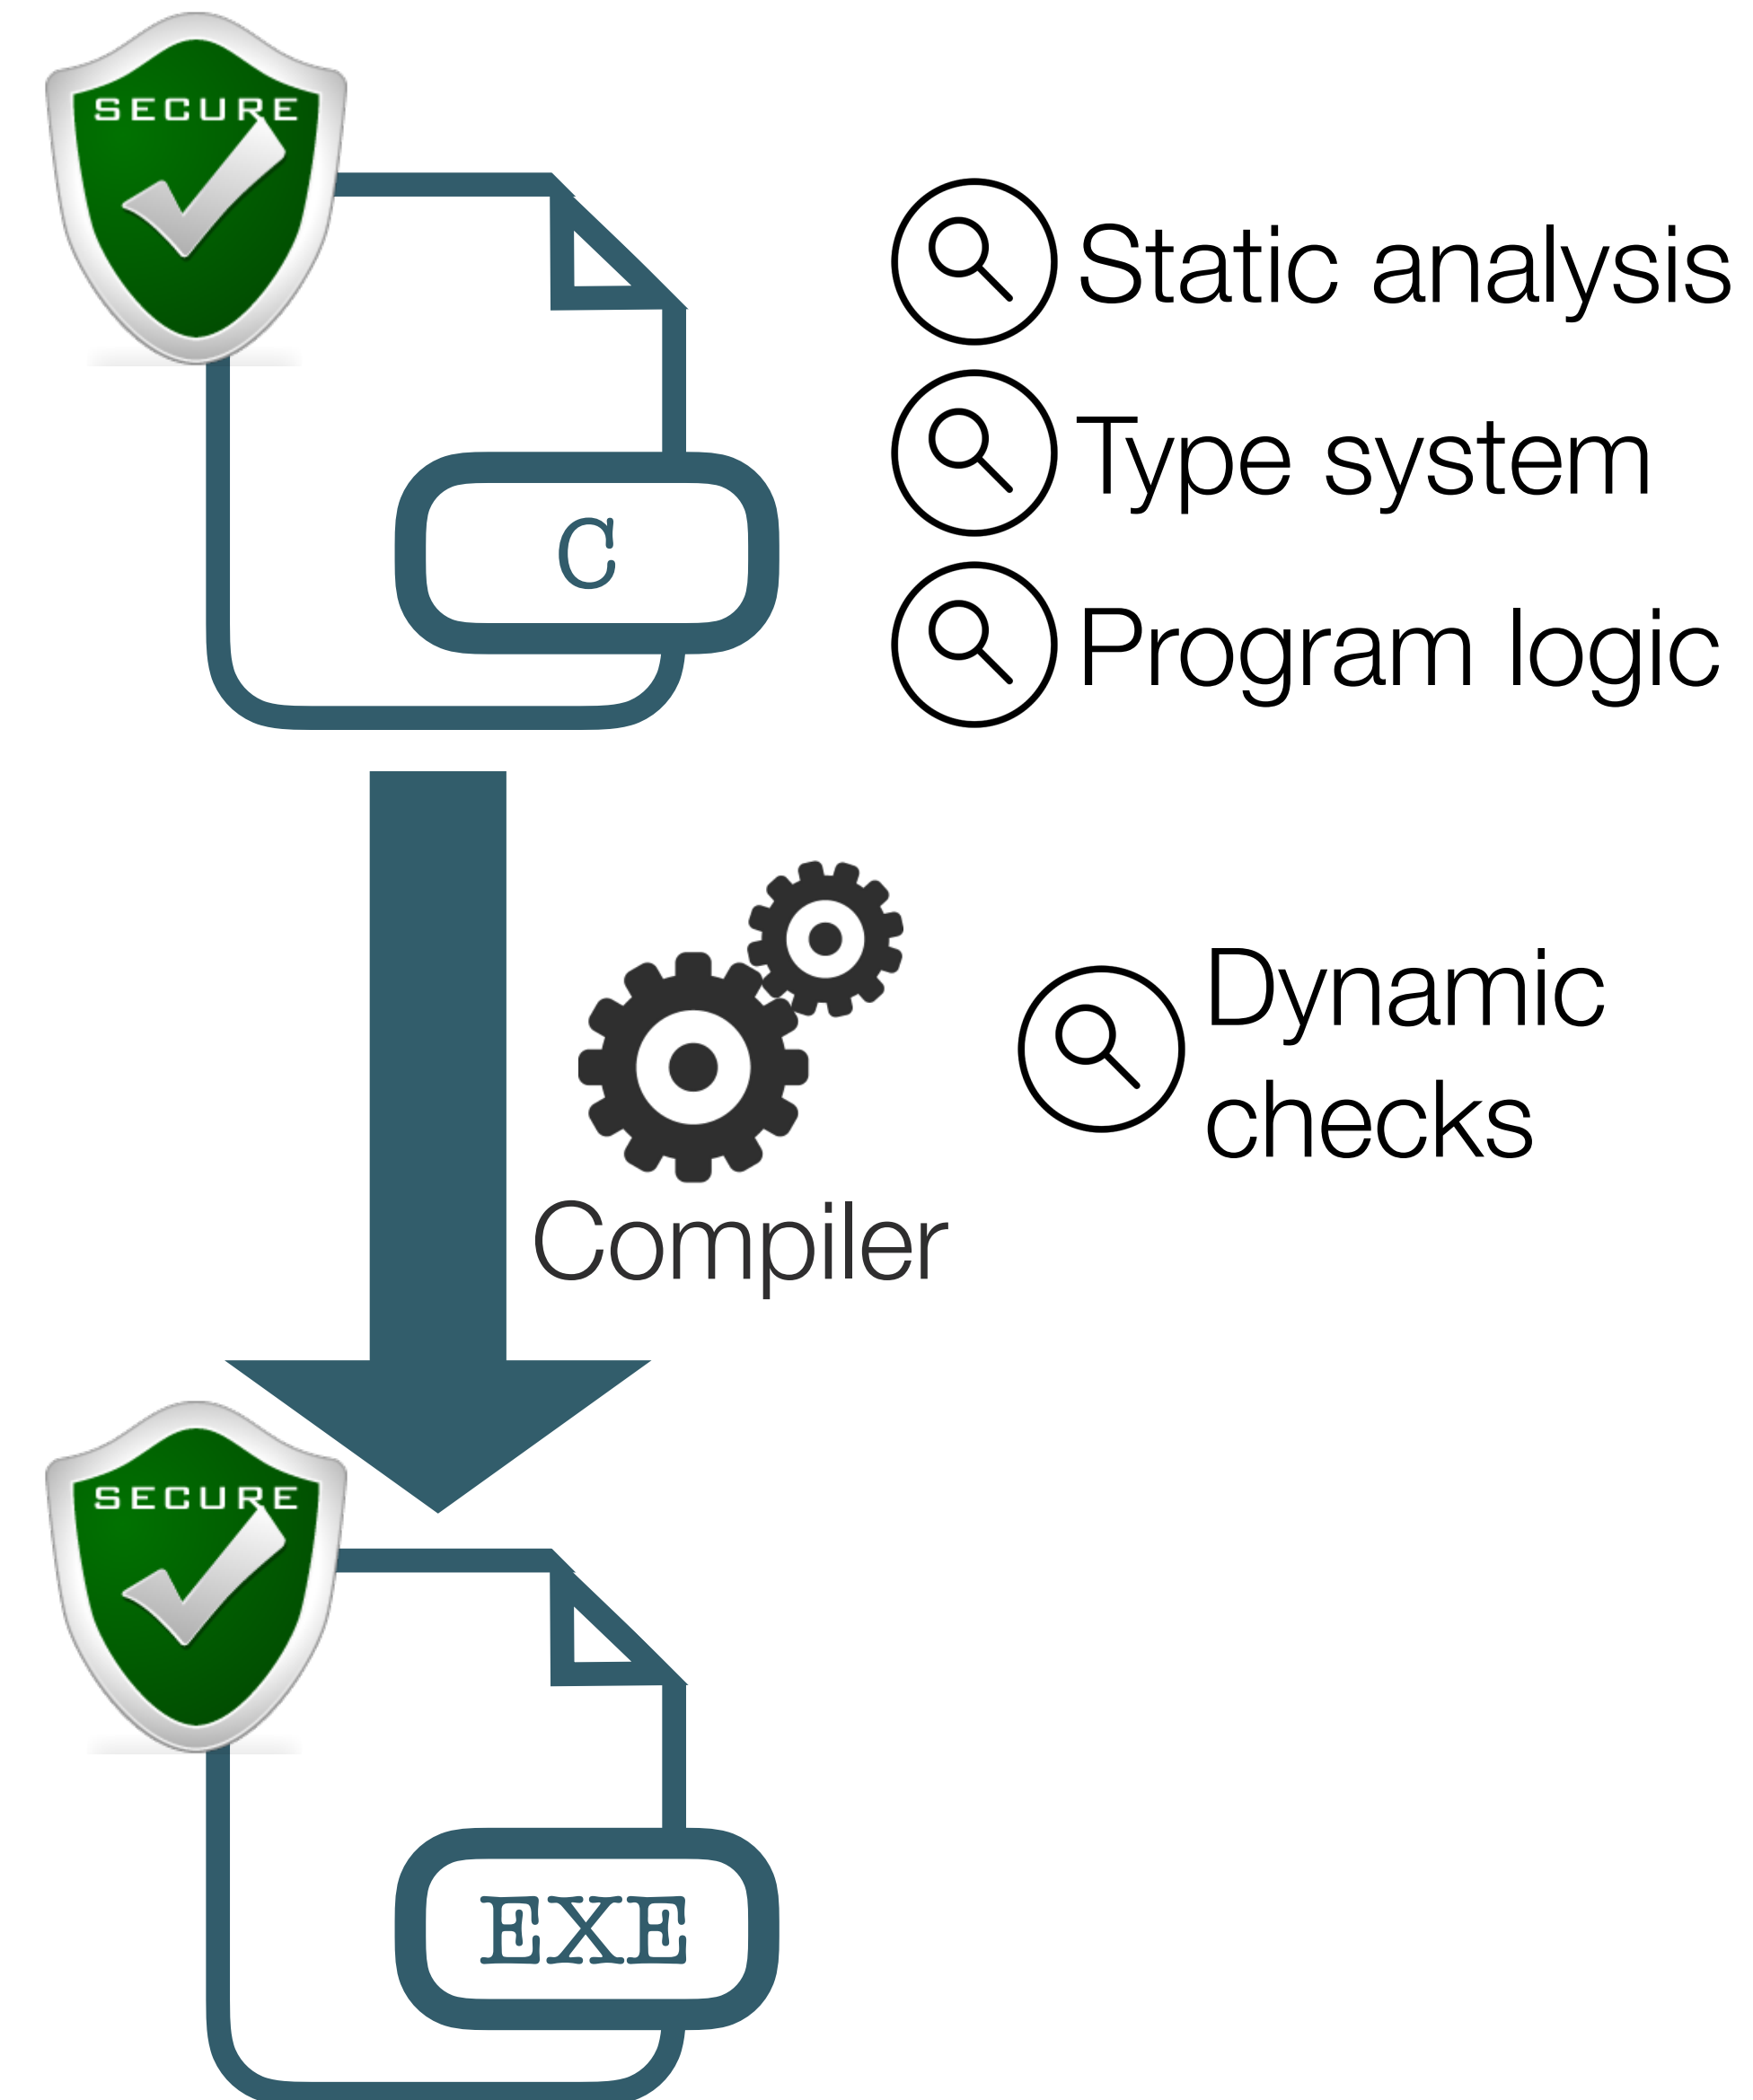
- Build secure programming abstractions at source level (C-like)



- 🔍 Static analysis
- 🔍 Type system
- 🔍 Program logic

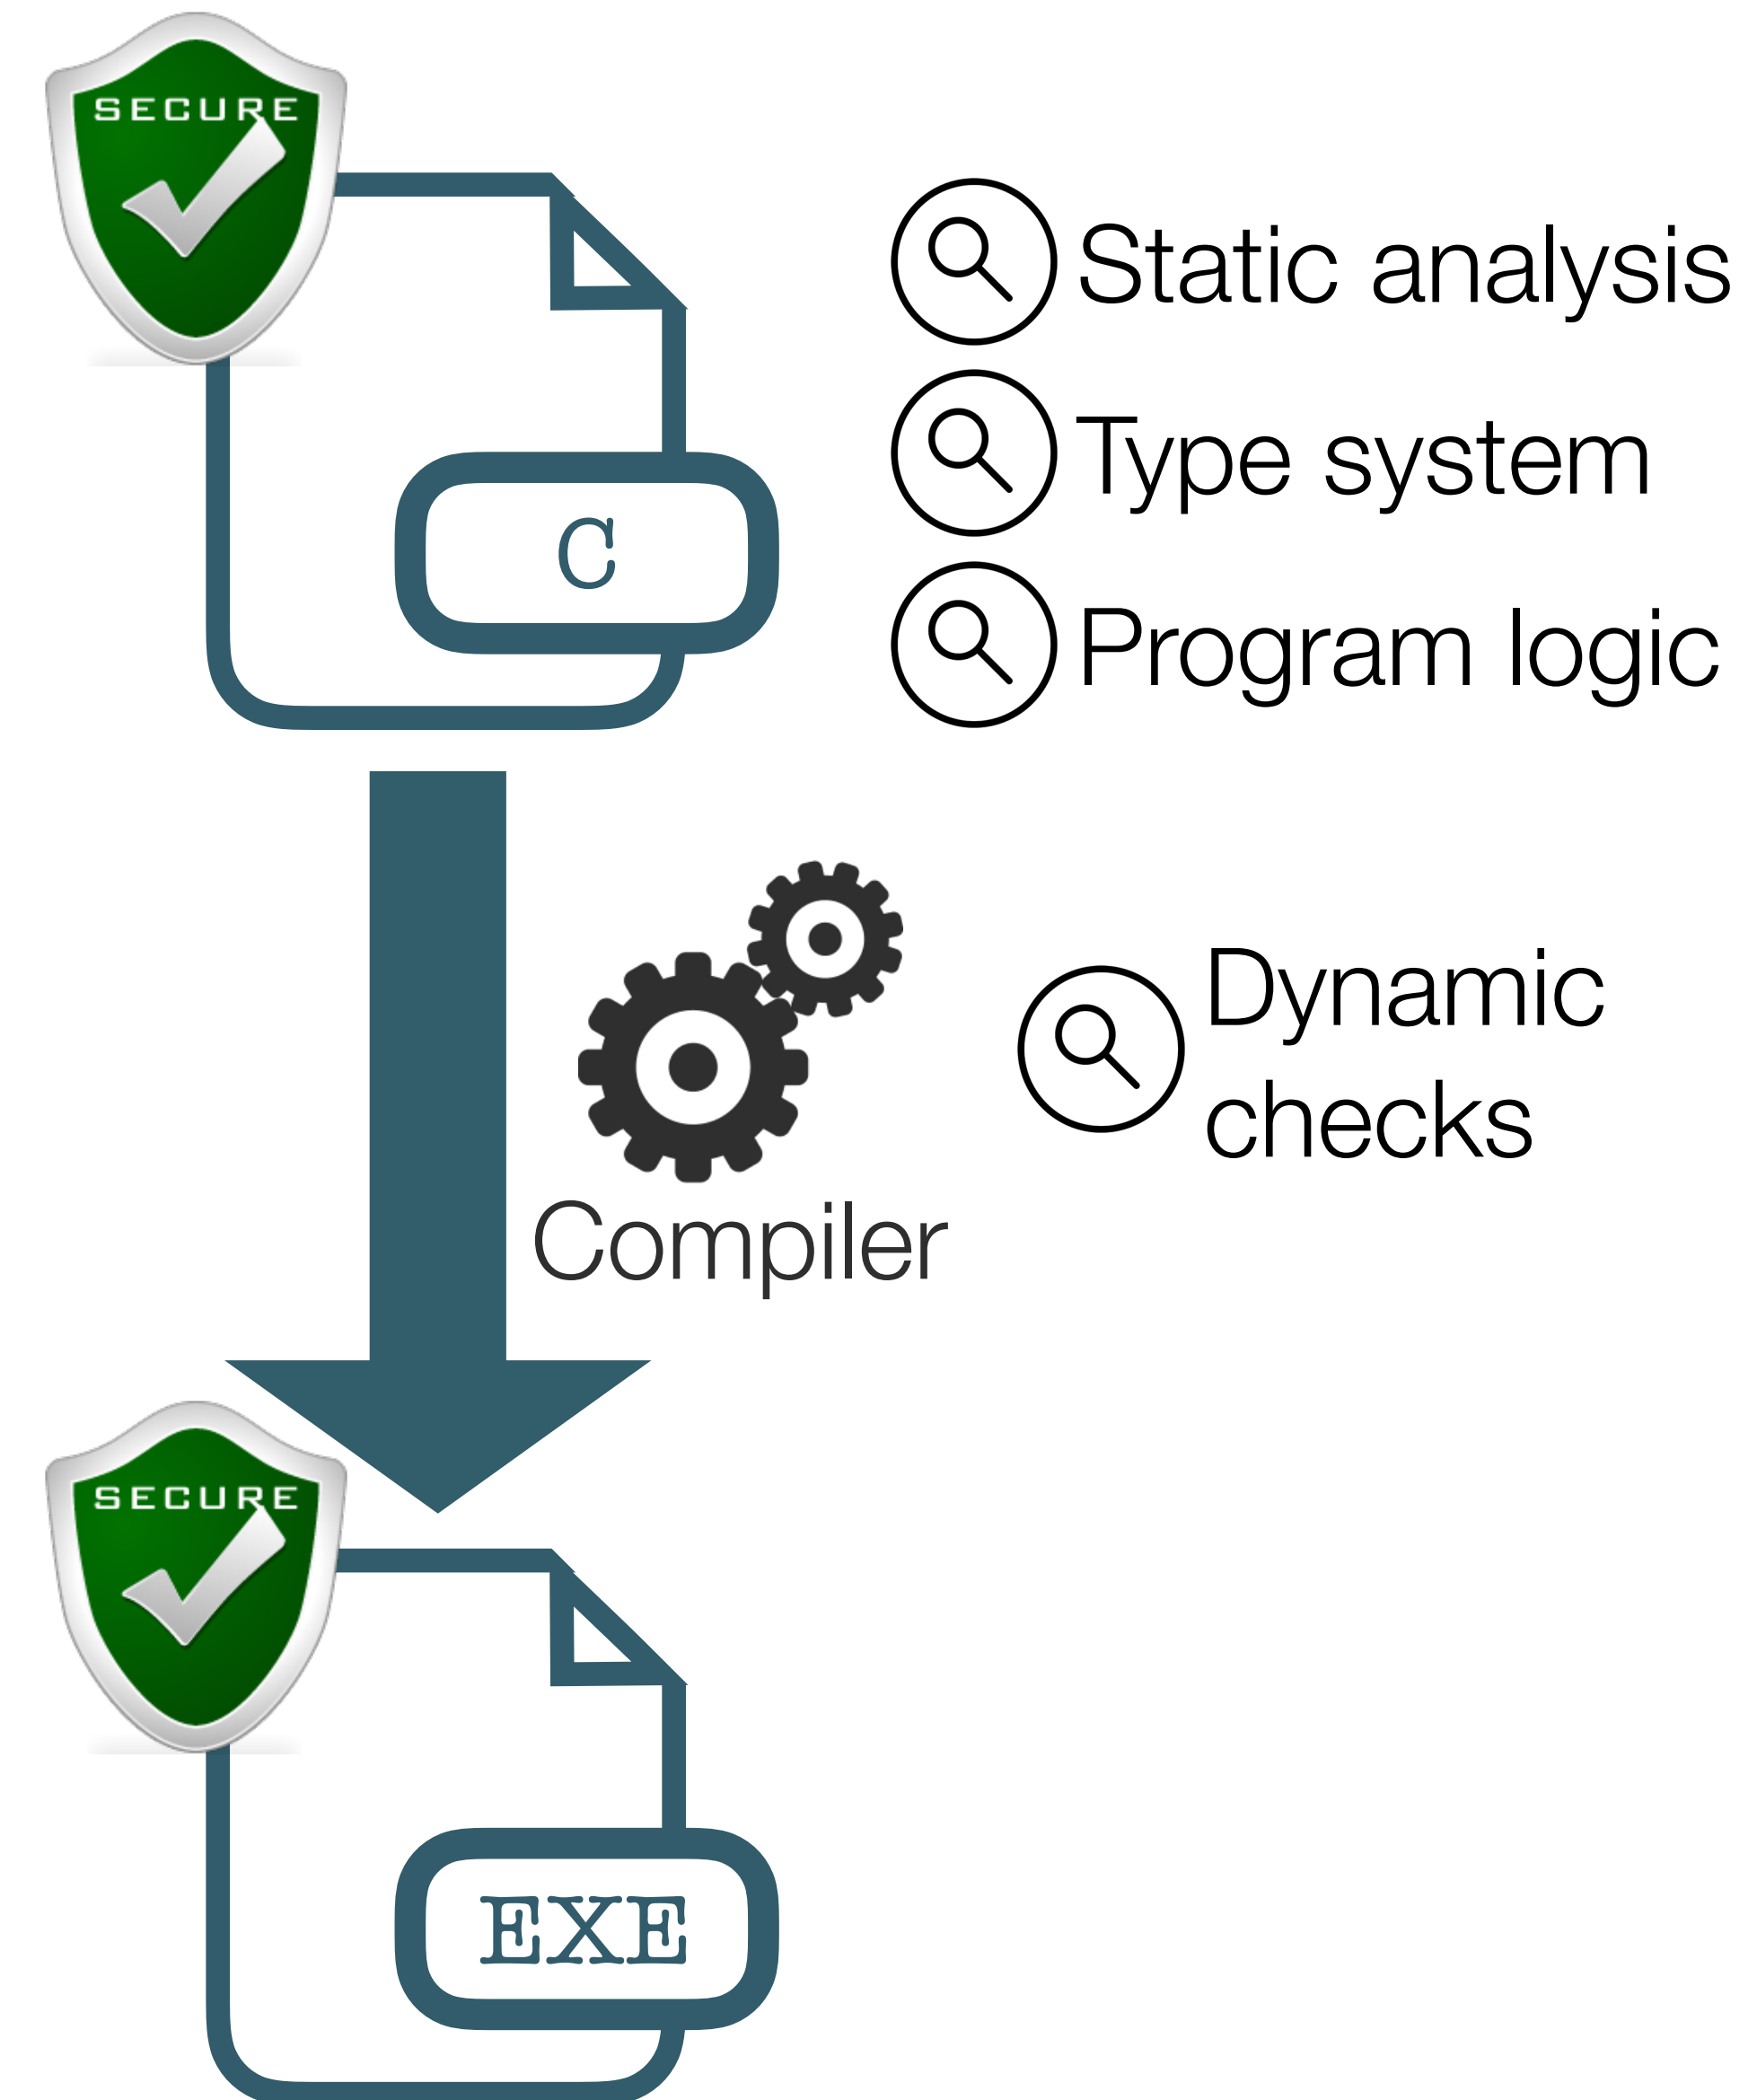
Our Research Program

- Build secure programming abstractions at source level (C-like)
- Make sure the compiler will generate executables that are as secure



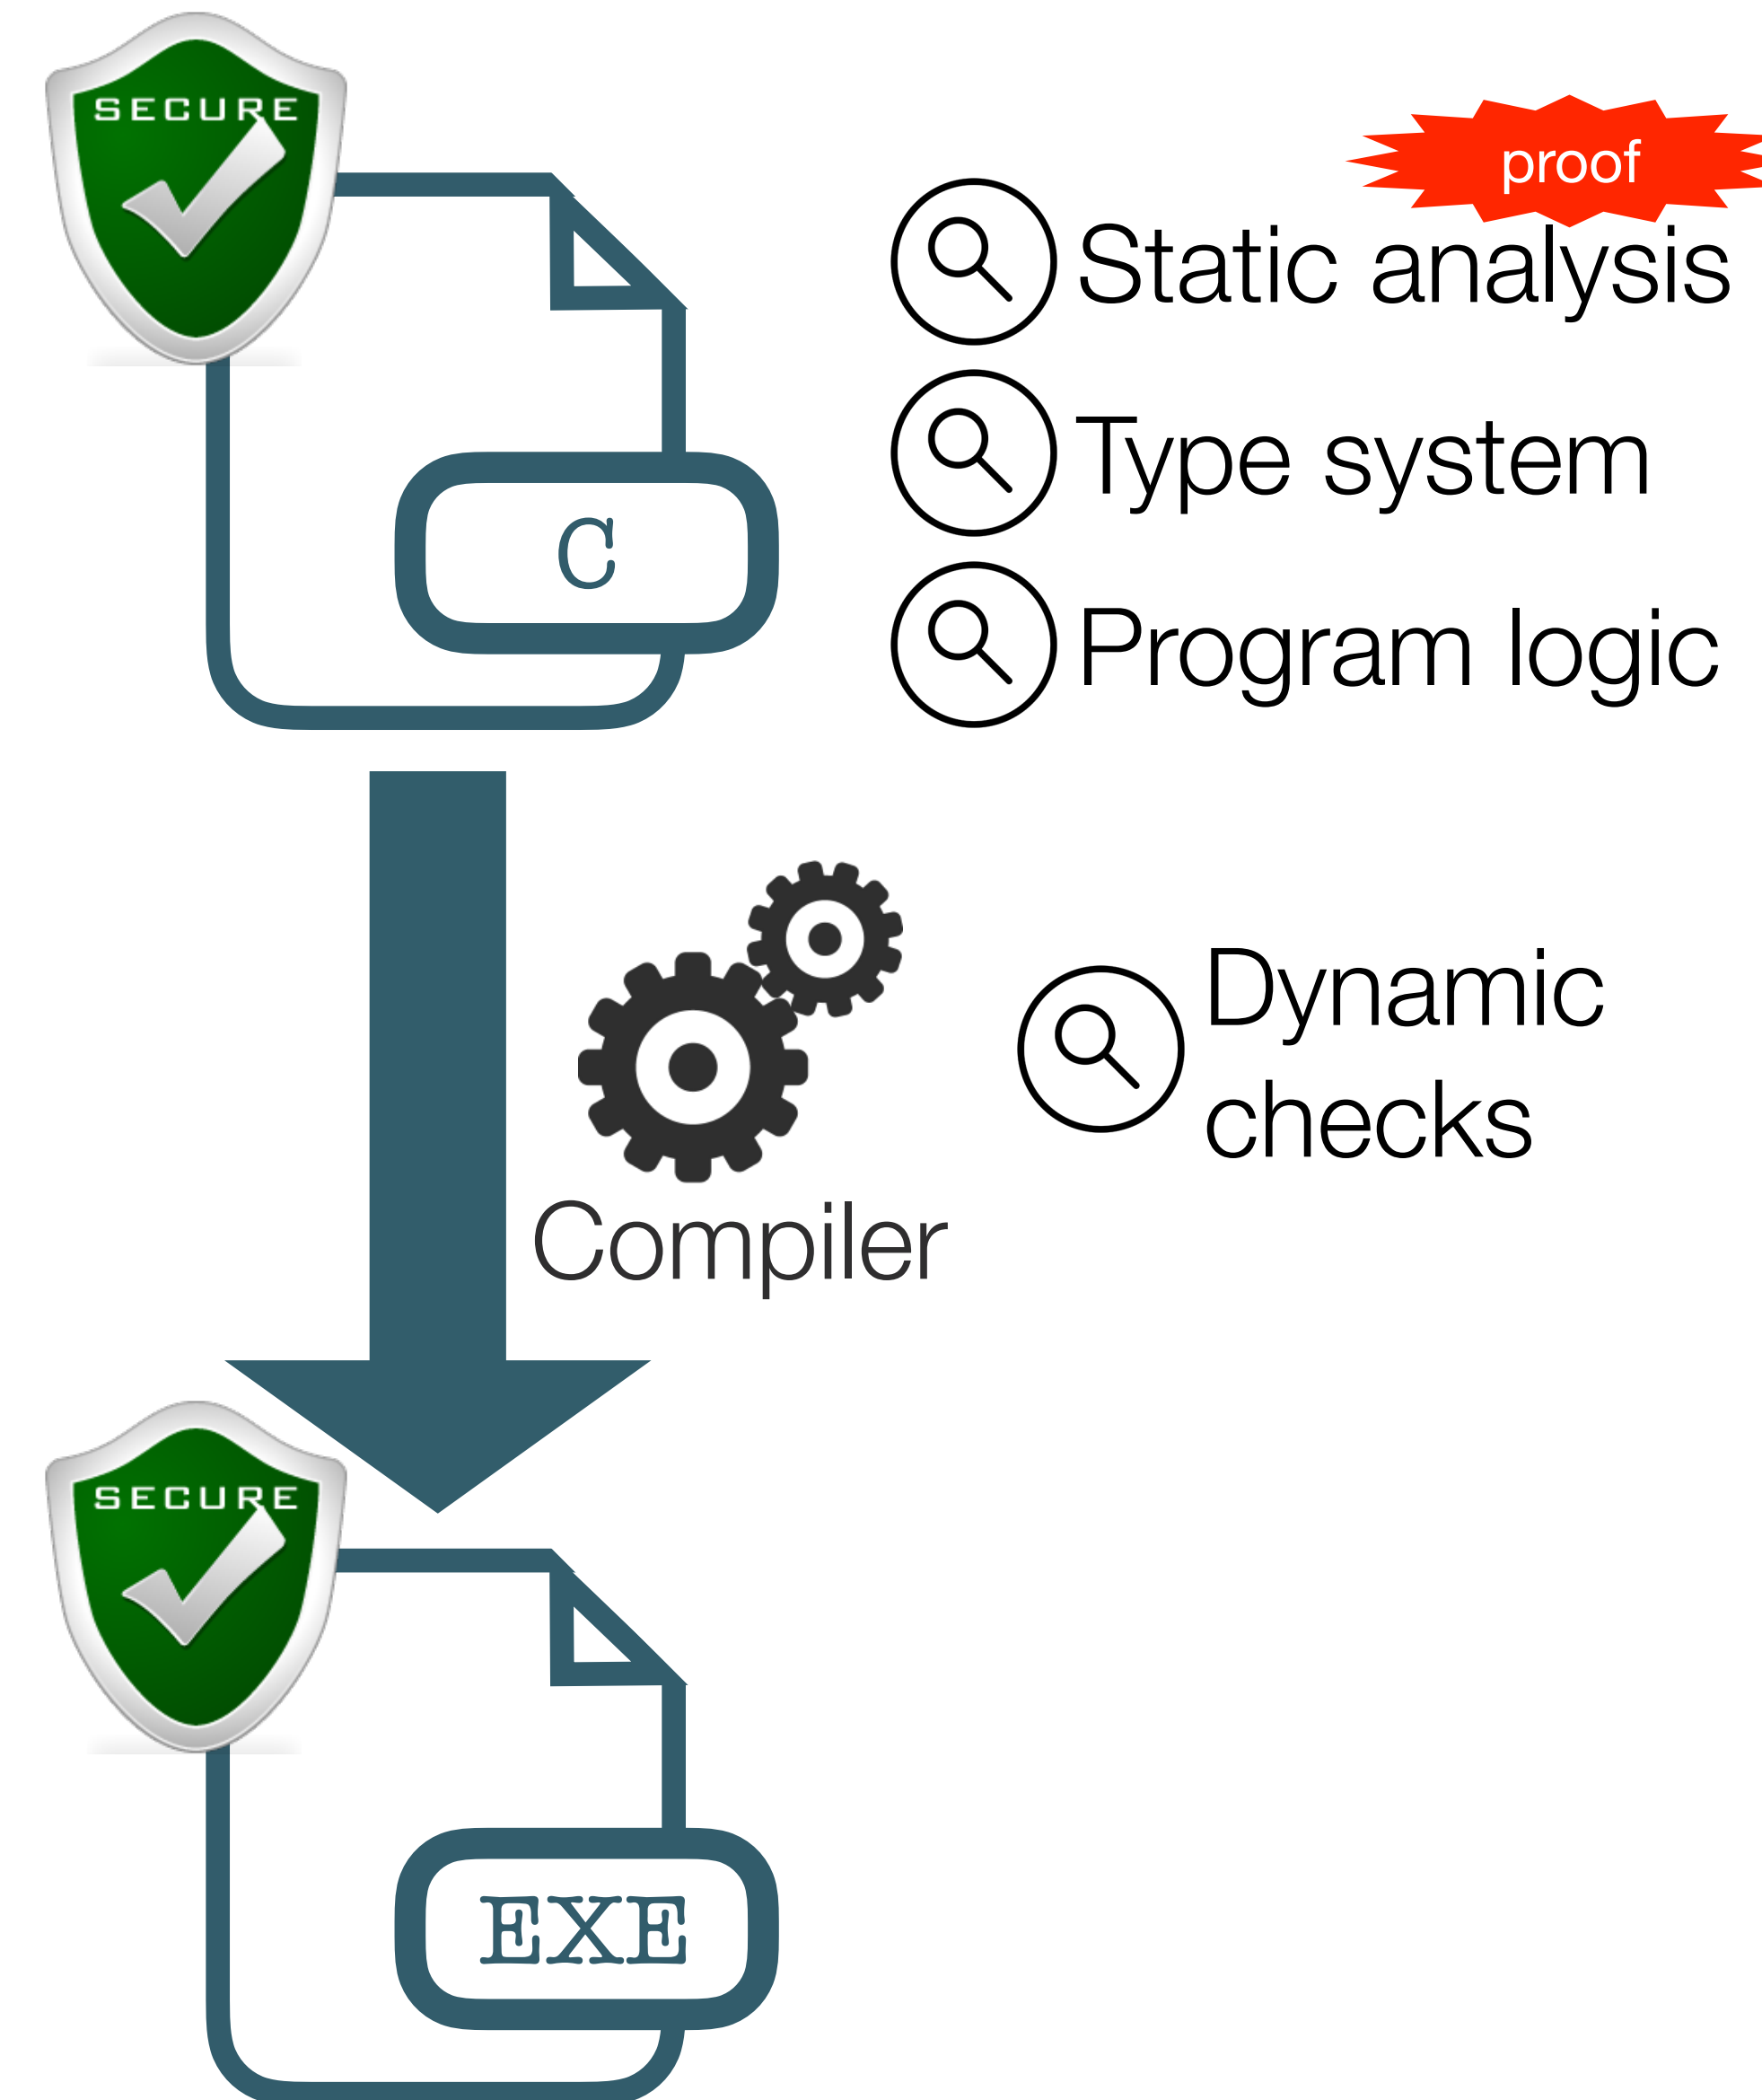
Our Research Program

- Build secure programming abstractions at source level (C-like)
- Make sure the compiler will generate executables that are as secure
- Reduce as much as possible the TCB (Trusted Computing Base) with formal proofs



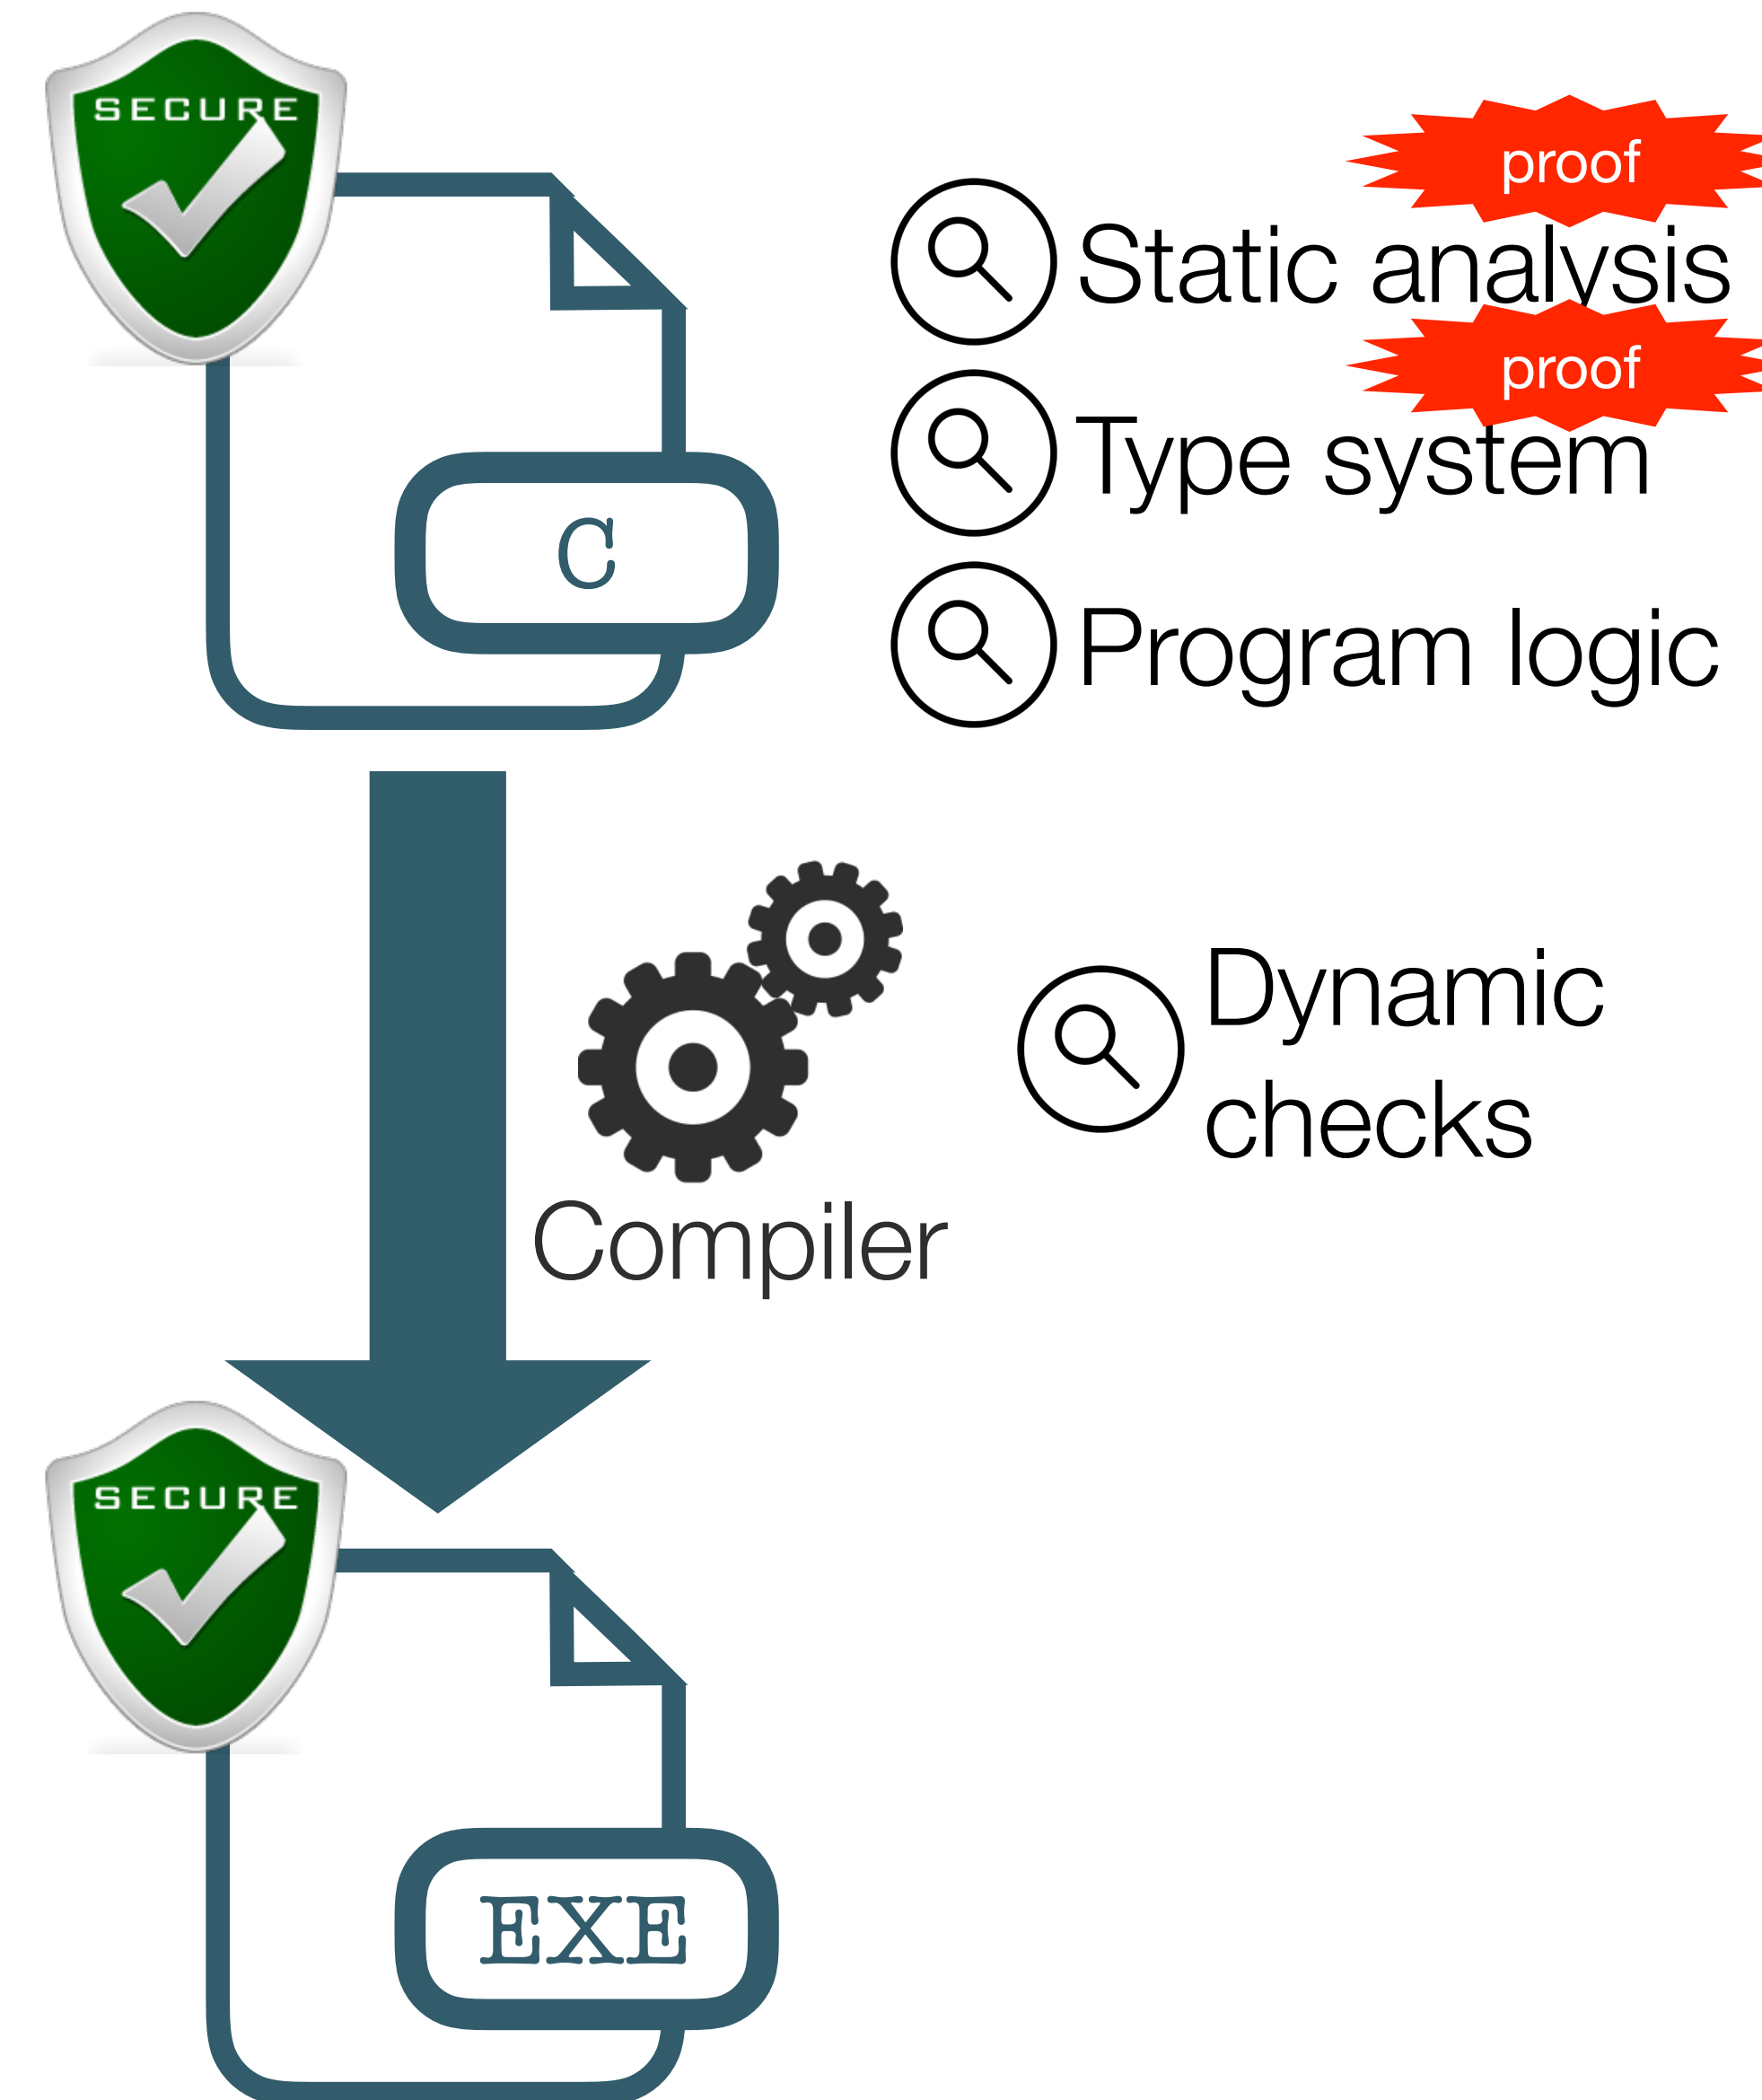
Our Research Program

- Build secure programming abstractions at source level (C-like)
- Make sure the compiler will generate executables that are as secure
- Reduce as much as possible the TCB (Trusted Computing Base) with formal proofs



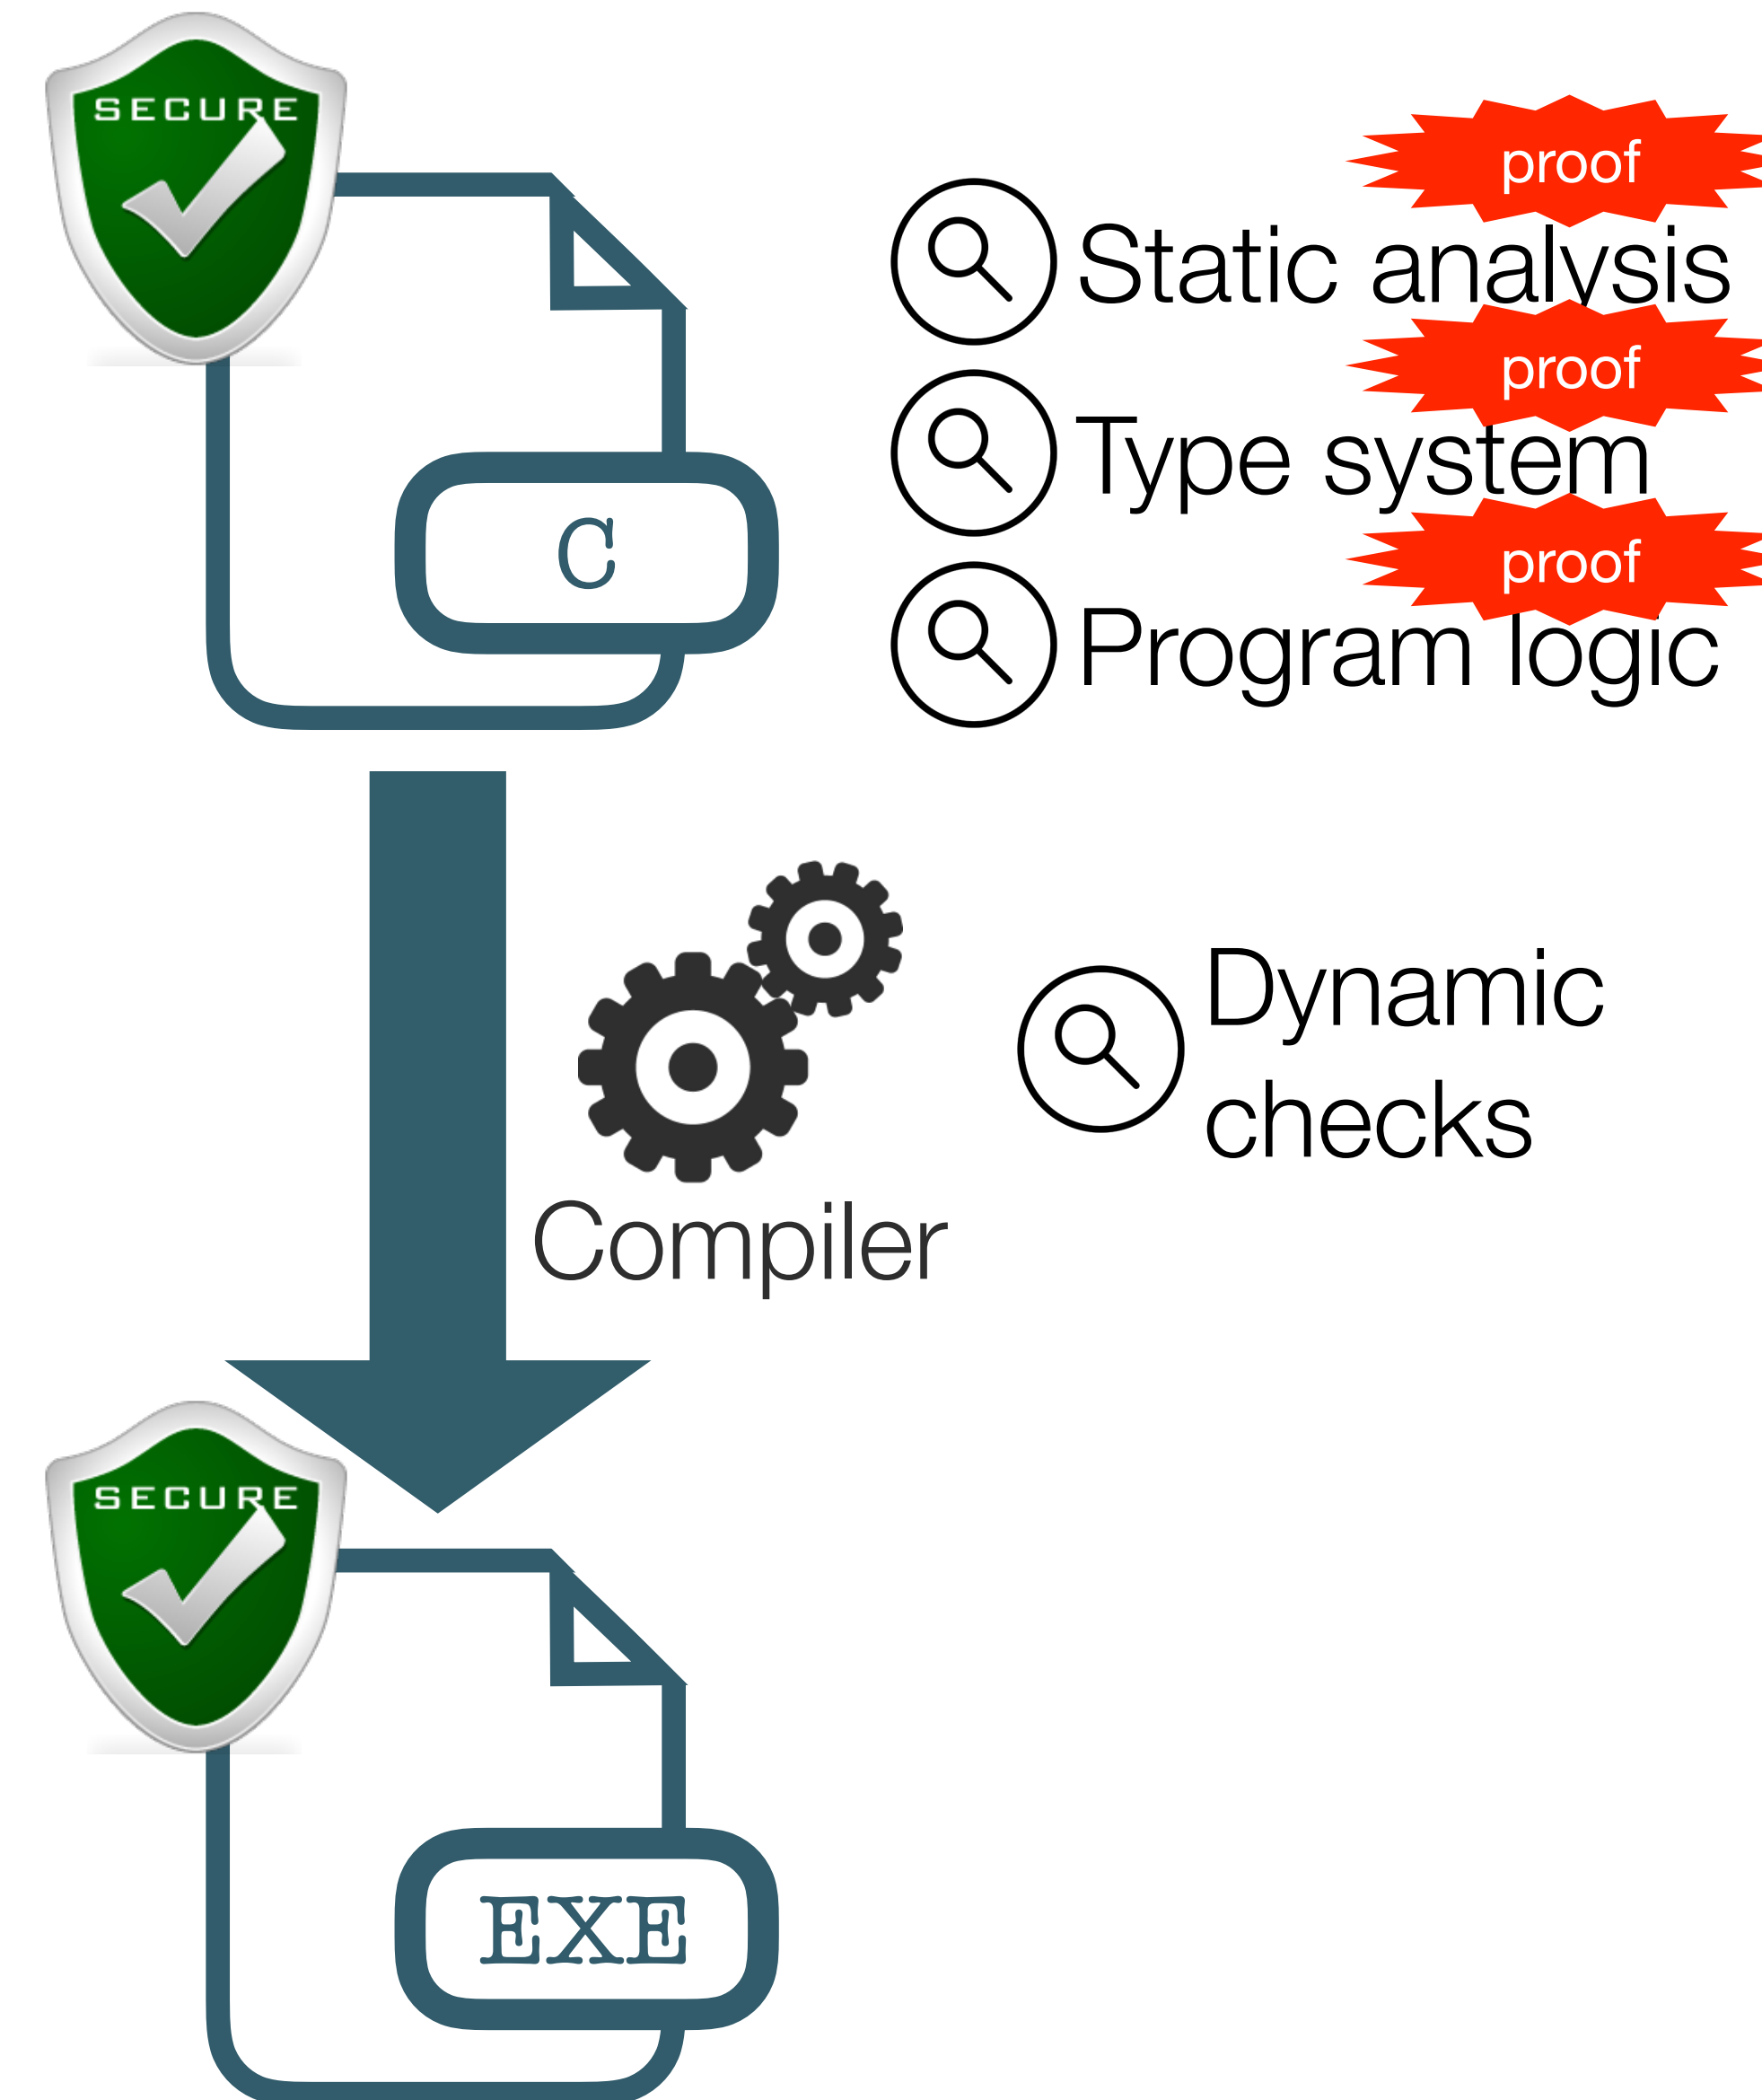
Our Research Program

- Build secure programming abstractions at source level (C-like)
- Make sure the compiler will generate executables that are as secure
- Reduce as much as possible the TCB (Trusted Computing Base) with formal proofs



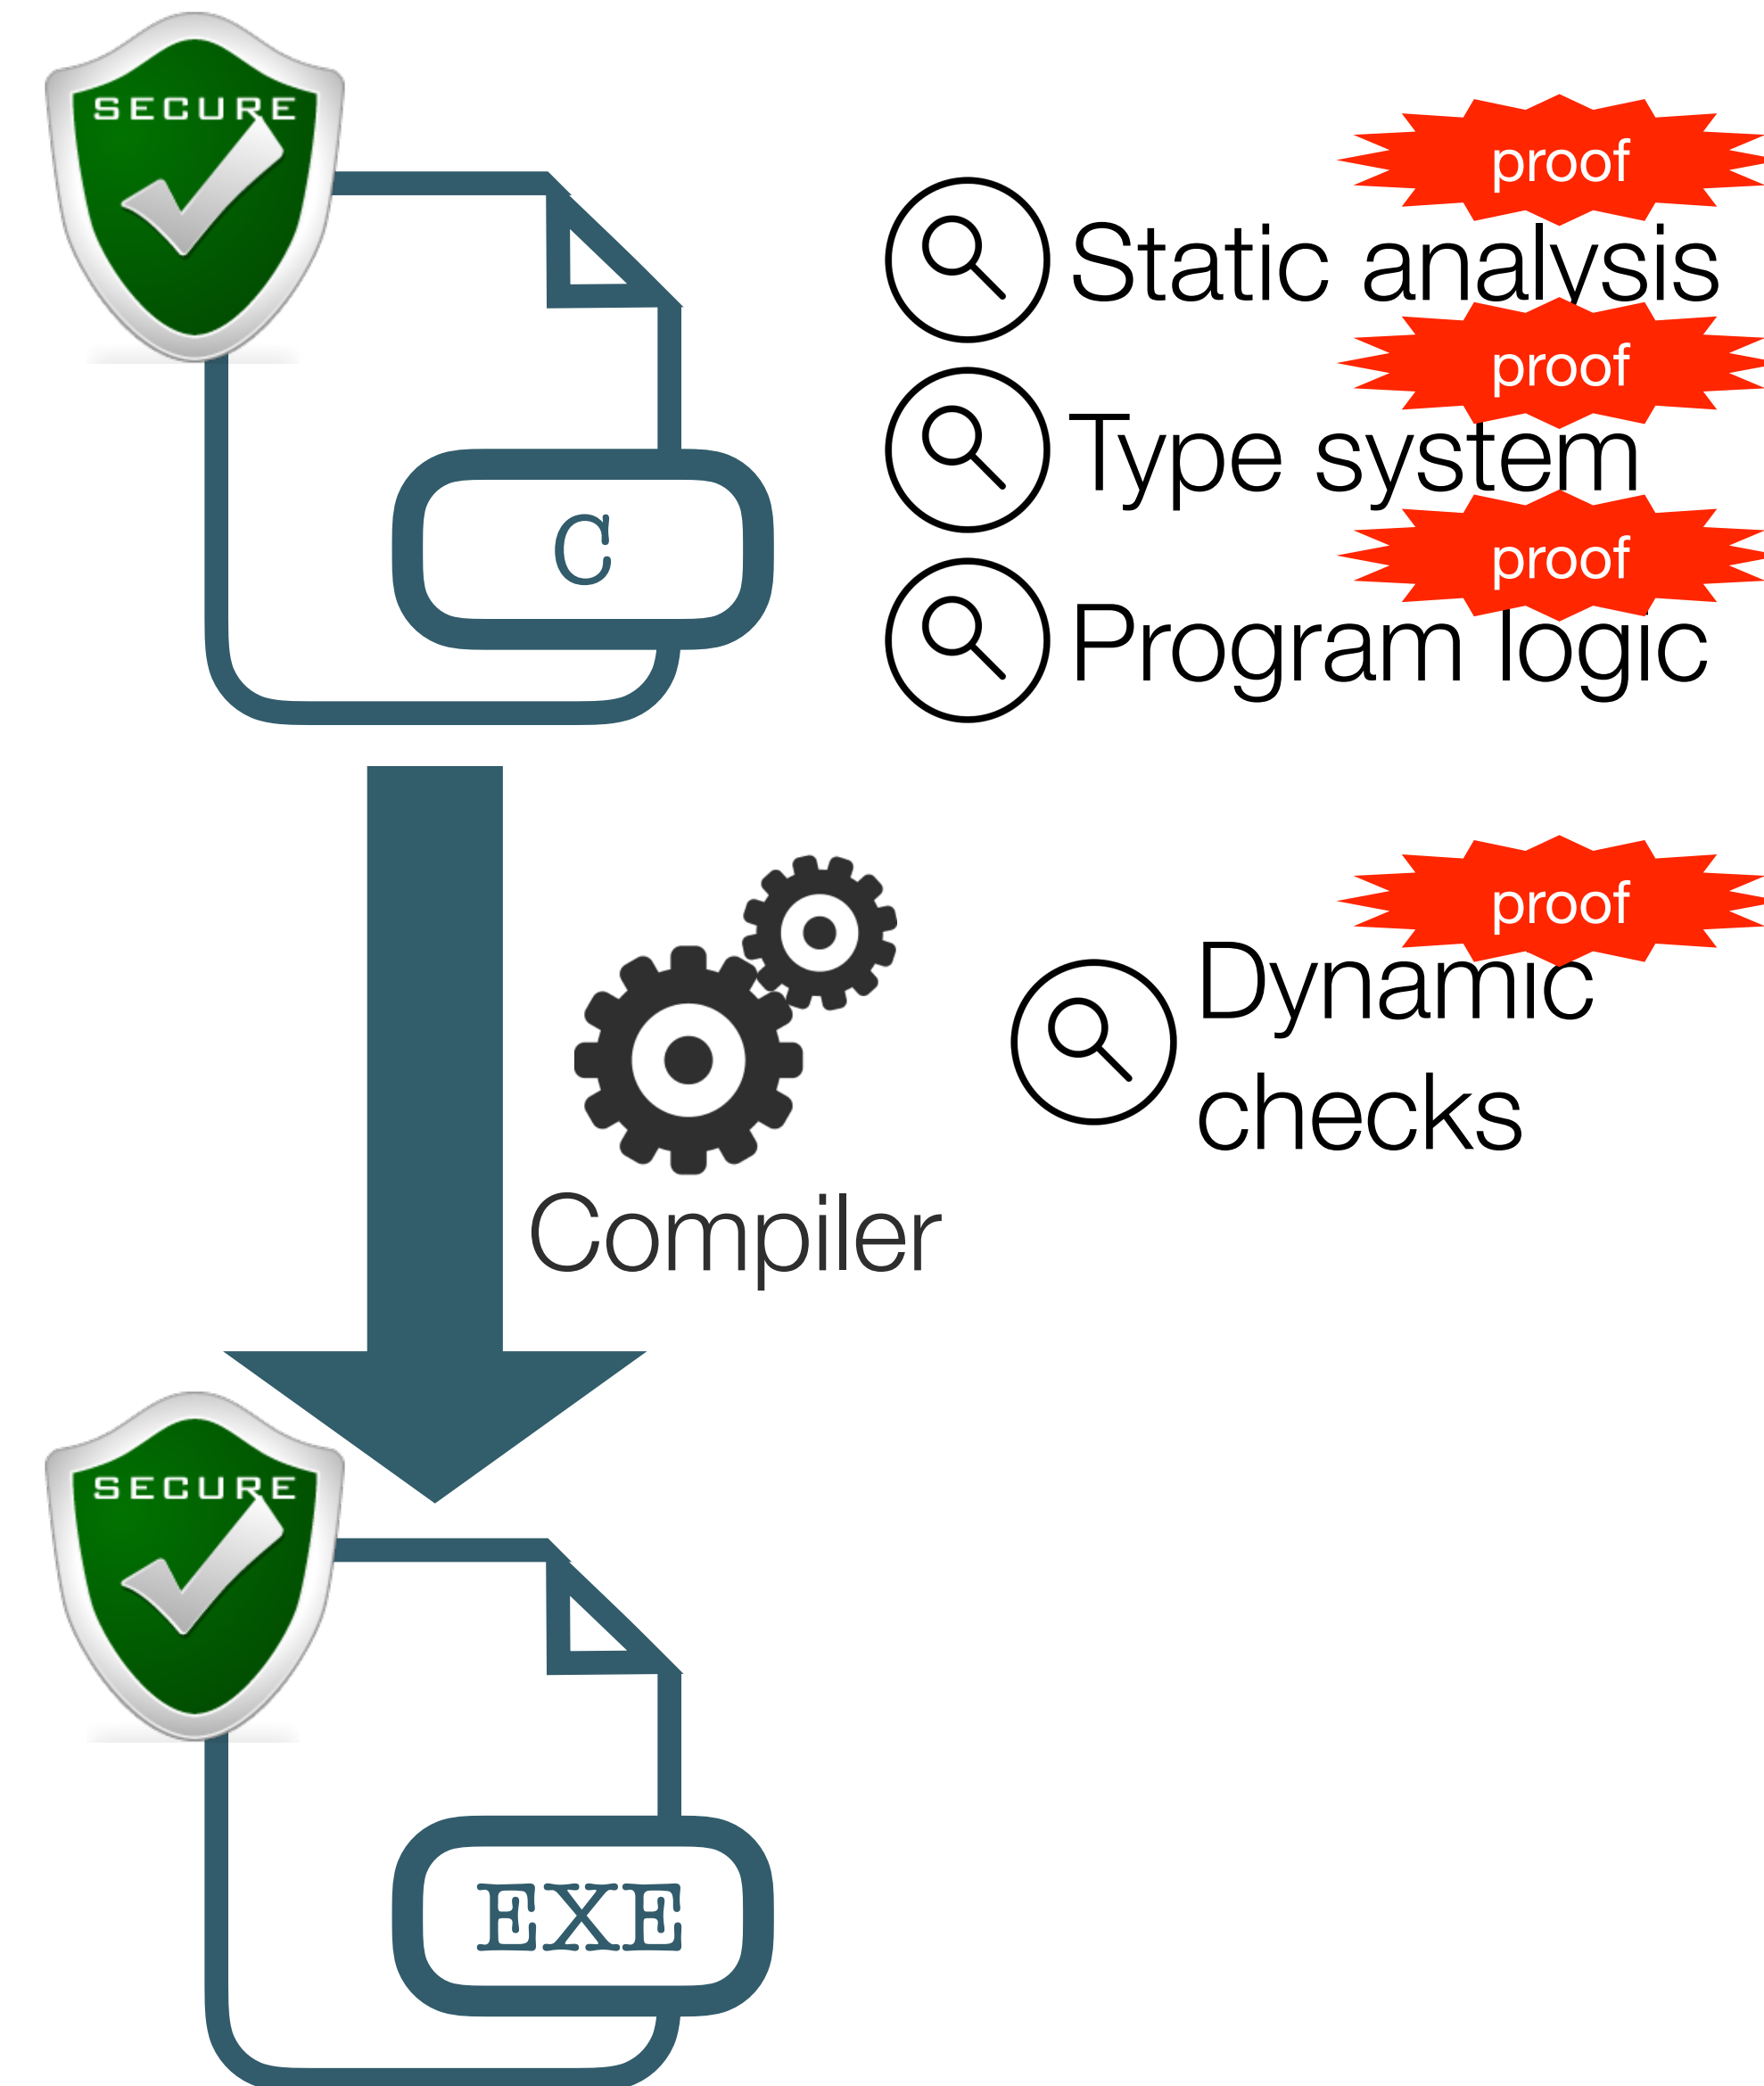
Our Research Program

- Build secure programming abstractions at source level (C-like)
- Make sure the compiler will generate executables that are as secure
- Reduce as much as possible the TCB (Trusted Computing Base) with formal proofs



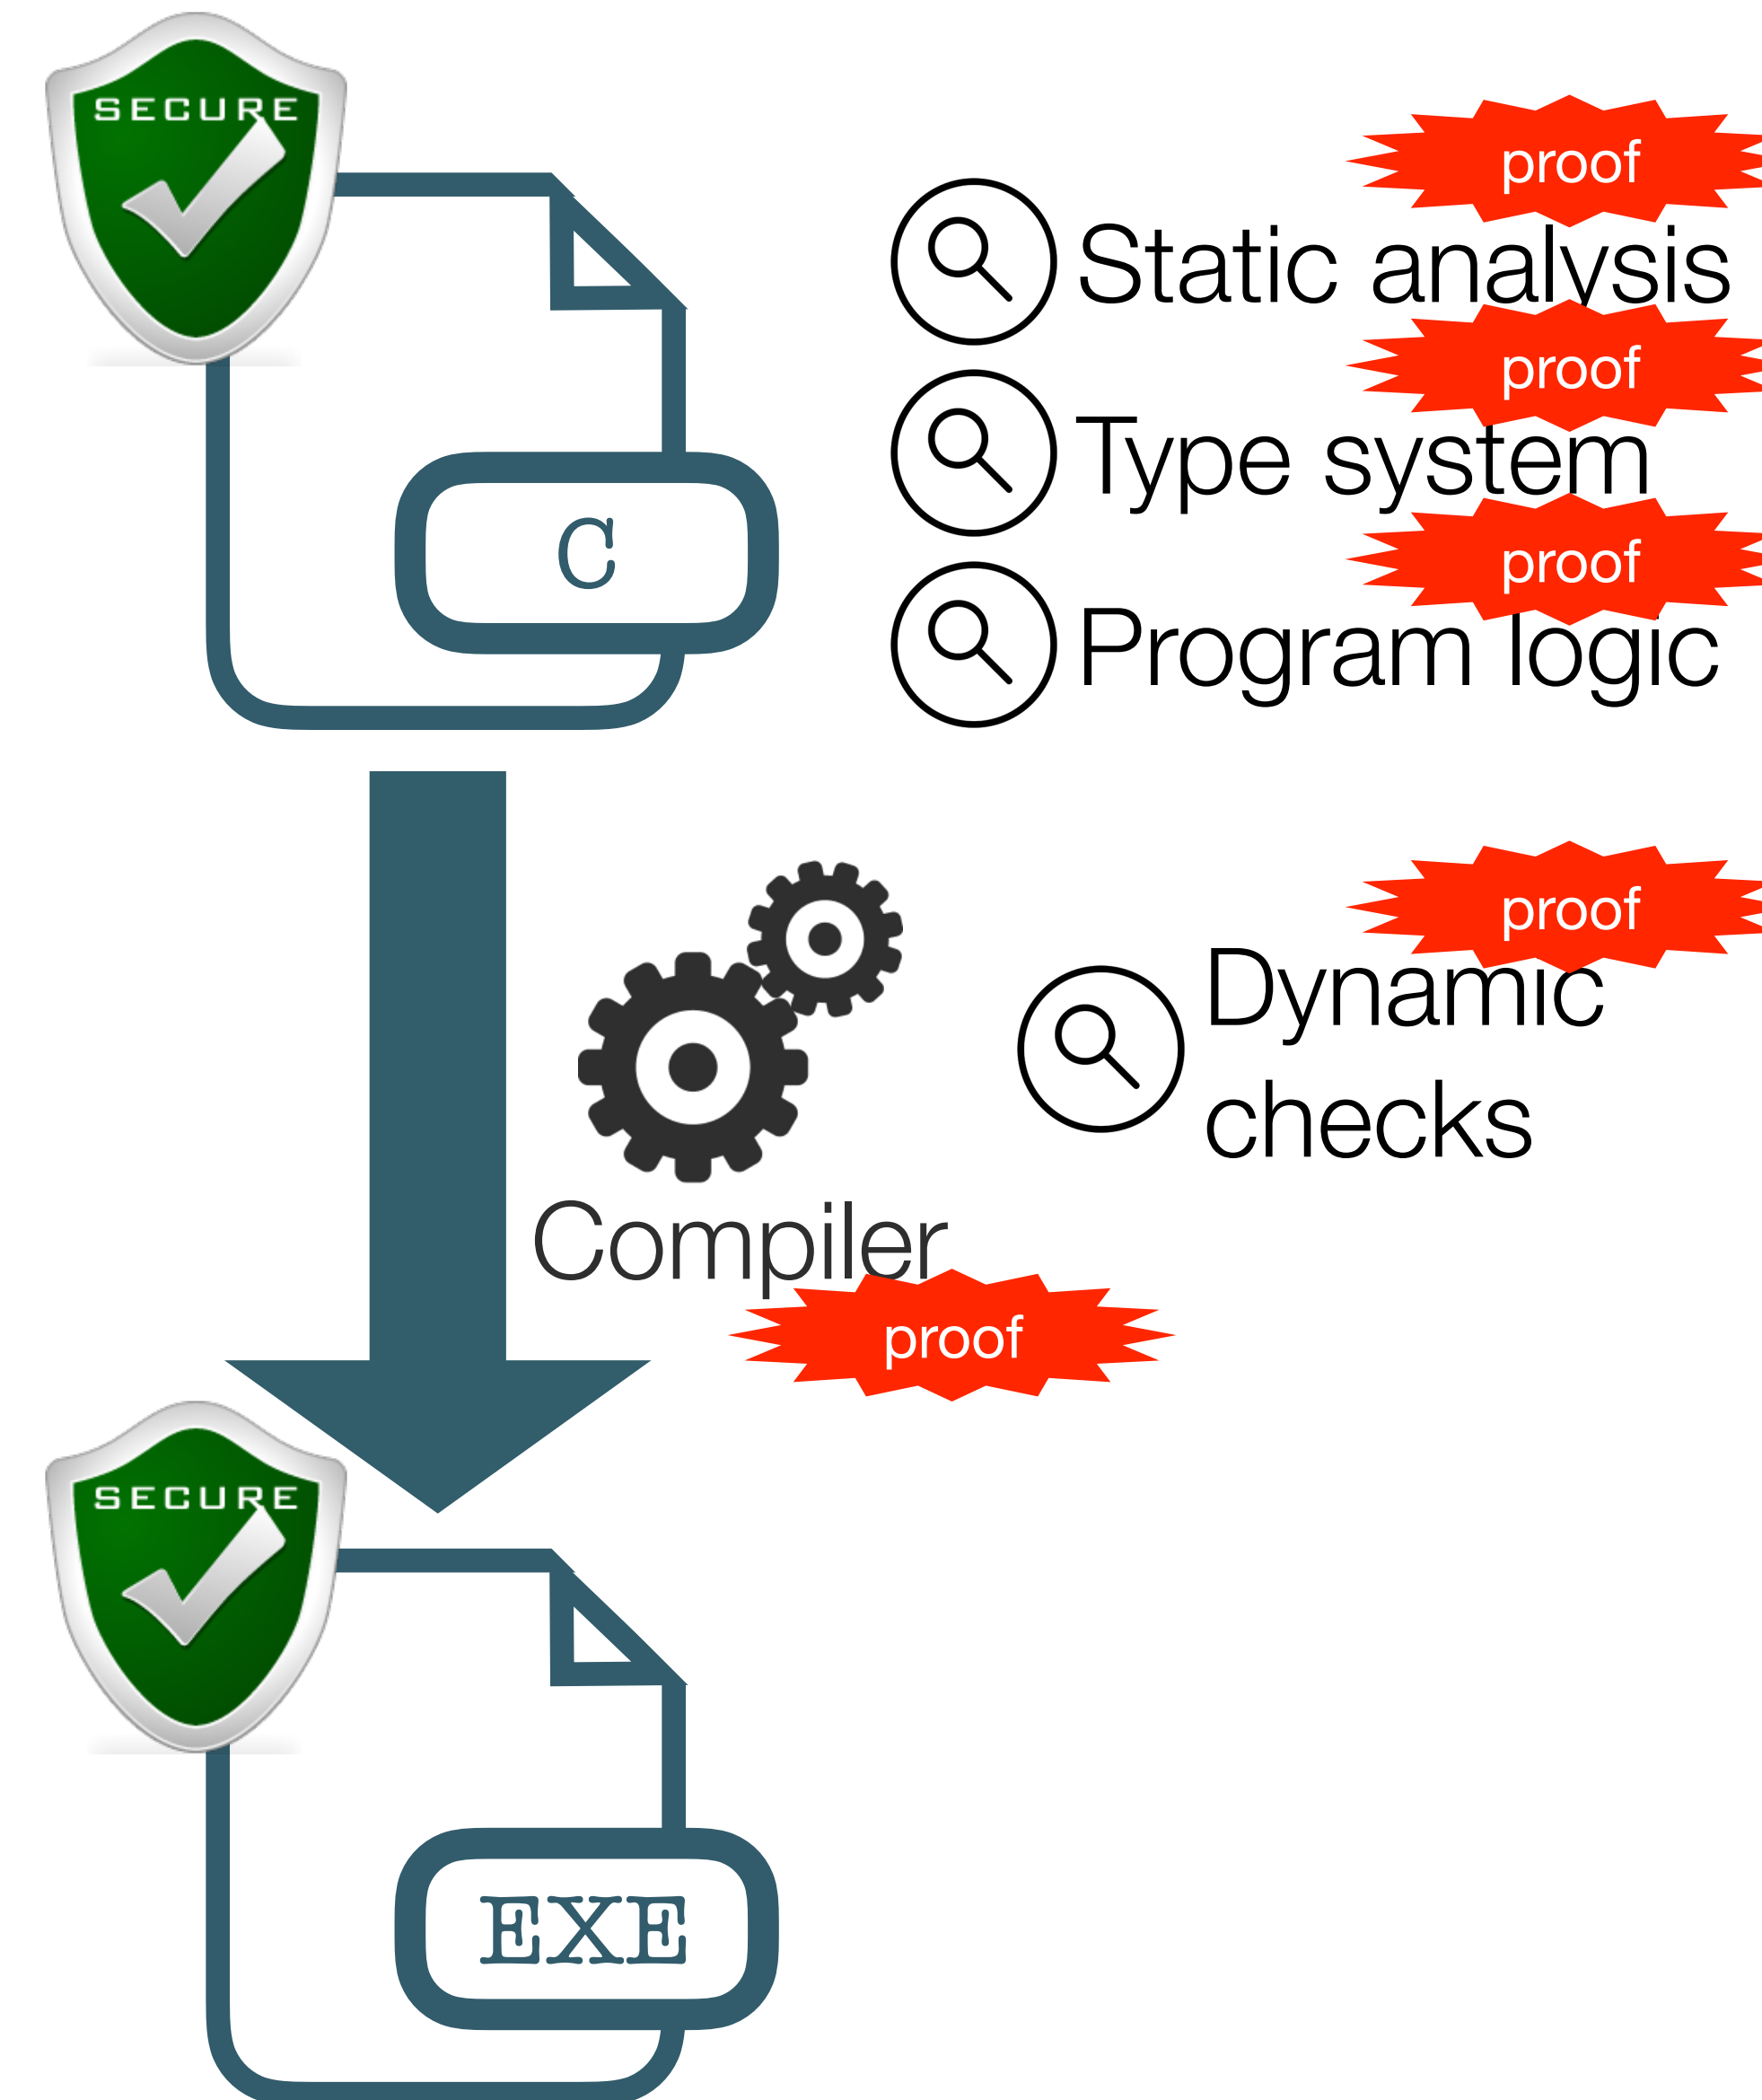
Our Research Program

- Build secure programming abstractions at source level (C-like)
- Make sure the compiler will generate executables that are as secure
- Reduce as much as possible the TCB (Trusted Computing Base) with formal proofs



Our Research Program

- Build secure programming abstractions at source level (C-like)
- Make sure the compiler will generate executables that are as secure
- Reduce as much as possible the TCB (Trusted Computing Base) with formal proofs



Cryptographic constant-time verification

- In [ESORICS'17] we provide a verification tool at C source level
 - it tracks taints in memory and checks the constant-time property
 - it is based on the Verasco C abstract interpreter [POPL'15]
- In this work [POPL'20]
 - we prove the CompCert compiler preserves the constant-time property




S. Blazy, D. Pichardie, A. Trieu.
Verifying Constant-Time Implementations by Abstract Interpretation.
ESORICS 2017 & Journal of Computer Security 2019.

J.-H. Jourdan, V. Laporte, S. Blazy, X. Leroy, and D. Pichardie.
A formally-verified C static analyzer.
POPL'15.

G. Barthe, S. Blazy, B. Grégoire, R. Hutin, V. Laporte, D. Pichardie, A. Trieu.
Formal verification of a constant-time preserving C compiler.
POPL'20

$$\forall p, \text{ConstantTime}(p) \Rightarrow \text{ConstantTime}(\text{compile}(p))$$

Verified Compilation

- Proving semantic properties on non-toy compilers requires a machine-checked proof
- CompCert [Leroy06] is a milestone in this area
 - a moderately optimizing compiler for C
 - programmed and verified with the Coq proof assistant
 - now being used in commercial settings and for software certification [Kästner18]
- CompCert theorems show
 - it preserves memory safety 
 - it preserves observable behaviors 
 - but they says nothing about side channels attacks 

This work

- Makes precise what secure compilation means for cryptographic constant-time
- Provides a machine checked-proof that a mildly modified version of the CompCert compiler preserves cryptographic constant-time
- Explains how to turn a pre-existing formally-verified compiler into a formally-verified secure compiler
- Provides a proof toolkit for proving security preservation with simulation diagrams

Some background on CompCert

Background: verifying a compiler

CompCert, a moderately optimizing C compiler usable for critical embedded software

= compiler + proof that the compiler does not introduce bugs

Using the Coq proof assistant, X. Leroy proves the following semantic preservation property:

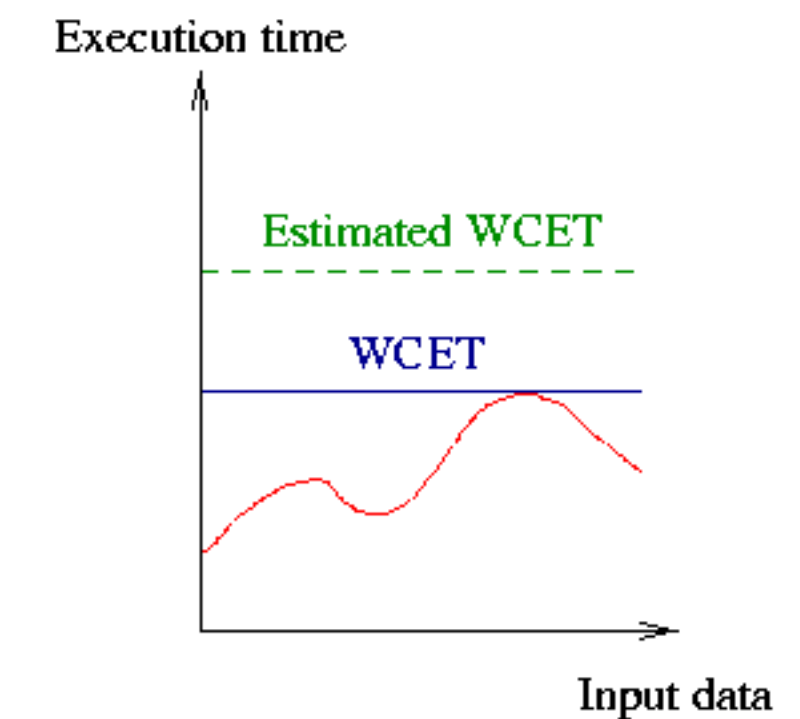
For all source programs S and compiler-generated code C , if the compiler generates machine code C from source S , without reporting a compilation error, then « C behaves like S ».

Compiler written from scratch, along with its proof; not trying to prove an existing compiler

Compcert meets the industrial world

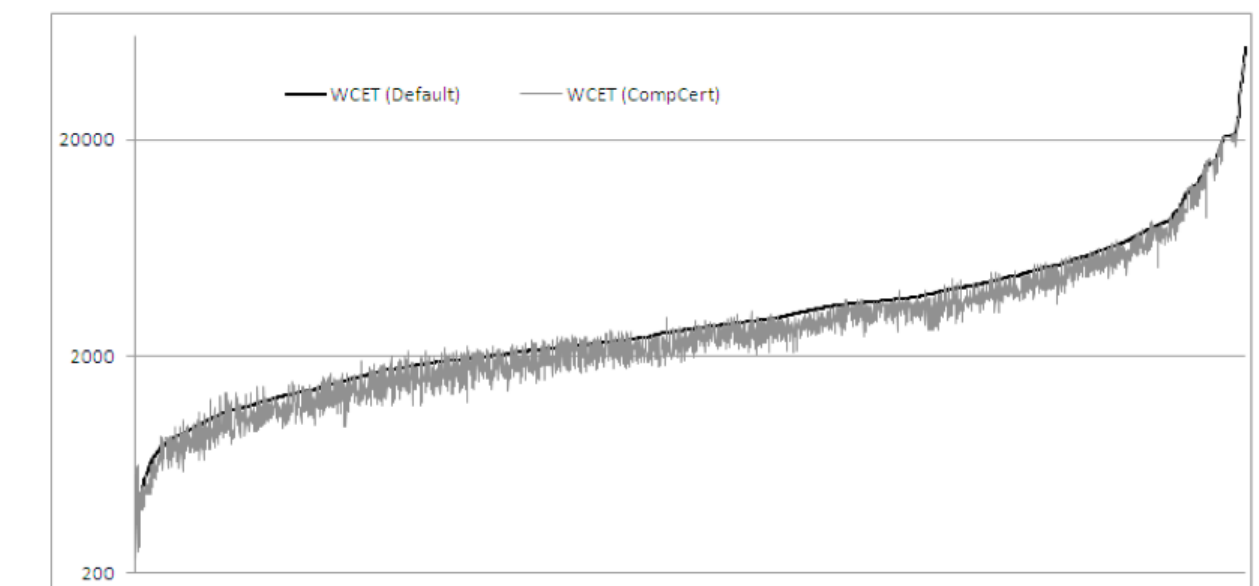
Fly-by-wire software, for recent Airbus planes

- control-command code generated from block diagrams (3600 files, 3.96 MB of assembly code)
- minimalistic OS



Results

- Estimated WCET for each file
- Average improvement per file: 14%
- Compiled with CompCert 2.3, May 2014



Conformance to the certification process (DO-178)

- Trade-off between traceability guarantees and efficiency of the generated code

Fly-by-wire software

- control-command
- (3600 files, 3.96
- minimalistic OS

Results

- Estimated WCET
- Average improve
- Compiled with O

Conformance to the

- Trade-off b

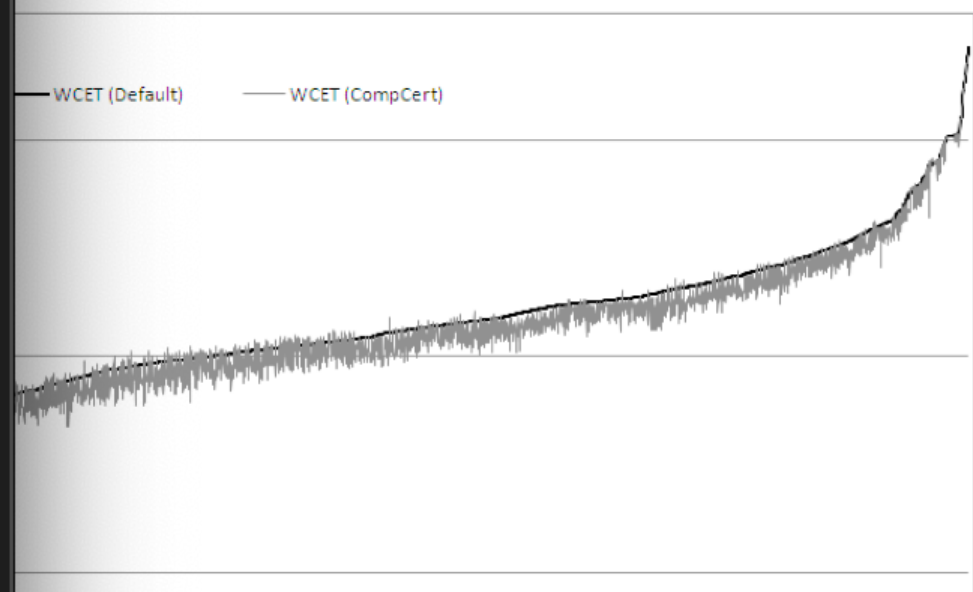
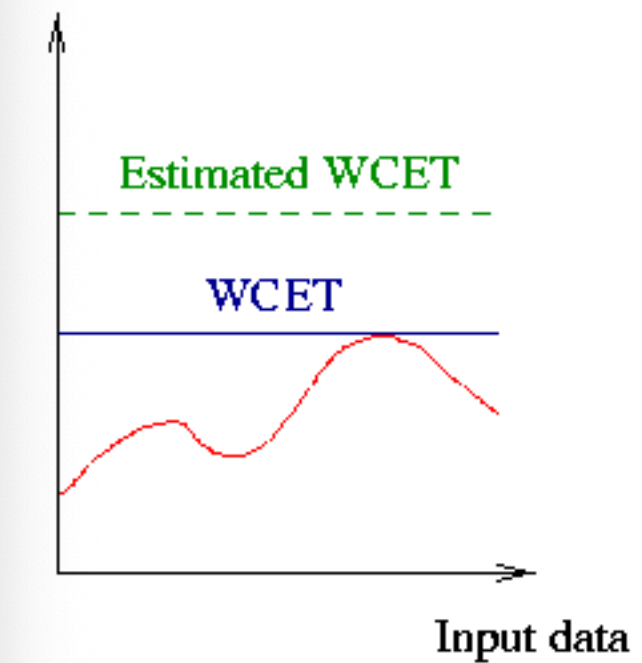
<https://www.absint.com/compcert/>

ted code

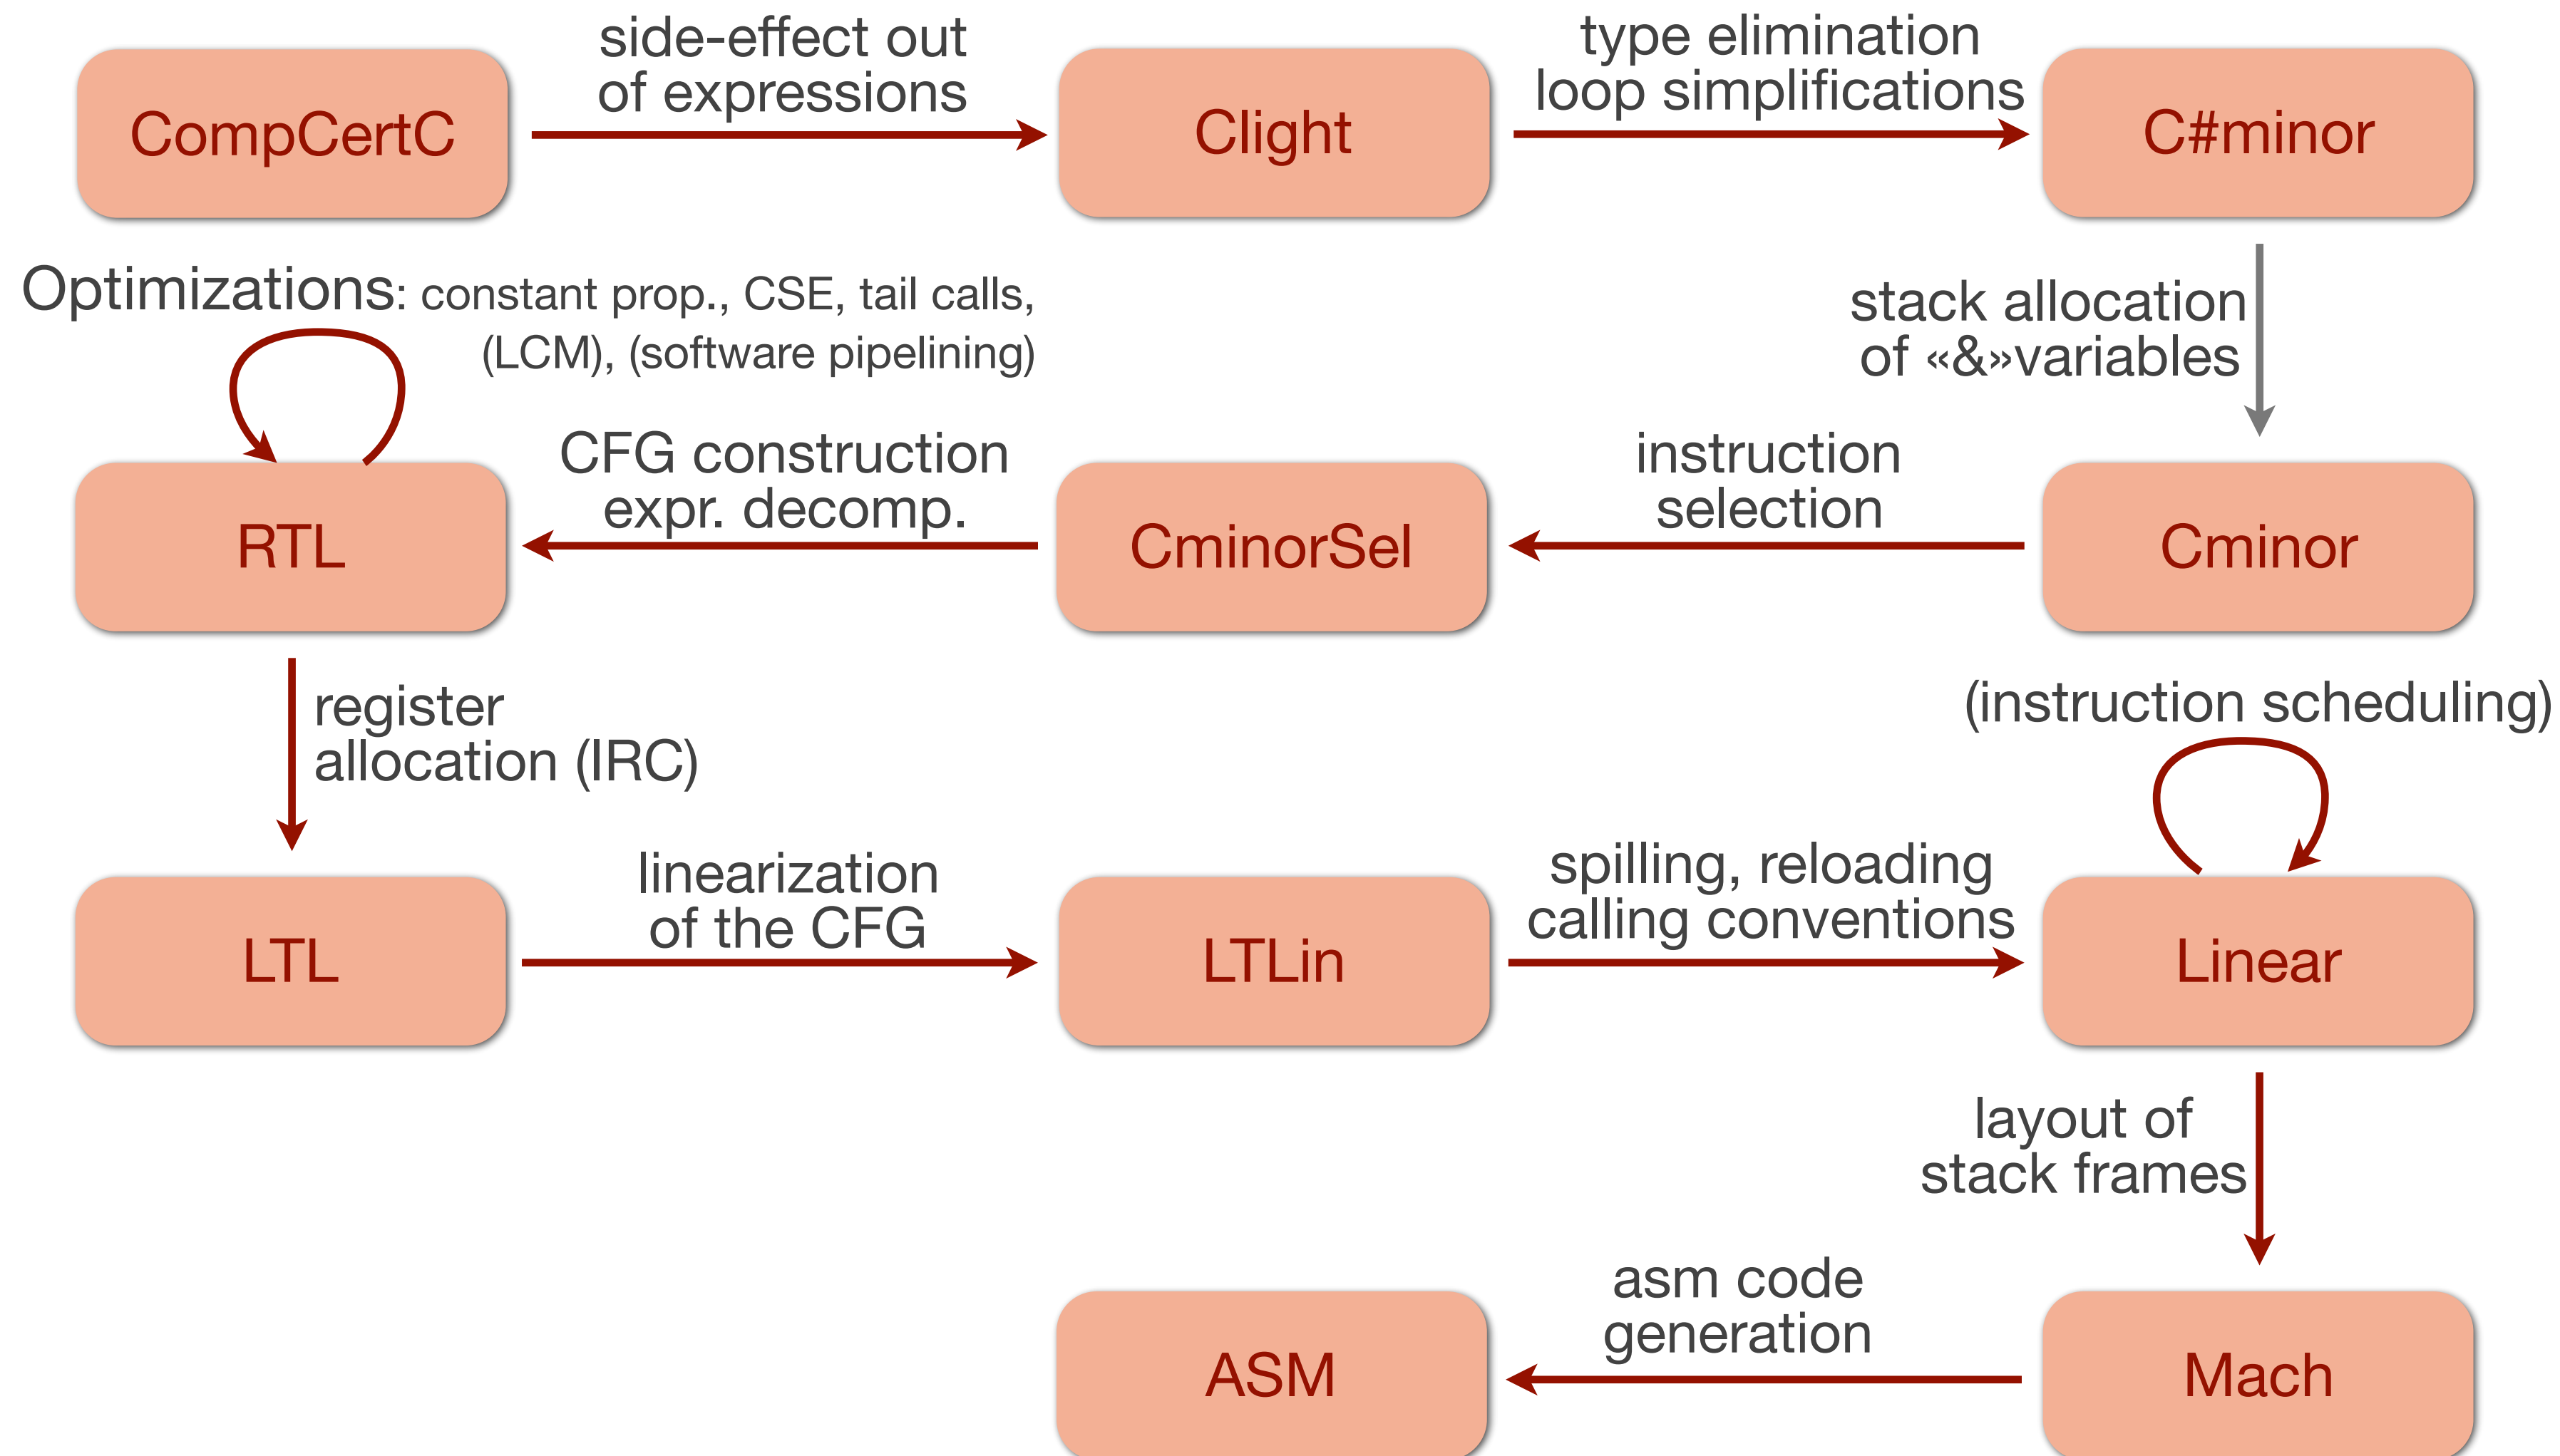
The screenshot shows the Absint website's 'Formally verified compilation' page. The page title is 'Formally verified compilation' and the sub-header is 'CompCert'. The main text describes CompCert as a formally verified optimizing C compiler. It lists supported architectures: ARM, PowerPC, x86, and RISC-V. A diagram illustrates the compilation pipeline from C source code through various stages: Clight, Cfminor, Cminor, RTL, LTL, Linear, Mach, PPC, and PowerPC assembly. A graph on the right shows 'Execution time' vs 'Input data', comparing 'Estimated WCET' (a horizontal dashed line) and 'WCET' (a fluctuating line). A quote from a 2011 study by Regehr, Yang et al. is included, praising the reliability of CompCert. A list of 'Formally verified optimizations' is provided at the bottom, including register allocation, instruction selection, and constant propagation.

ld

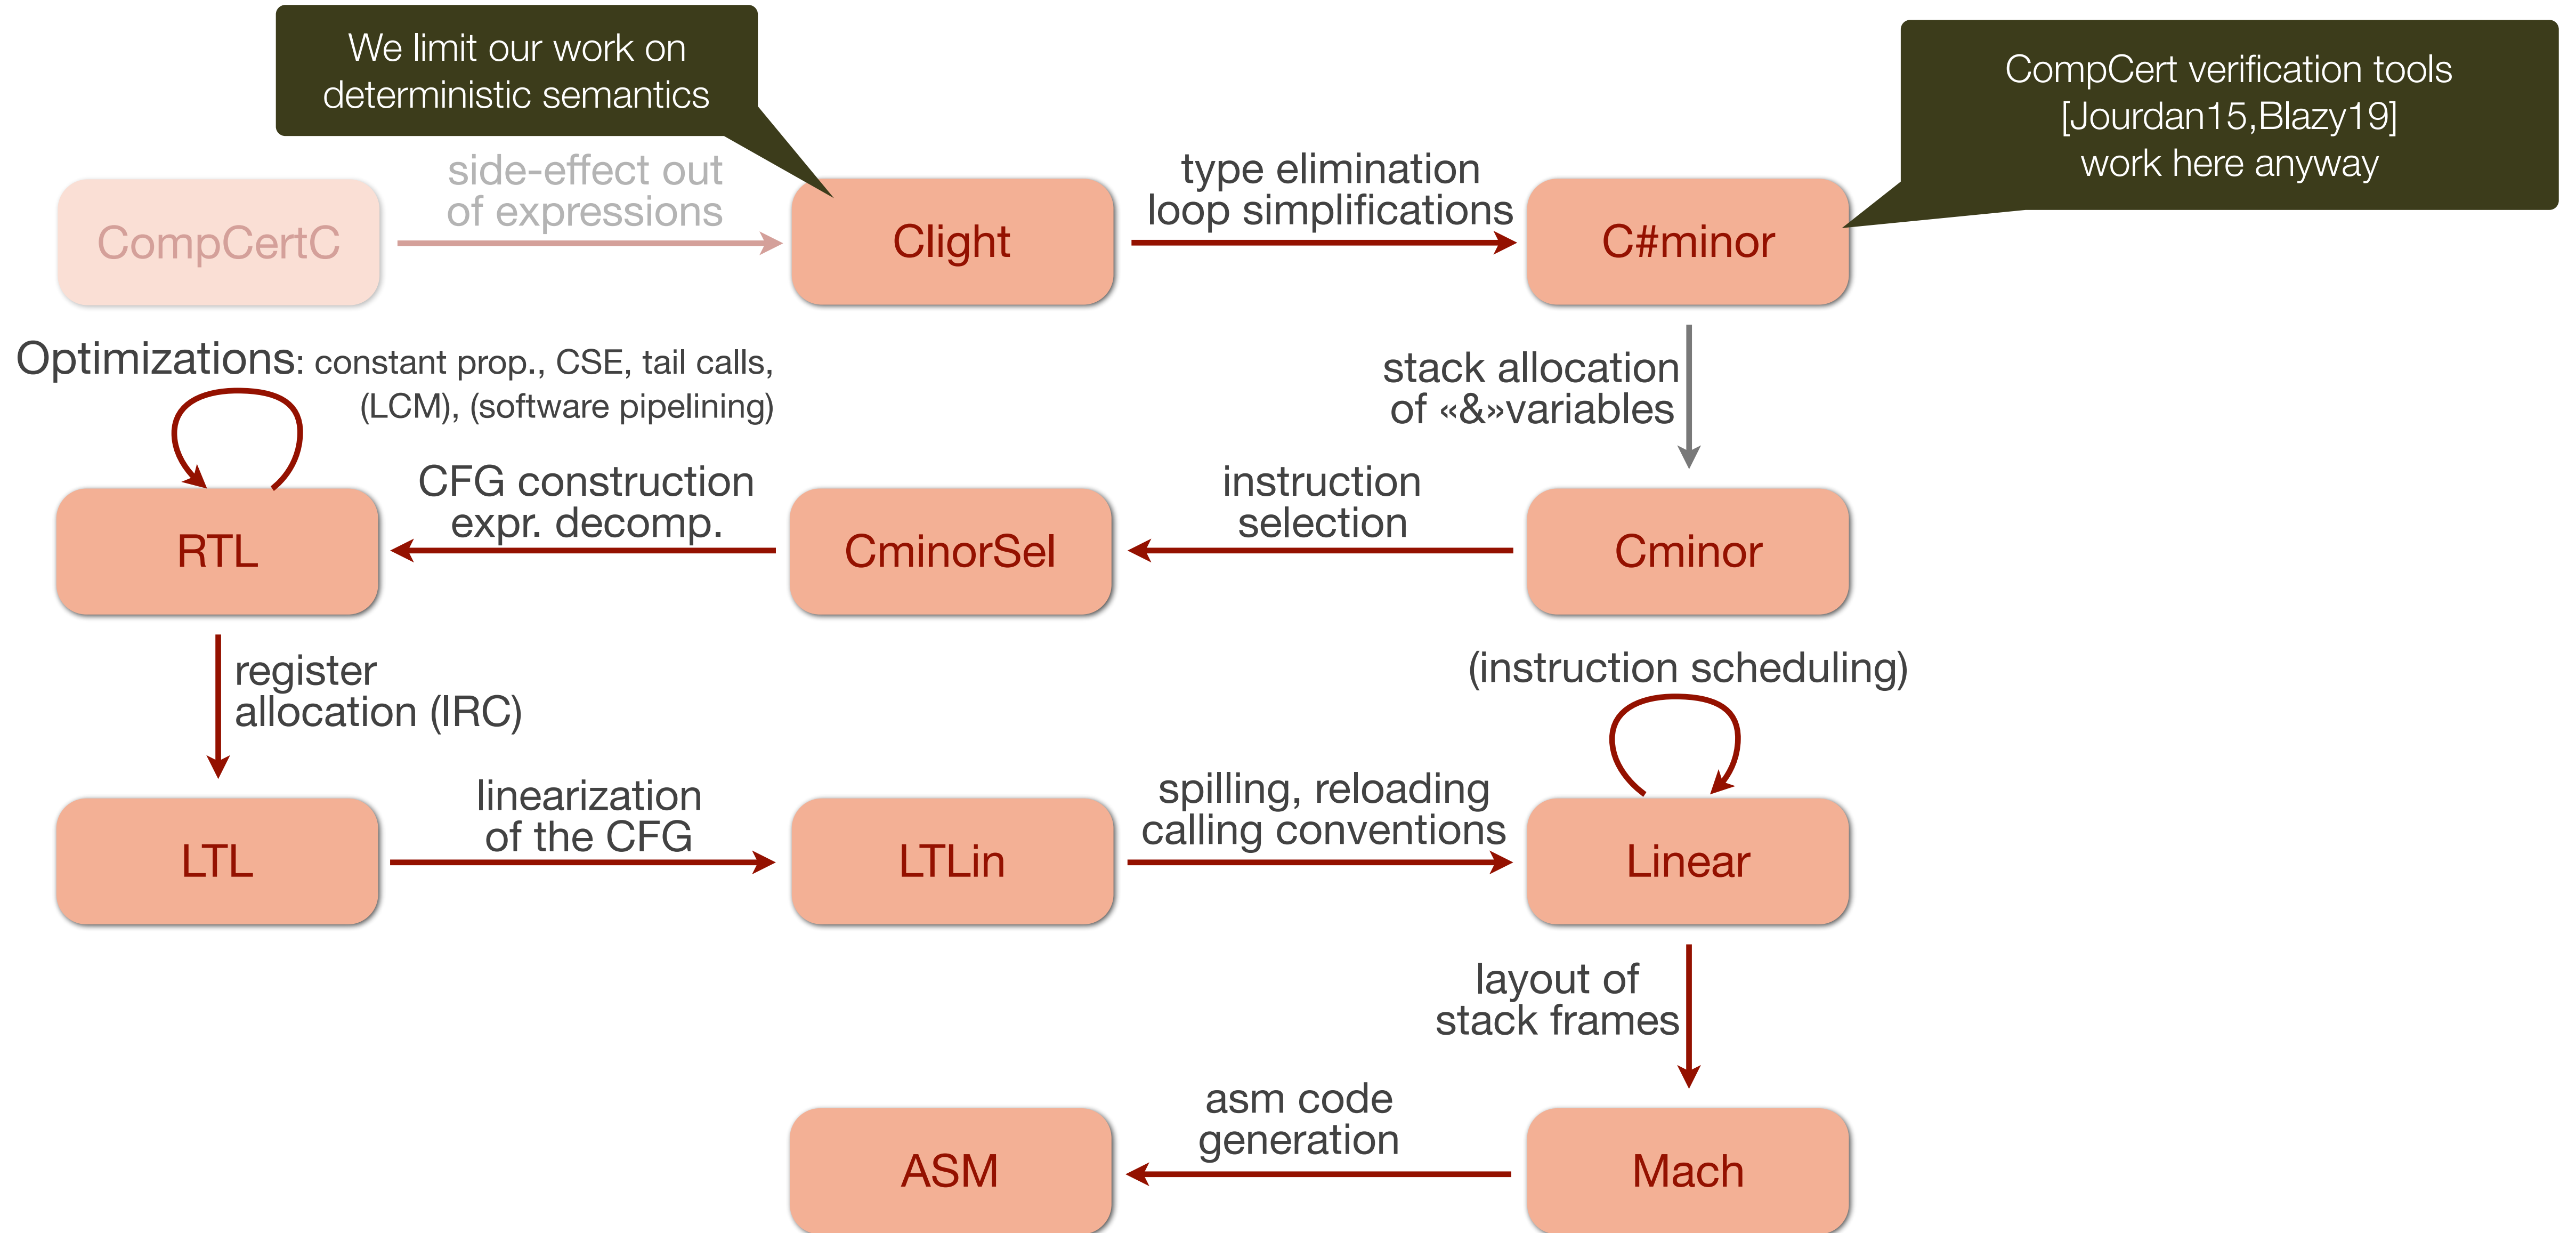
Execution time



CompCert: 1 compiler, 11 languages...



CompCert: 1 compiler, 11 languages...



CompCert: ... and 17 preservations proofs

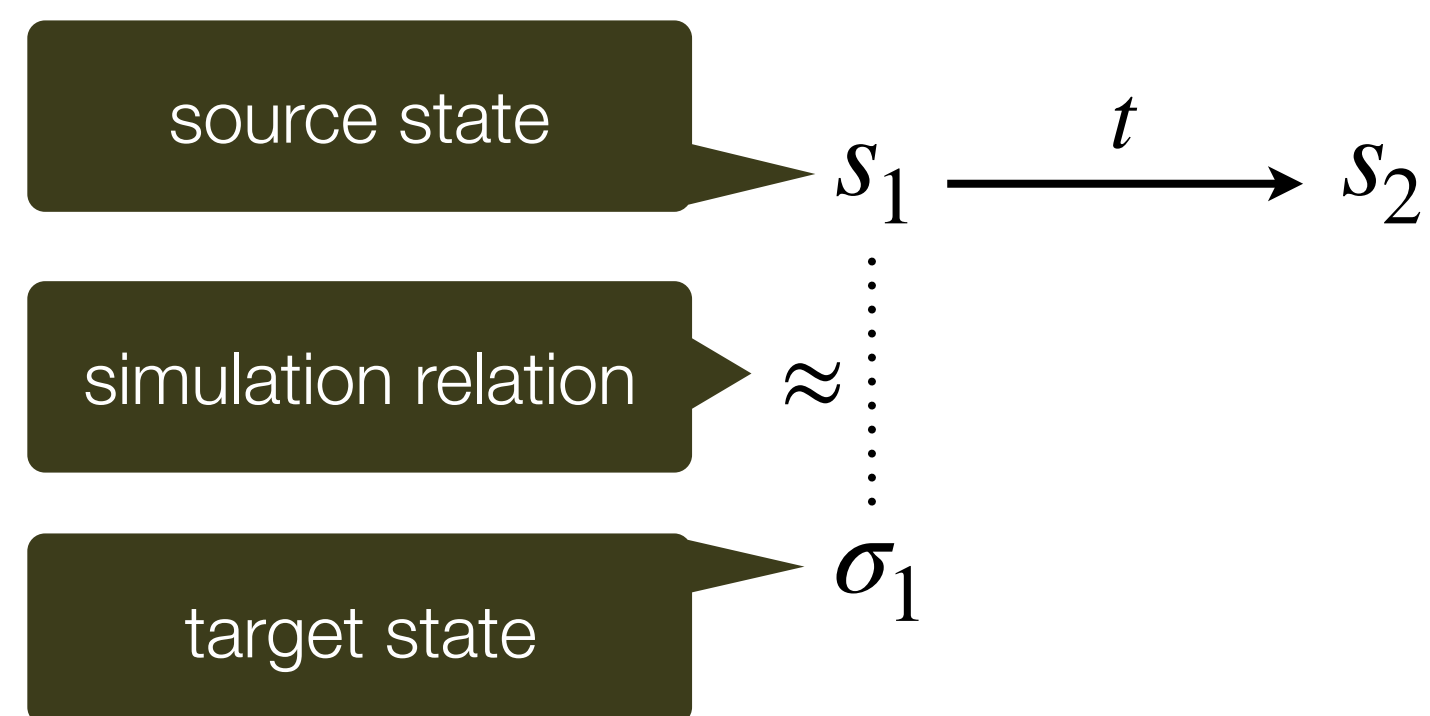
Compiler pass	Explanation on the pass
Cshngen	Type elaboration, simplification of control
Cminorgen	Stack allocation
Selection	Recognition of operators and addr. modes
RTLgen	Generation of CFG and 3-address code
Tailcall	Tailcall recognition
Inlining	Function inlining
Renumber	Renumbering CFG nodes
ConstProp	Constant propagation
CSE	Common subexpression elimination
Deadcode	Redundancy elimination
Allocation	Register allocation
Tunneling	Branch tunneling
Linearize	Linearization of CFG
CleanupLabels	Removal of unreferenced labels
Debugvar	Synthesis of debugging information
Stacking	Laying out stack frames
Asmggen	Emission of assembly code

CompCert preservation proof methodology

- Each language is given an **operational semantics** $s \xrightarrow{t} s'$ that models a small step transition from a state s to a state s' by emitting a trace of external events t .
- From this stems a notion of **program behavior** (event trace) for complete (possibly infinite) executions.
- Behavior preservation is proved via backward and forward simulation, but thanks to language determinism, **forward simulation** is enough.

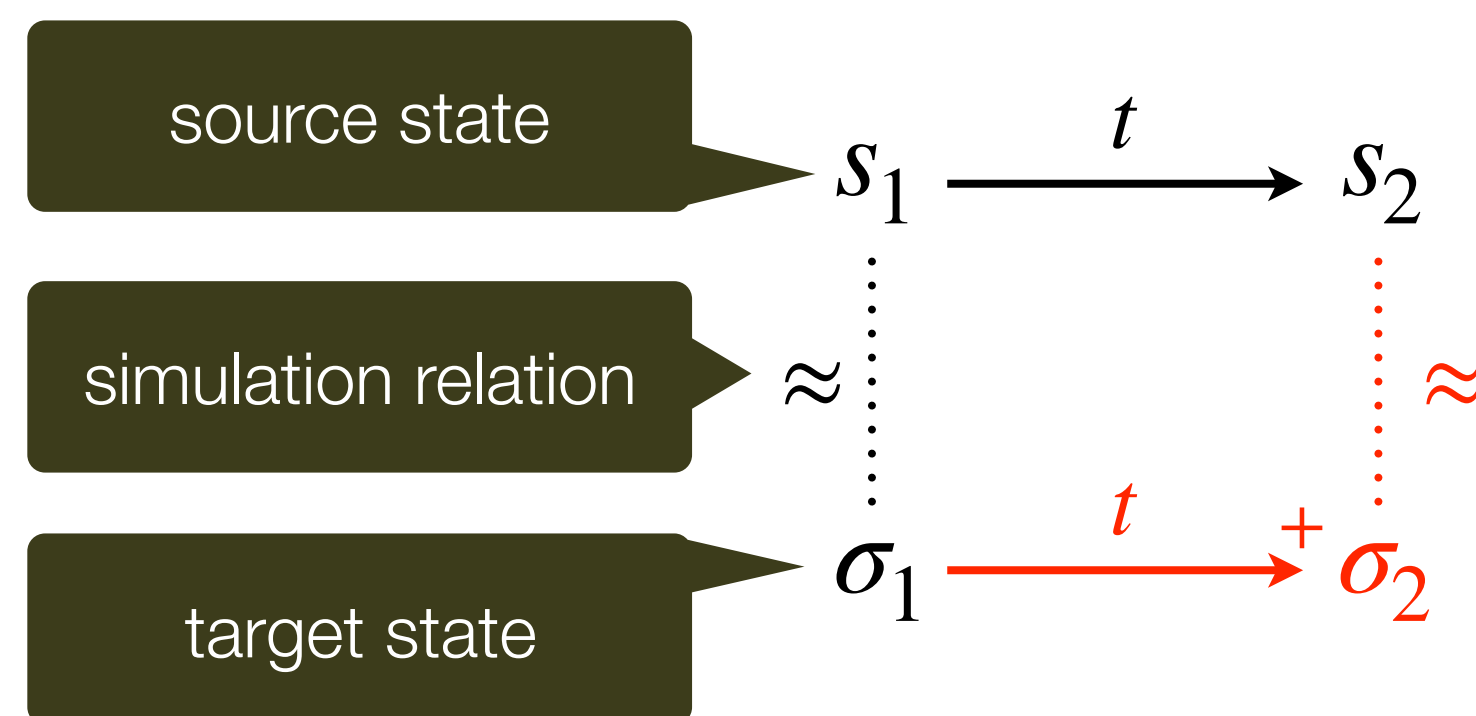
CompCert preservation proof methodology

- Each language is given an **operational semantics** $s \xrightarrow{t} s'$ that models a small step transition from a state s to a state s' by emitting a trace of external events t .
- From this stems a notion of **program behavior** (event trace) for complete (possibly infinite) executions.
- Behavior preservation is proved via backward and forward simulation, but thanks to language determinism, **forward simulation** is enough.




CompCert preservation proof methodology

- Each language is given an **operational semantics** $s \xrightarrow{t} s'$ that models a small step transition from a state s to a state s' by emitting a trace of external events t .
- From this stems a notion of **program behavior** (event trace) for complete (possibly infinite) executions.
- Behavior preservation is proved via backward and forward simulation, but thanks to language determinism, **forward simulation** is enough.



CompCert preservation proof methodology

- Each language is given an **operational semantics** $s \xrightarrow{t} s'$ that models a small step transition from a state s to a state s' by emitting a trace of external events t .
- From this stems a notion of **program behavior** (event trace) for complete (possibly infinite) executions.
- Behavior preservation is proved via backward and forward simulation, but thanks to language determinism, **forward simulation** is enough.



Verified Static Analysis meets CompCert

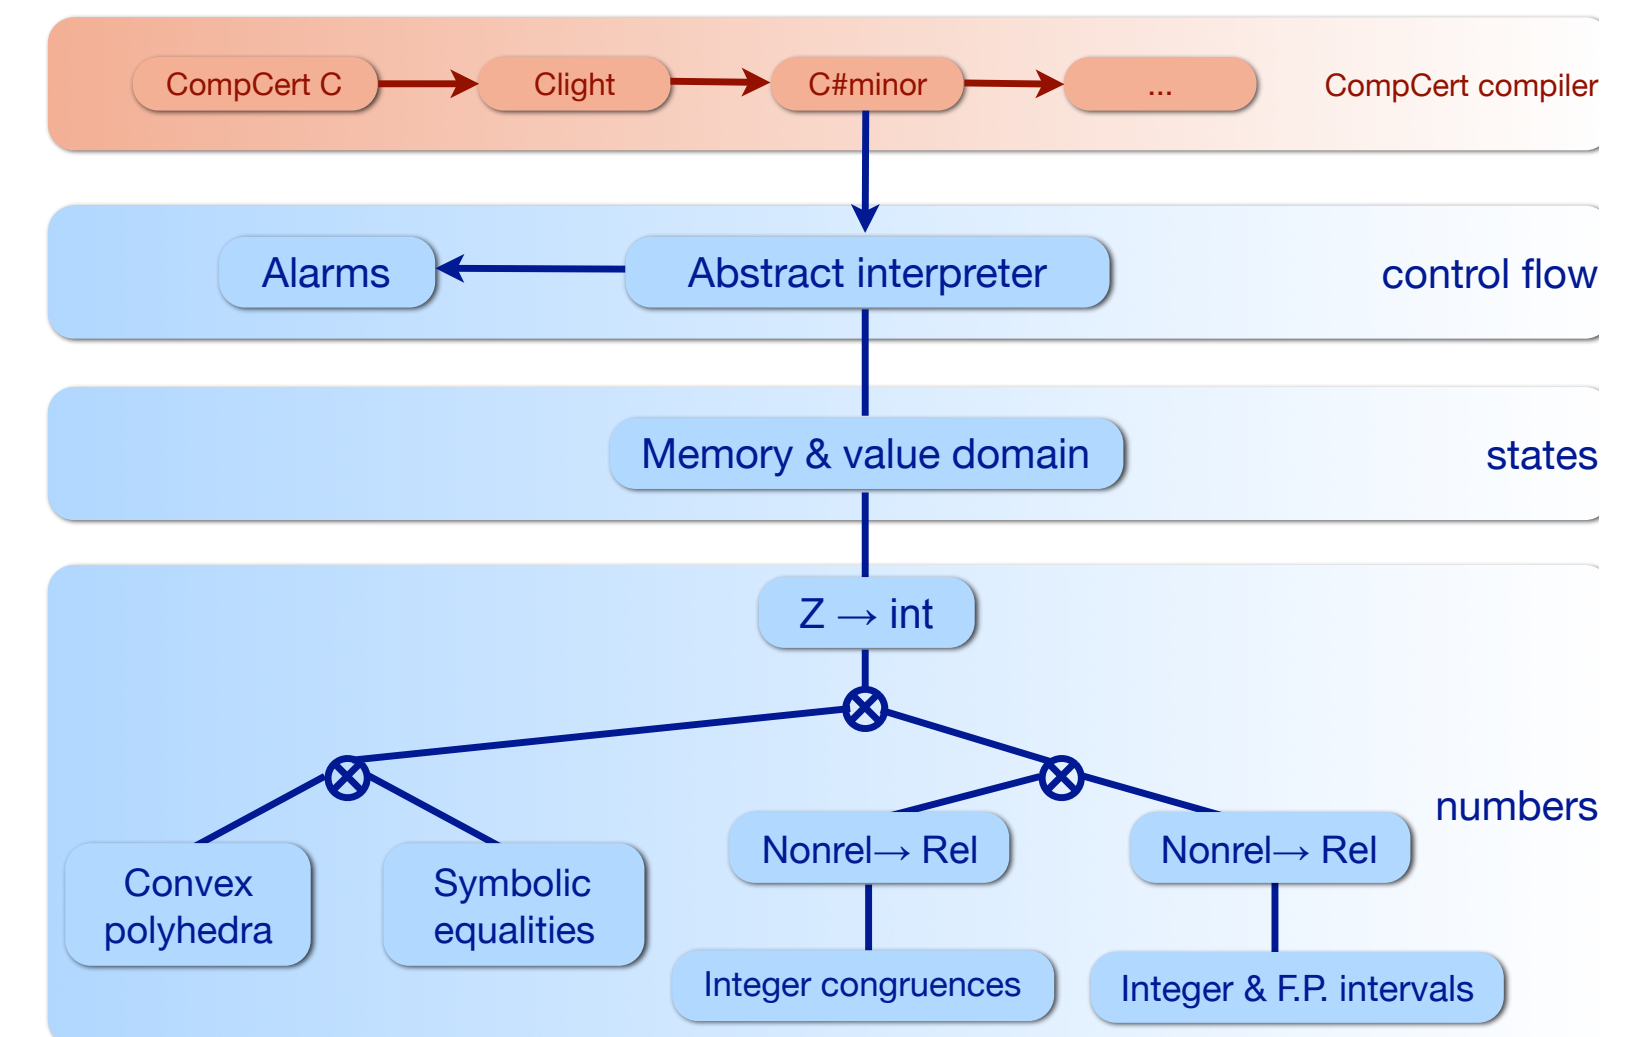
The verified C static analyzer Verasco [POPL'15]

Goal: develop and verify in Coq a realistic static analyzer by abstract interpretation

- language analyzed: the CompCert subset of C
- nontrivial abstract domains, including relational domains
- modular architecture inspired from Astrée's
- to prove the absence of undefined behaviors in C source programs

Slogan:

- if « CompCert \approx 1/10th of GCC but formally verified »,
- likewise « Verasco \approx 1/10th of Astrée but formally verified »



Verasco architecture


Defining Cryptographic Constant-Time Preservation

Cryptographic constant-time property: defining leakages

- We enrich the CompCert traces of events with **leakages** of two types
 - either the truth value of a condition,
 - or a pointer representing the address of
 - either a memory access (i.e., a load or a store)
 - or a called function
- Using **event erasure**, from $s \xrightarrow{t} s'$ we can extract
 - the compile-only judgment $s \xrightarrow{t} \text{comp} s'$
 - the leak-only judgment $s \xrightarrow{t} \text{leak} s'$
- **Program leakage** is defined as the behavior of the $\rightarrow \text{leak}$ semantics

Cryptographic constant-time property: preservation

- We note $\varphi(s, s')$ the fact that two initial states s and s' share the same values for public inputs, but may differ on the values of secret inputs
- A program is **constant-time secure w.r.t. φ** if for two initial states s and s' such that $\varphi(s, s')$ holds, then both leak-only executions starting from s and s' observe the same leakage



Cryptographic constant-time property: preservation

- We note $\varphi(s, s')$ the fact that two initial states s and s' share the same values for public inputs, but may differ on the values of secret inputs
- A program is **constant-time secure w.r.t. φ** if for two initial states s and s' such that $\varphi(s, s')$ holds, then both leak-only executions starting from s and s' observe the same leakage



Main Theorem (Constant-Time security preservation): Let P be a safe Clight source program that is compiled into an x86 assembly program P' . If P is constant-time w.r.t. φ , then so is P' .

Proving Cryptographic Constant-Time Preservation

Proving cryptographic constant-time preservation

A proof engineering perspective

- Cryptographic constant-time preservation is a property about the leak-only semantics \rightarrow leak
- But existing CompCert simulation diagrams deal with the compile-only semantics \rightarrow comp
- Our proof engineering strategy is to benefit as much as possible from the **proof scripts** of these diagrams

Proving cryptographic constant-time preservation

A proof engineering perspective

- Cryptographic constant-time preservation is a property about the leak-only semantics $\rightarrow\text{leak}$
- But existing CompCert simulation diagrams deal with the compile-only semantics $\rightarrow\text{comp}$
- Our proof engineering strategy is to benefit as much as possible from the **proof scripts** of these diagrams

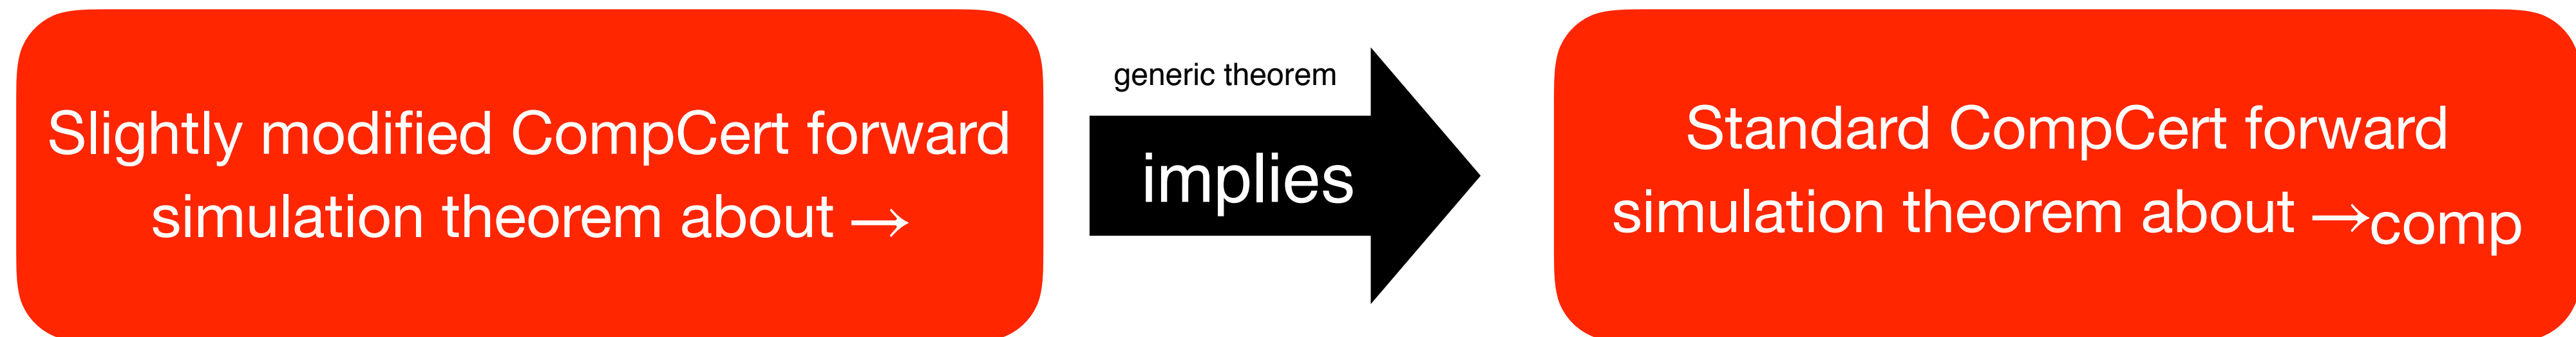
Standard CompCert forward
simulation theorem about $\rightarrow\text{comp}$

Standard CompCert forward
simulation proof script

Proving cryptographic constant-time preservation

A proof engineering perspective

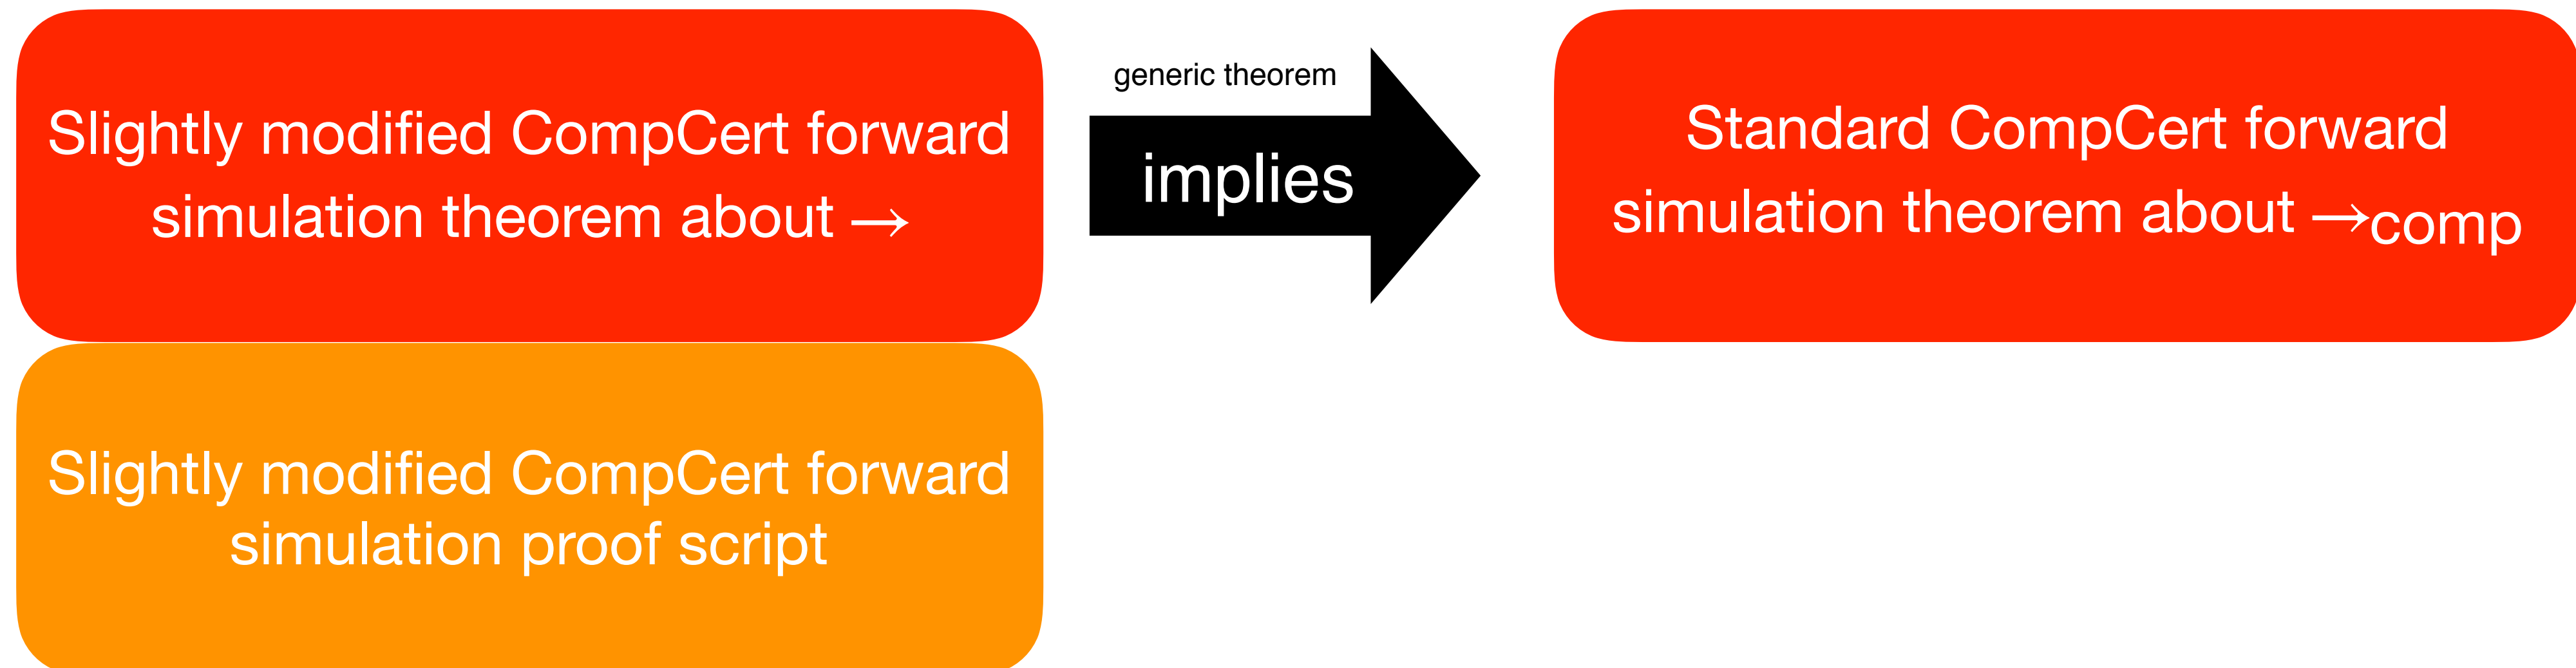
- Cryptographic constant-time preservation is a property about the leak-only semantics \rightarrow leak
- But existing CompCert simulation diagrams deal with the compile-only semantics \rightarrow comp
- Our proof engineering strategy is to benefit as much as possible from the **proof scripts** of these diagrams



Proving cryptographic constant-time preservation

A proof engineering perspective

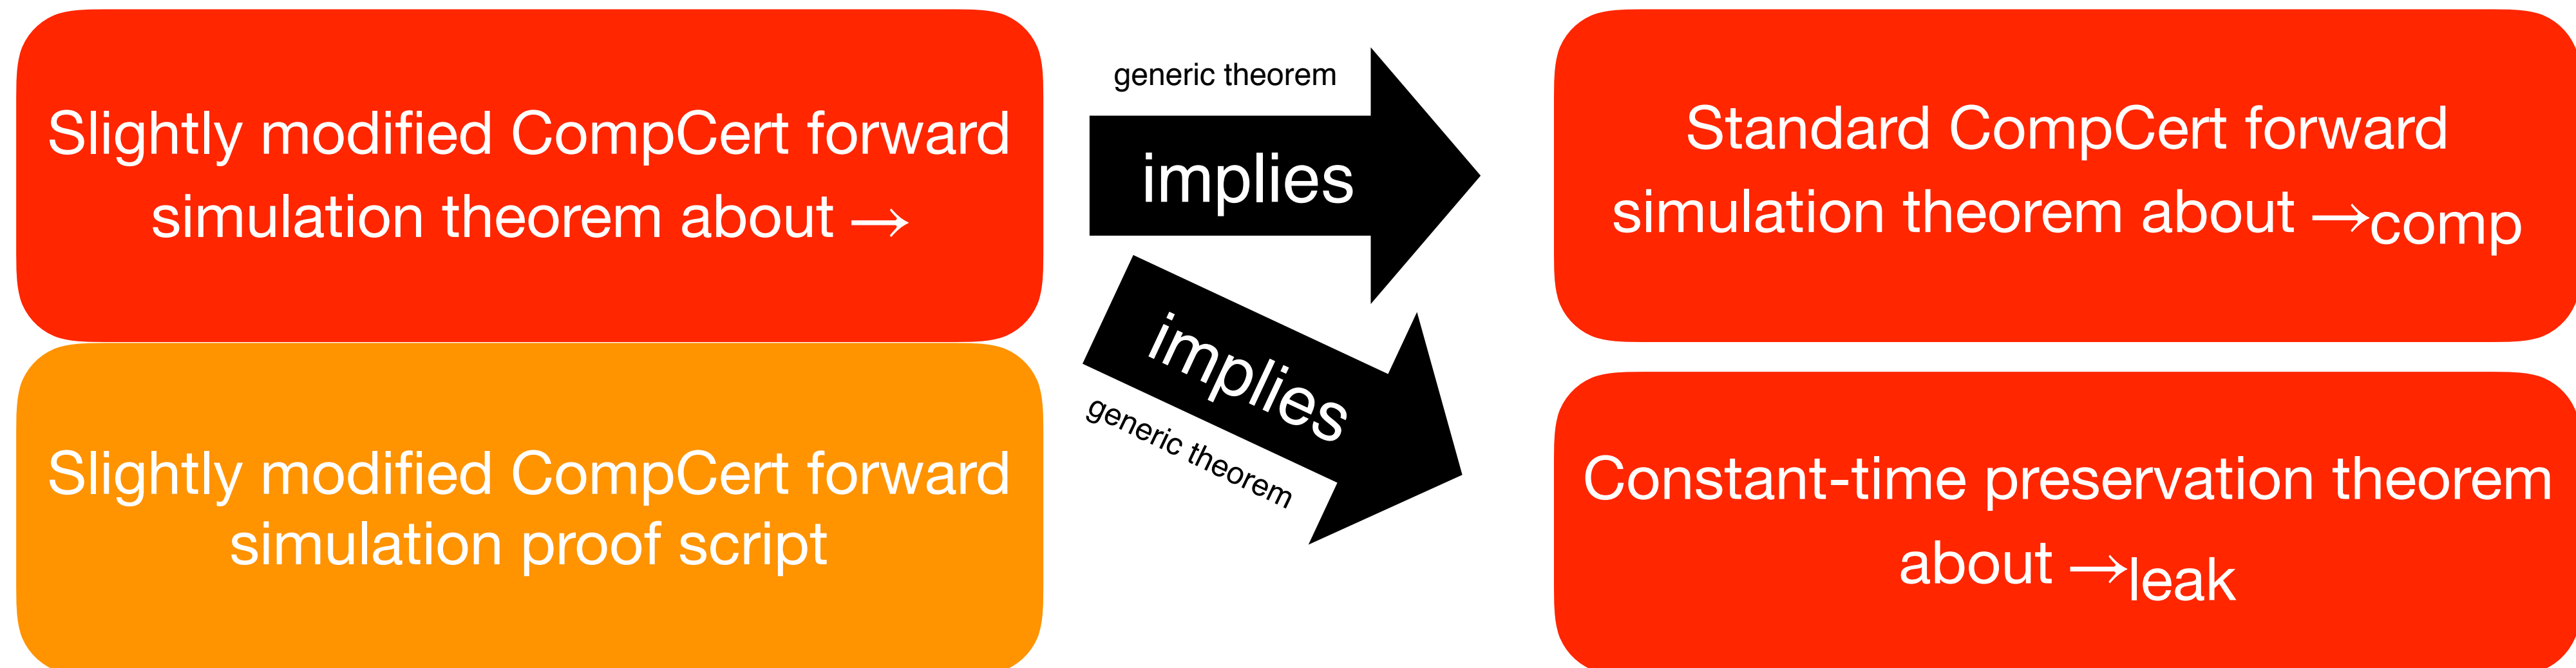
- Cryptographic constant-time preservation is a property about the leak-only semantics $\rightarrow \text{leak}$
- But existing CompCert simulation diagrams deal with the compile-only semantics $\rightarrow \text{comp}$
- Our proof engineering strategy is to benefit as much as possible from the **proof scripts** of these diagrams



Proving cryptographic constant-time preservation

A proof engineering perspective

- Cryptographic constant-time preservation is a property about the leak-only semantics \rightarrow leak
- But existing CompCert simulation diagrams deal with the compile-only semantics \rightarrow comp
- Our proof engineering strategy is to benefit as much as possible from the **proof scripts** of these diagrams



Four proof techniques

- Each technique provides a specific tradeoff between generality and proof tractability
- The first three are slight relaxations of the classical forward diagram and reuse existing scripts

Trace preservation

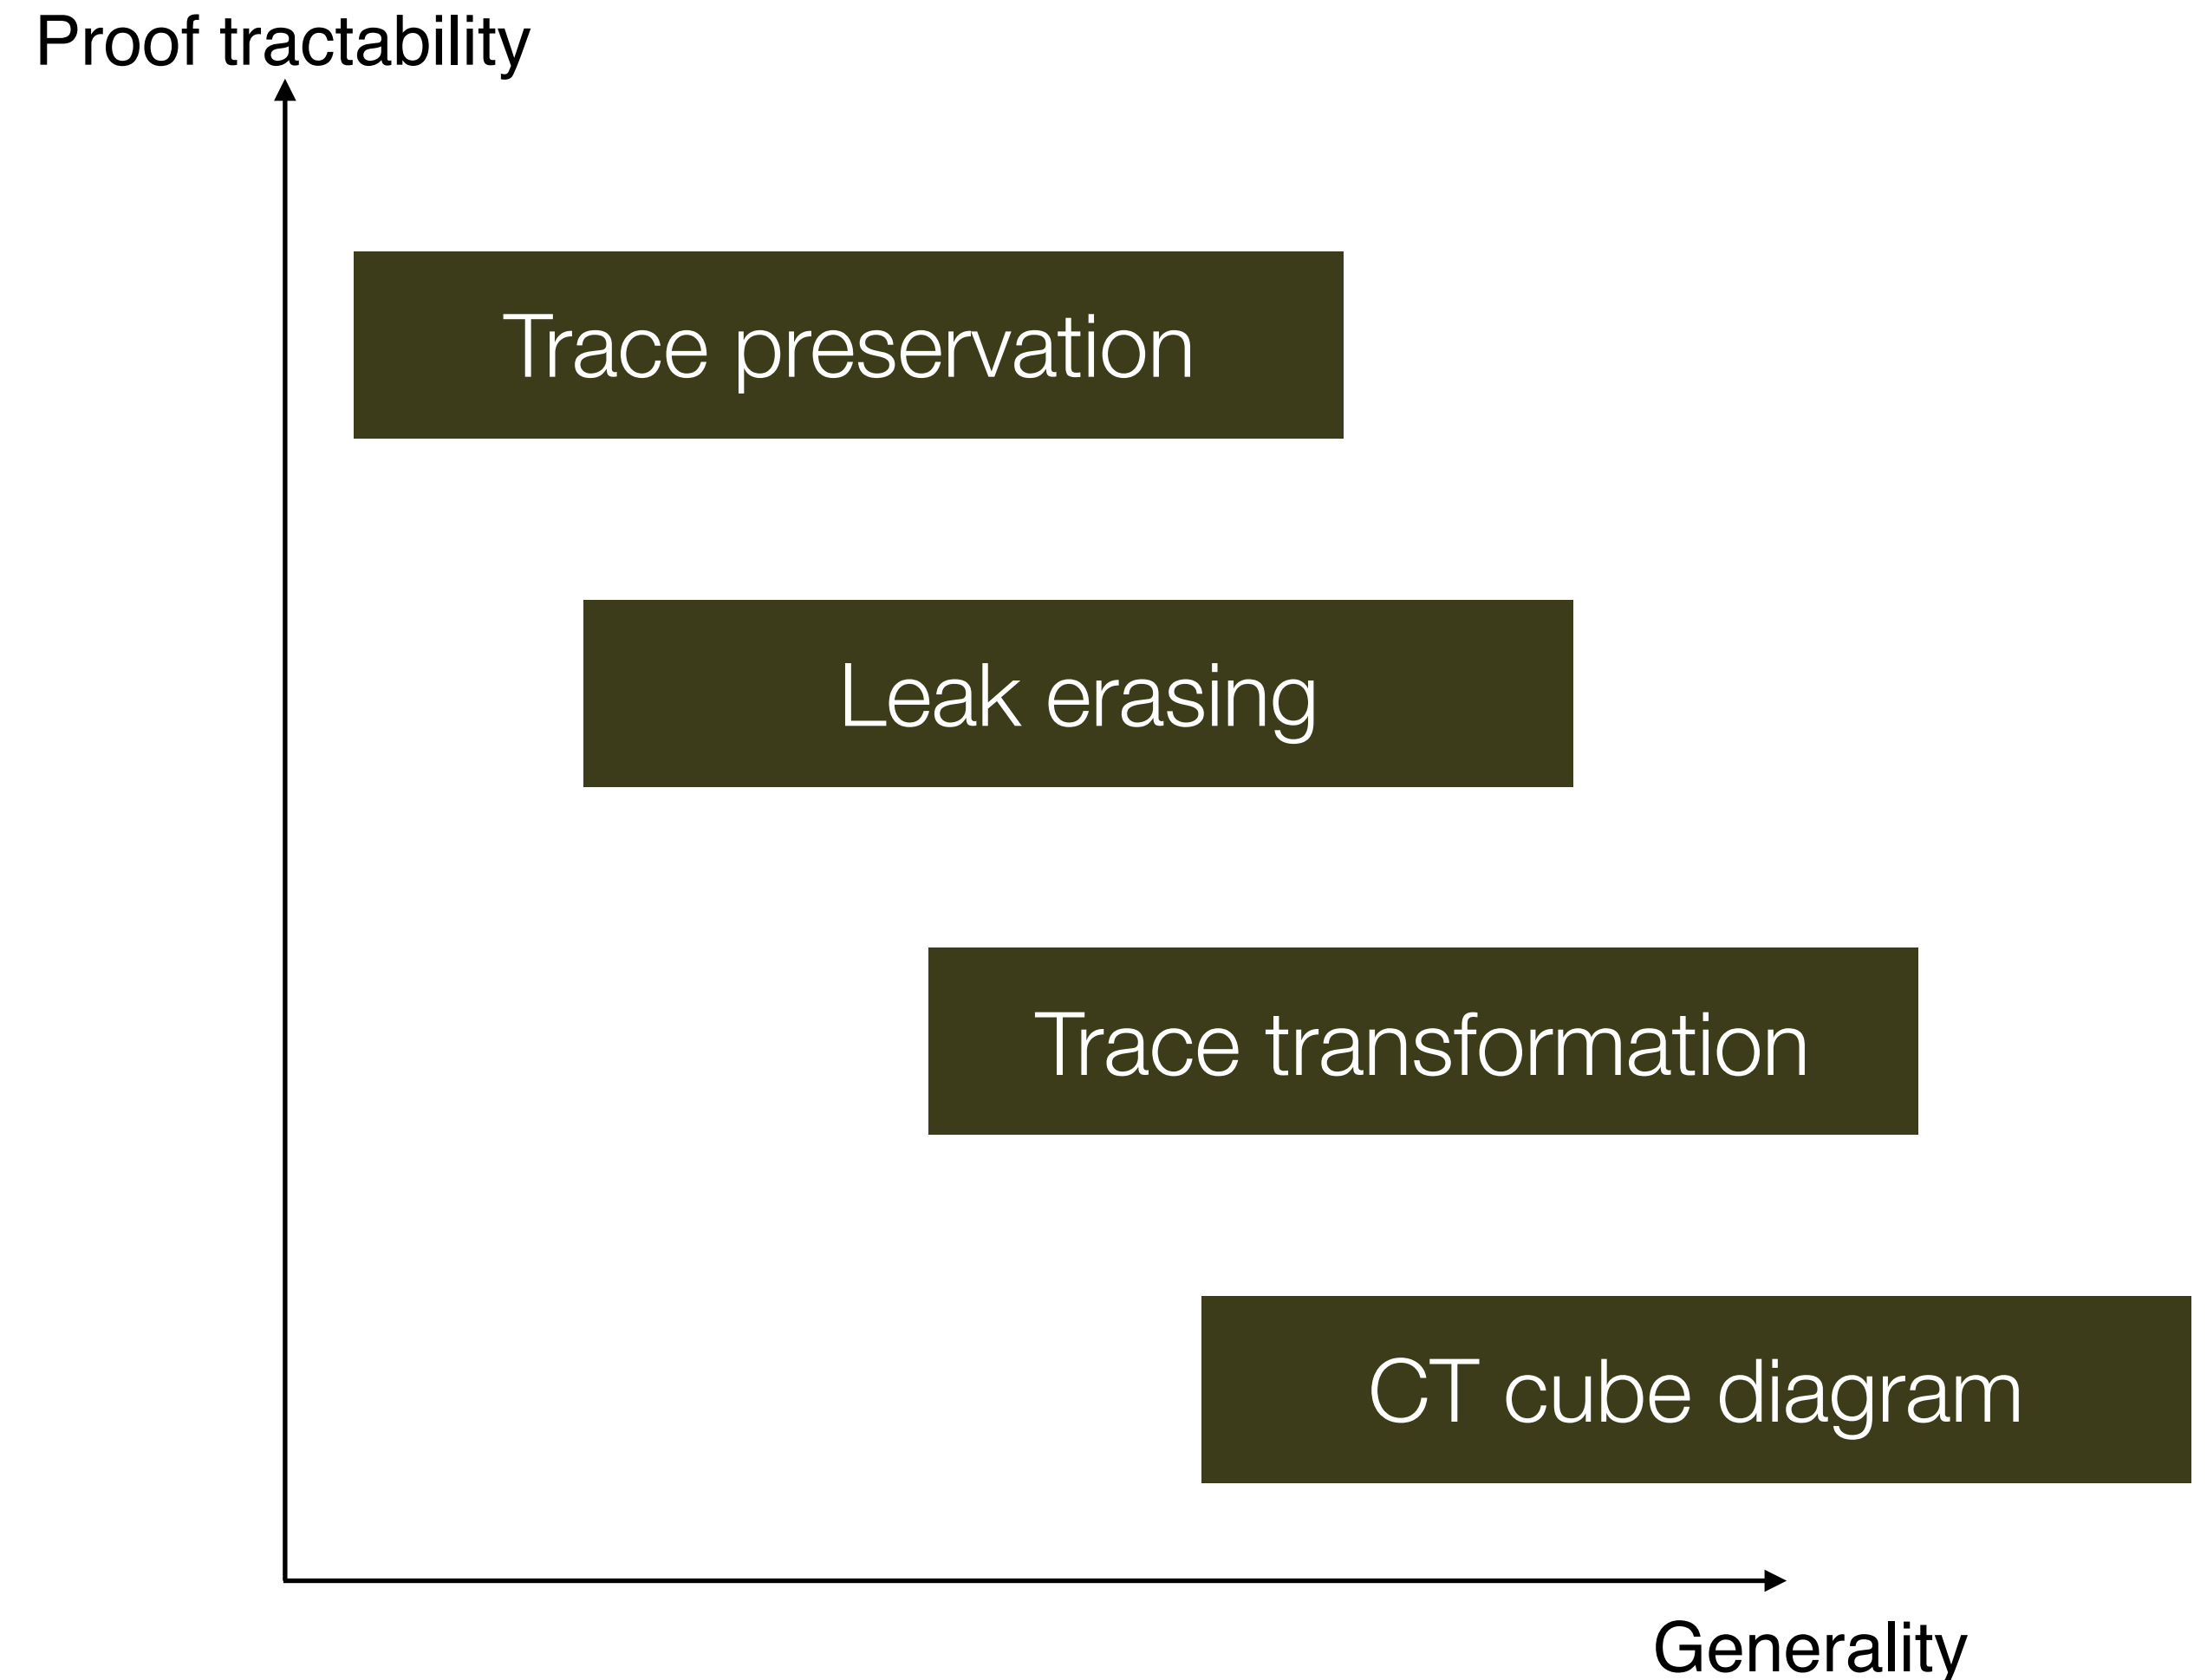
Leak erasing

Trace transformation

CT cube diagram

Four proof techniques

- Each technique provides a specific tradeoff between generality and proof tractability
- The first three are slight relaxations of the classical forward diagram and reuse existing scripts



A palette of proof methods

Trace preservation

Leak erasing

Trace transformation

CT cube diagram

Compiler pass	Diagram used	Explanation on the pass
Cshmgen		Type elaboration, simplification of control
Cminorgen		Stack allocation
Selection		Recognition of operators and addr. modes
RTLgen		Generation of CFG and 3-address code
Tailcall		Tailcall recognition
Inlining		Function inlining
Renumber		Renumbering CFG nodes
ConstProp		Constant propagation
CSE		Common subexpression elimination
Deadcode		Redundancy elimination
Allocation		Register allocation
Tunneling		Branch tunneling
Linearize		Linearization of CFG
CleanupLabels		Removal of unreferenced labels
Debugvar		Synthesis of debugging information
Stacking		Laying out stack frames
Asmgen		Emission of assembly code

A palette of proof methods

- Trace preservation
- Leak erasing
- Trace transformation
- CT cube diagram

6/17

Compiler pass	Diagram used	Explanation on the pass
Cshmgen	Trace preservation	Type elaboration, simplification of control
Cminorgen		Stack allocation
Selection		Recognition of operators and addr. modes
RTLgen	Trace preservation	Generation of CFG and 3-address code
Tailcall	Trace preservation	Tailcall recognition
Inlining		Function inlining
Renumber	Trace preservation	Renumbering CFG nodes
ConstProp		Constant propagation
CSE		Common subexpression elimination
Deadcode		Redundancy elimination
Allocation		Register allocation
Tunneling		Branch tunneling
Linearize		Linearization of CFG
CleanupLabels	Trace preservation	Removal of unreferenced labels
Debugvar	Trace preservation	Synthesis of debugging information
Stacking		Laying out stack frames
Asmgen		Emission of assembly code

A palette of proof methods

- Trace preservation
- Leak erasing
- Trace transformation
- CT cube diagram

6/17

5/17

Compiler pass	Diagram used	Explanation on the pass
Cshmgen	Trace preservation	Type elaboration, simplification of control
Cminorgen		Stack allocation
Selection	Leak erasing	Recognition of operators and addr. modes
RTLgen	Trace preservation	Generation of CFG and 3-address code
Tailcall	Trace preservation	Tailcall recognition
Inlining		Function inlining
Renumber	Trace preservation	Renumbering CFG nodes
ConstProp		Constant propagation
CSE	Leak erasing	Common subexpression elimination
Deadcode	Leak erasing	Redundancy elimination
Allocation	Leak erasing	Register allocation
Tunneling	Leak erasing	Branch tunneling
Linearize		Linearization of CFG
CleanupLabels	Trace preservation	Removal of unreferenced labels
Debugvar	Trace preservation	Synthesis of debugging information
Stacking		Laying out stack frames
Asmgen		Emission of assembly code

A palette of proof methods

Trace preservation
Leak erasing
Trace transformation
CT cube diagram

6/17
5/17
5/17

Compiler pass	Diagram used	Explanation on the pass
Cshmgen	Trace preservation	Type elaboration, simplification of control
Cminorgen	Trace transformation	Stack allocation
Selection	Leak erasing	Recognition of operators and addr. modes
RTLgen	Trace preservation	Generation of CFG and 3-address code
Tailcall	Trace preservation	Tailcall recognition
Inlining	Trace transformation	Function inlining
Renumber	Trace preservation	Renumbering CFG nodes
ConstProp	Trace transformation	Constant propagation
CSE	Leak erasing	Common subexpression elimination
Deadcode	Leak erasing	Redundancy elimination
Allocation	Leak erasing	Register allocation
Tunneling	Leak erasing	Branch tunneling
Linearize		Linearization of CFG
CleanupLabels	Trace preservation	Removal of unreferenced labels
Debugvar	Trace preservation	Synthesis of debugging information
Stacking	Trace transformation	Laying out stack frames
Asmggen	Trace transformation	Emission of assembly code

A palette of proof methods

		Compiler pass	Diagram used	Explanation on the pass
Trace preservation	6/17	Cshmgen	Trace preservation	Type elaboration, simplification of control
		Cminorgen	Trace transformation	Stack allocation
		Selection	Leak erasing	Recognition of operators and addr. modes
		RTLgen	Trace preservation	Generation of CFG and 3-address code
		Tailcall	Trace preservation	Tailcall recognition
		Inlining	Trace transformation	Function inlining
Leak erasing	5/17	Renumber	Trace preservation	Renumbering CFG nodes
		ConstProp	Trace transformation	Constant propagation
Trace transformation	5/17	CSE	Leak erasing	Common subexpression elimination
		Deadcode	Leak erasing	Redundancy elimination
CT cube diagram	1/17	Allocation	Leak erasing	Register allocation
		Tunneling	Leak erasing	Branch tunneling
		Linearize	CT cube diagram	Linearization of CFG
		CleanupLabels	Trace preservation	Removal of unreferenced labels
		Debugvar	Trace preservation	Synthesis of debugging information
		Stacking	Trace transformation	Laying out stack frames
		Asmgen	Trace transformation	Emission of assembly code

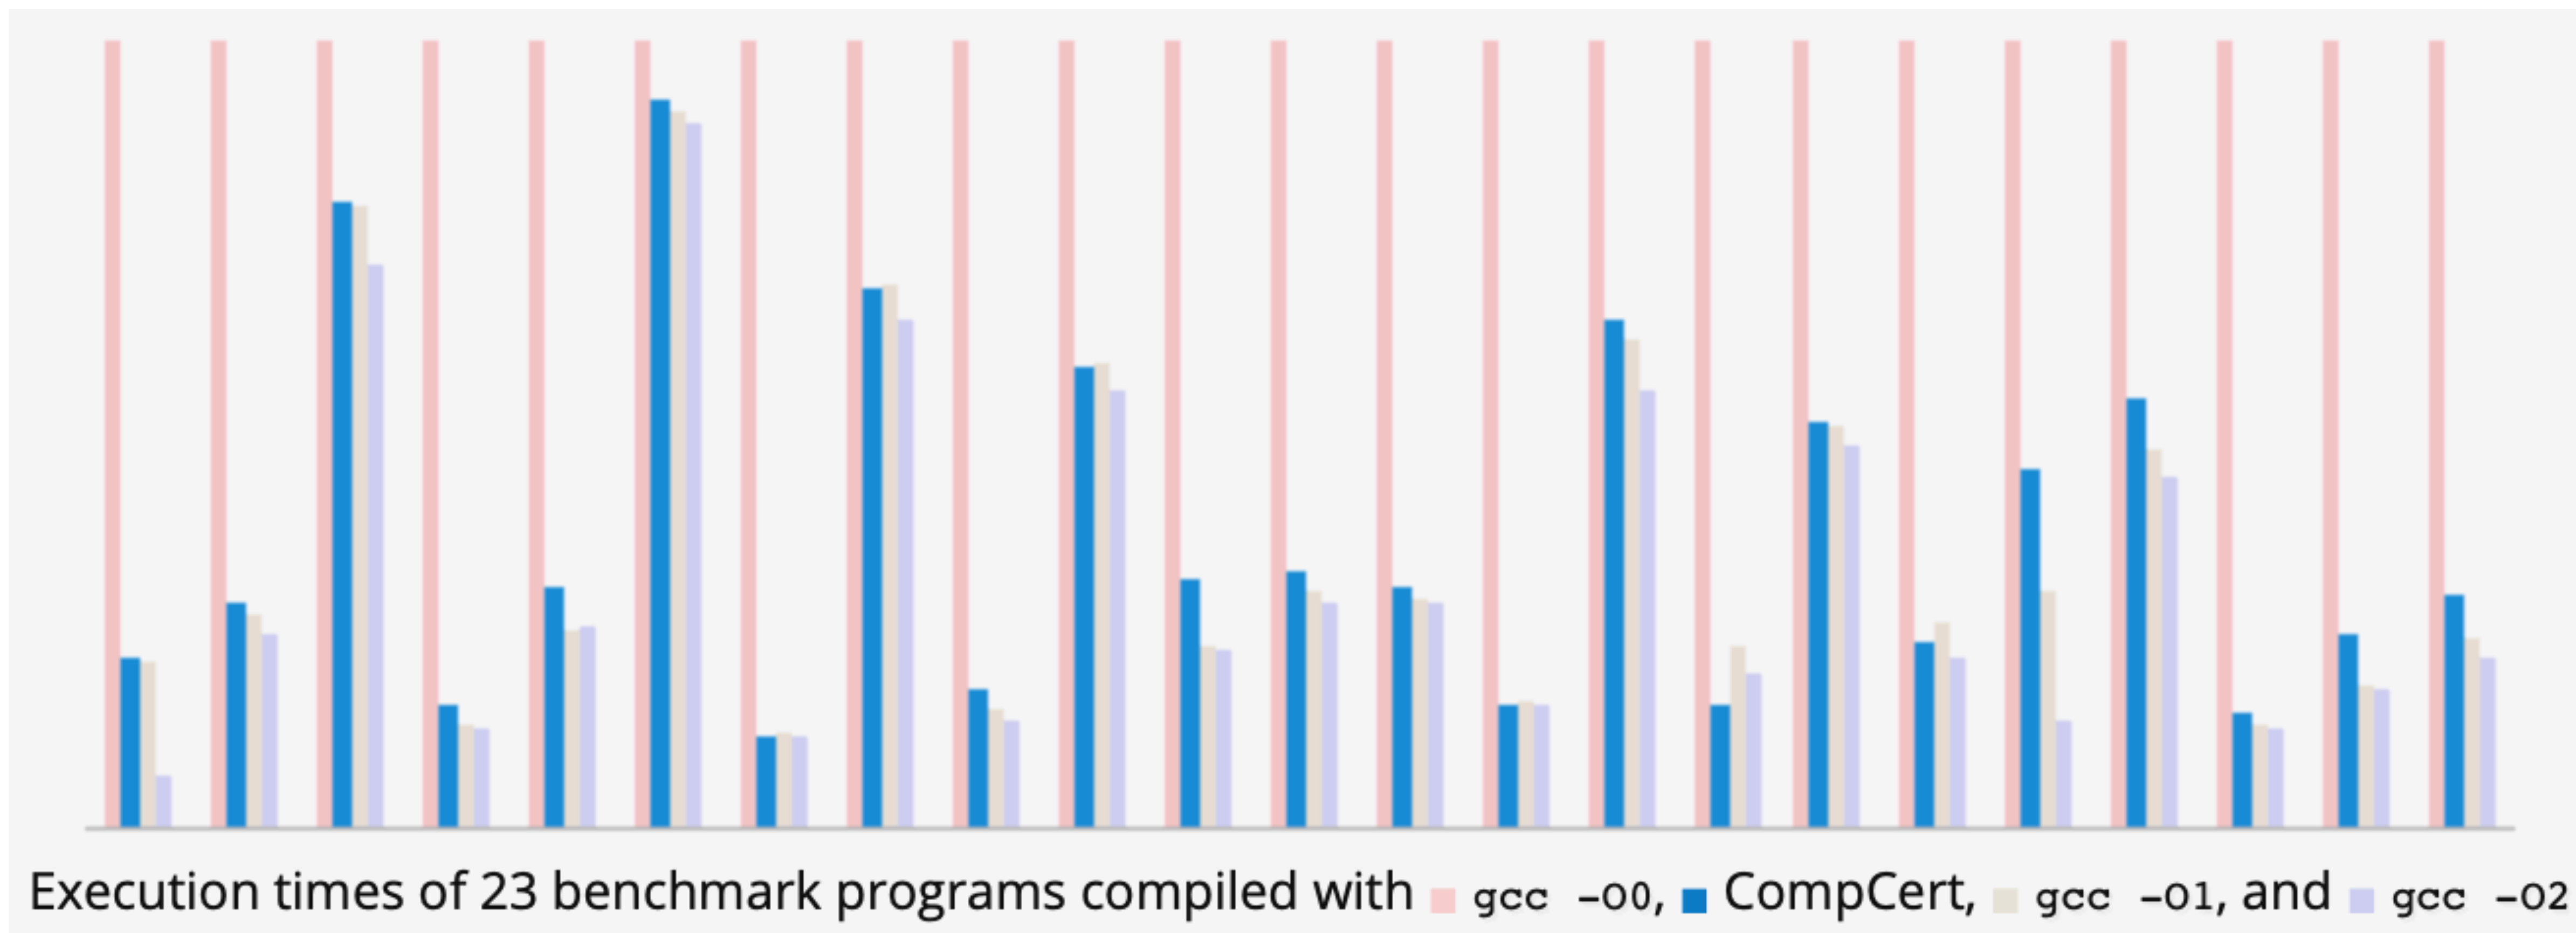
Conclusion and perspectives

Conclusion

- A machine checked-proof that a mildly modified version of the CompCert compiler preserves cryptographic constant-time
- A carefully crafted methodology that maximises proof reuse

Perspectives

- Make CompCert generate more efficient code for crypto programs (e.g. using SIMD instructions)
- Explore other observational information-flow policies and adapt CompCert



<https://www.absint.com/compcert/>

