

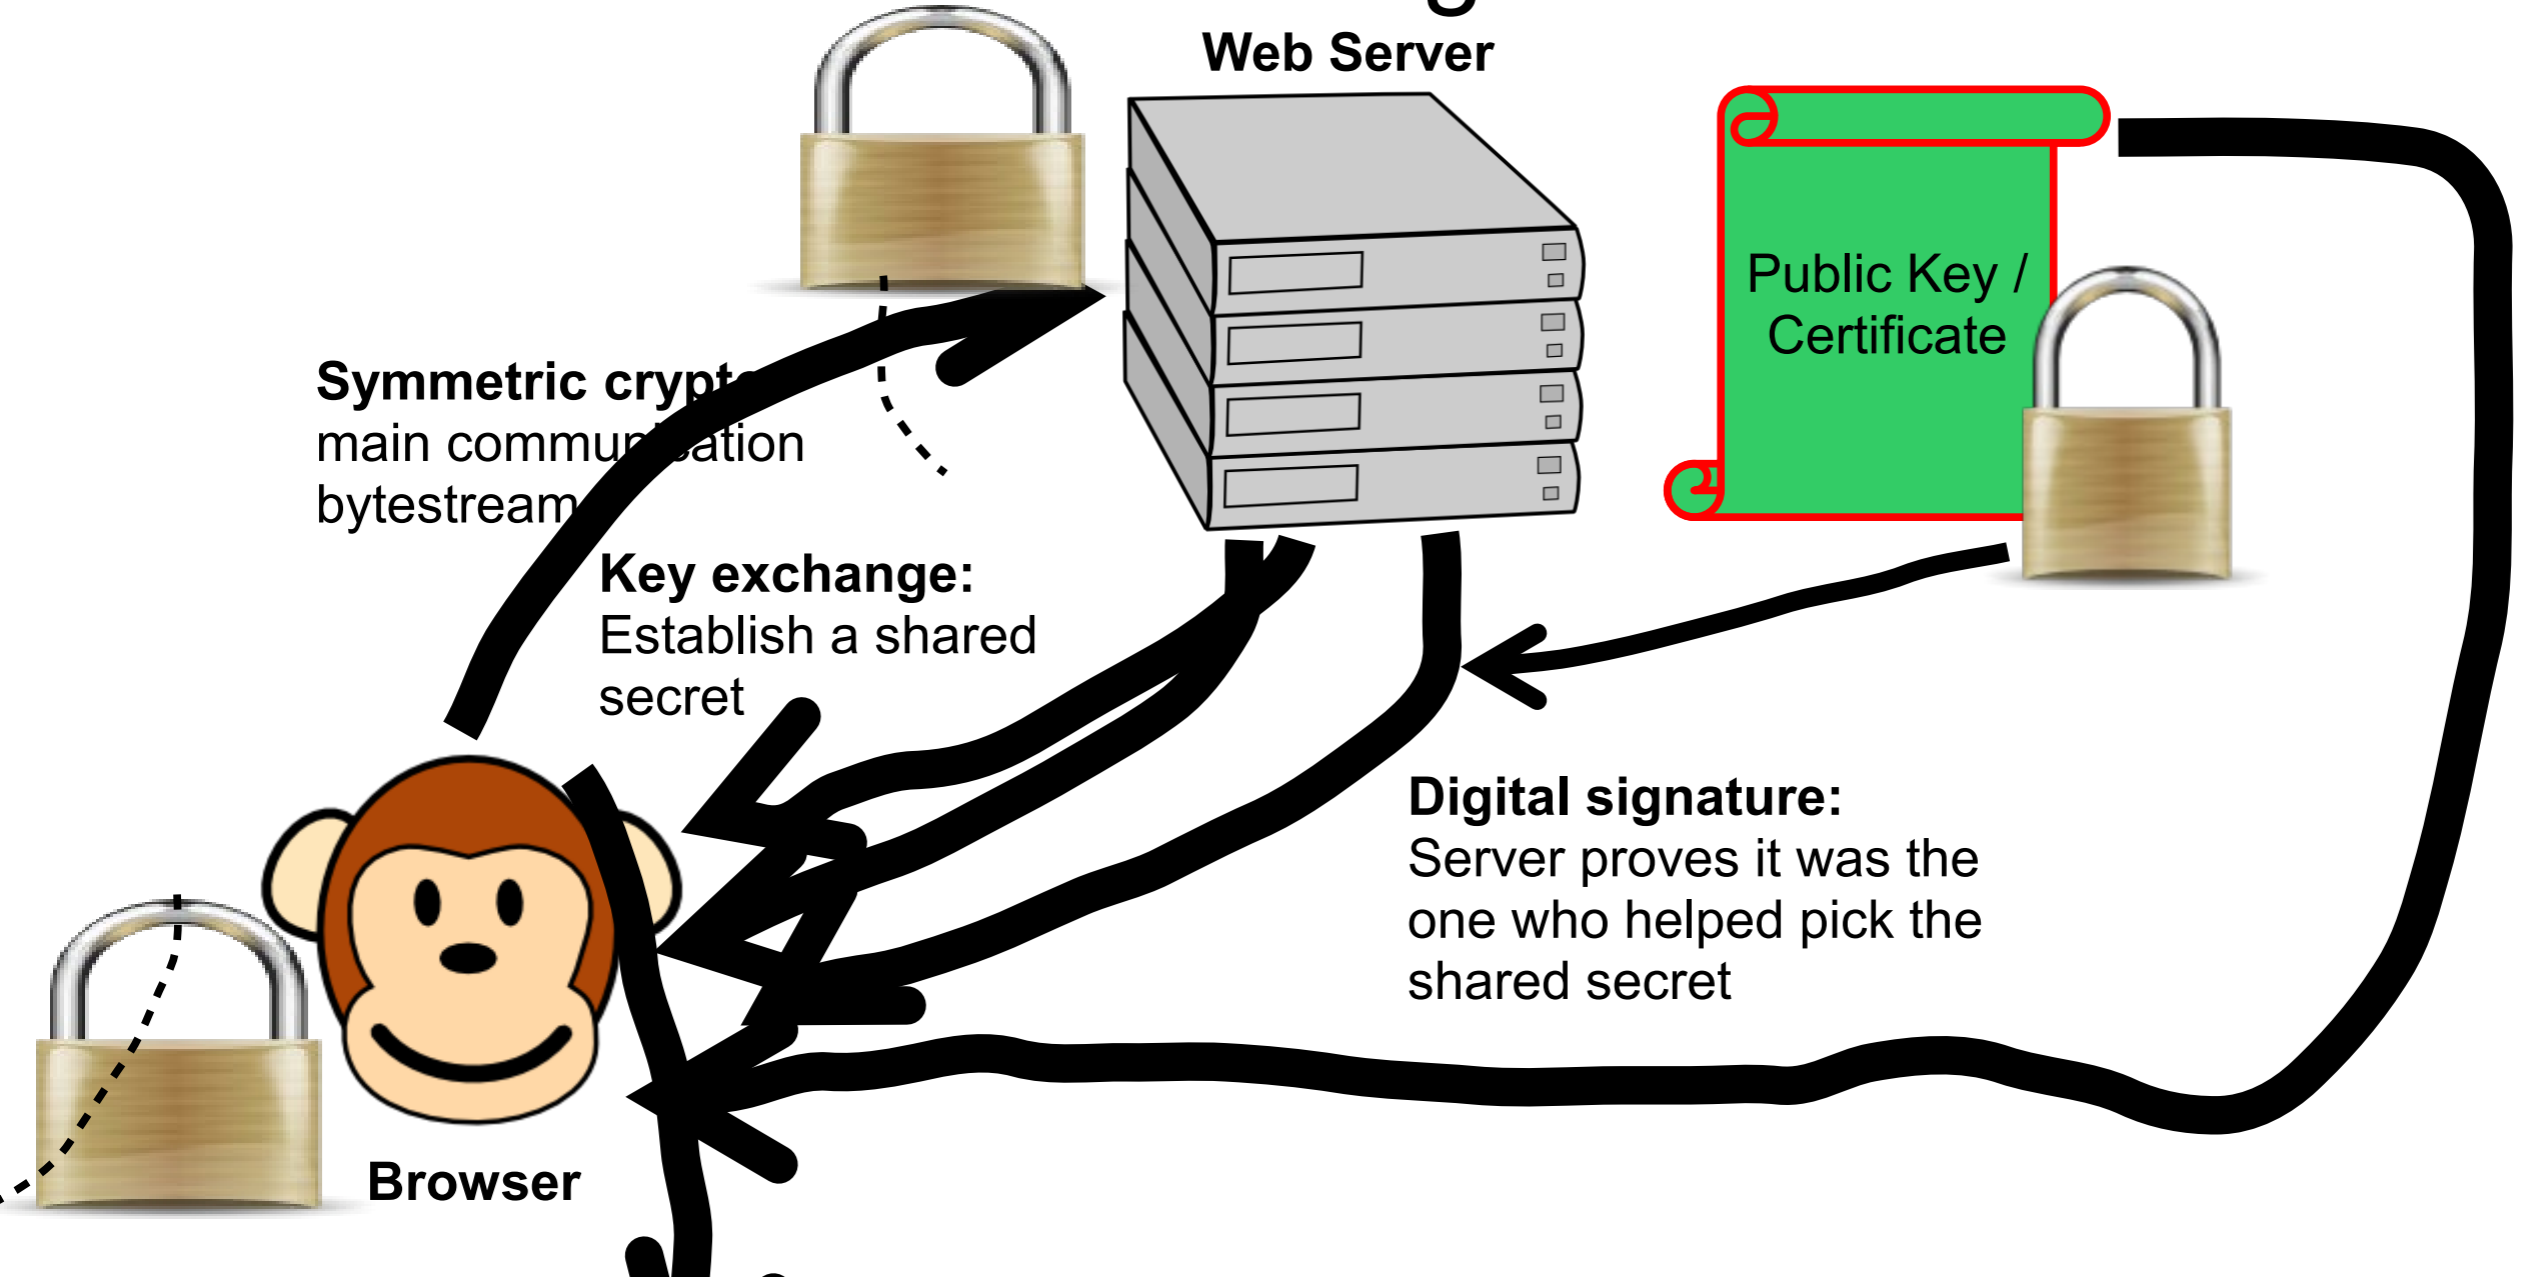
Correct-by-Construction Cryptography Without Performance Compromises

Adam Chlipala, MIT CSAIL
NUS Computer Science Research Week
January 2022

Joint work with:

Joonwon Choi, Andres Erbsen, Jason Gross, Jade Philipoom, Robert Sloan, Clark Wood


Web Browsing with SSL



About the First Two Stages (Public-Key Crypto)

- Public-key stages only run once per session, but, with many small HTTPS connections common in practice, their performance is still important.
- Balancing correctness and performance is also more challenging for the public-key algorithms.

But the experts know how to do all this, right?




Labor-intensive adaptation, with each combination taking significant expert effort.


We introduced *Fiat Cryptography*.

- An automatic generator for this kind of code,
- with correctness proofs in the Coq theorem prover.
- Adopted for small but important parts of TLS implementations in both Chrome and Firefox, plus a number of blockchain systems, etc.

The Real World
(maintainers of OpenSSL,
etc., live here)



The Ivory Tower
(formal-verification
experts live here)



Outline

- Catching up: formal verification in the 21st century
- More specific project motivation
- Classic Fiat Cryptography
- Towards correct-by-construction cryptographic appliances

Catching up: formal verification in the 21st century

Debugging: The Secret Essence of Programming

“By June 1949 people had begun to realize that it was not so easy to get programs right as at one time appeared.

[...] the realization came over me with full force that a good part of the remainder of my life was going to be spent in finding errors in my own programs.”

Maurice Wilkes, *Memoirs of a Computer Pioneer*, MIT Press, 1985, p. 145.



Crucial Substitutions

Debugging
exploring concrete executions



Proving
exploring symbolic arguments

Testing
describing concrete scenarios



Specifying
describing general requirements

Auditing code
algorithms in detail




Auditing specs
functionality without optimizations

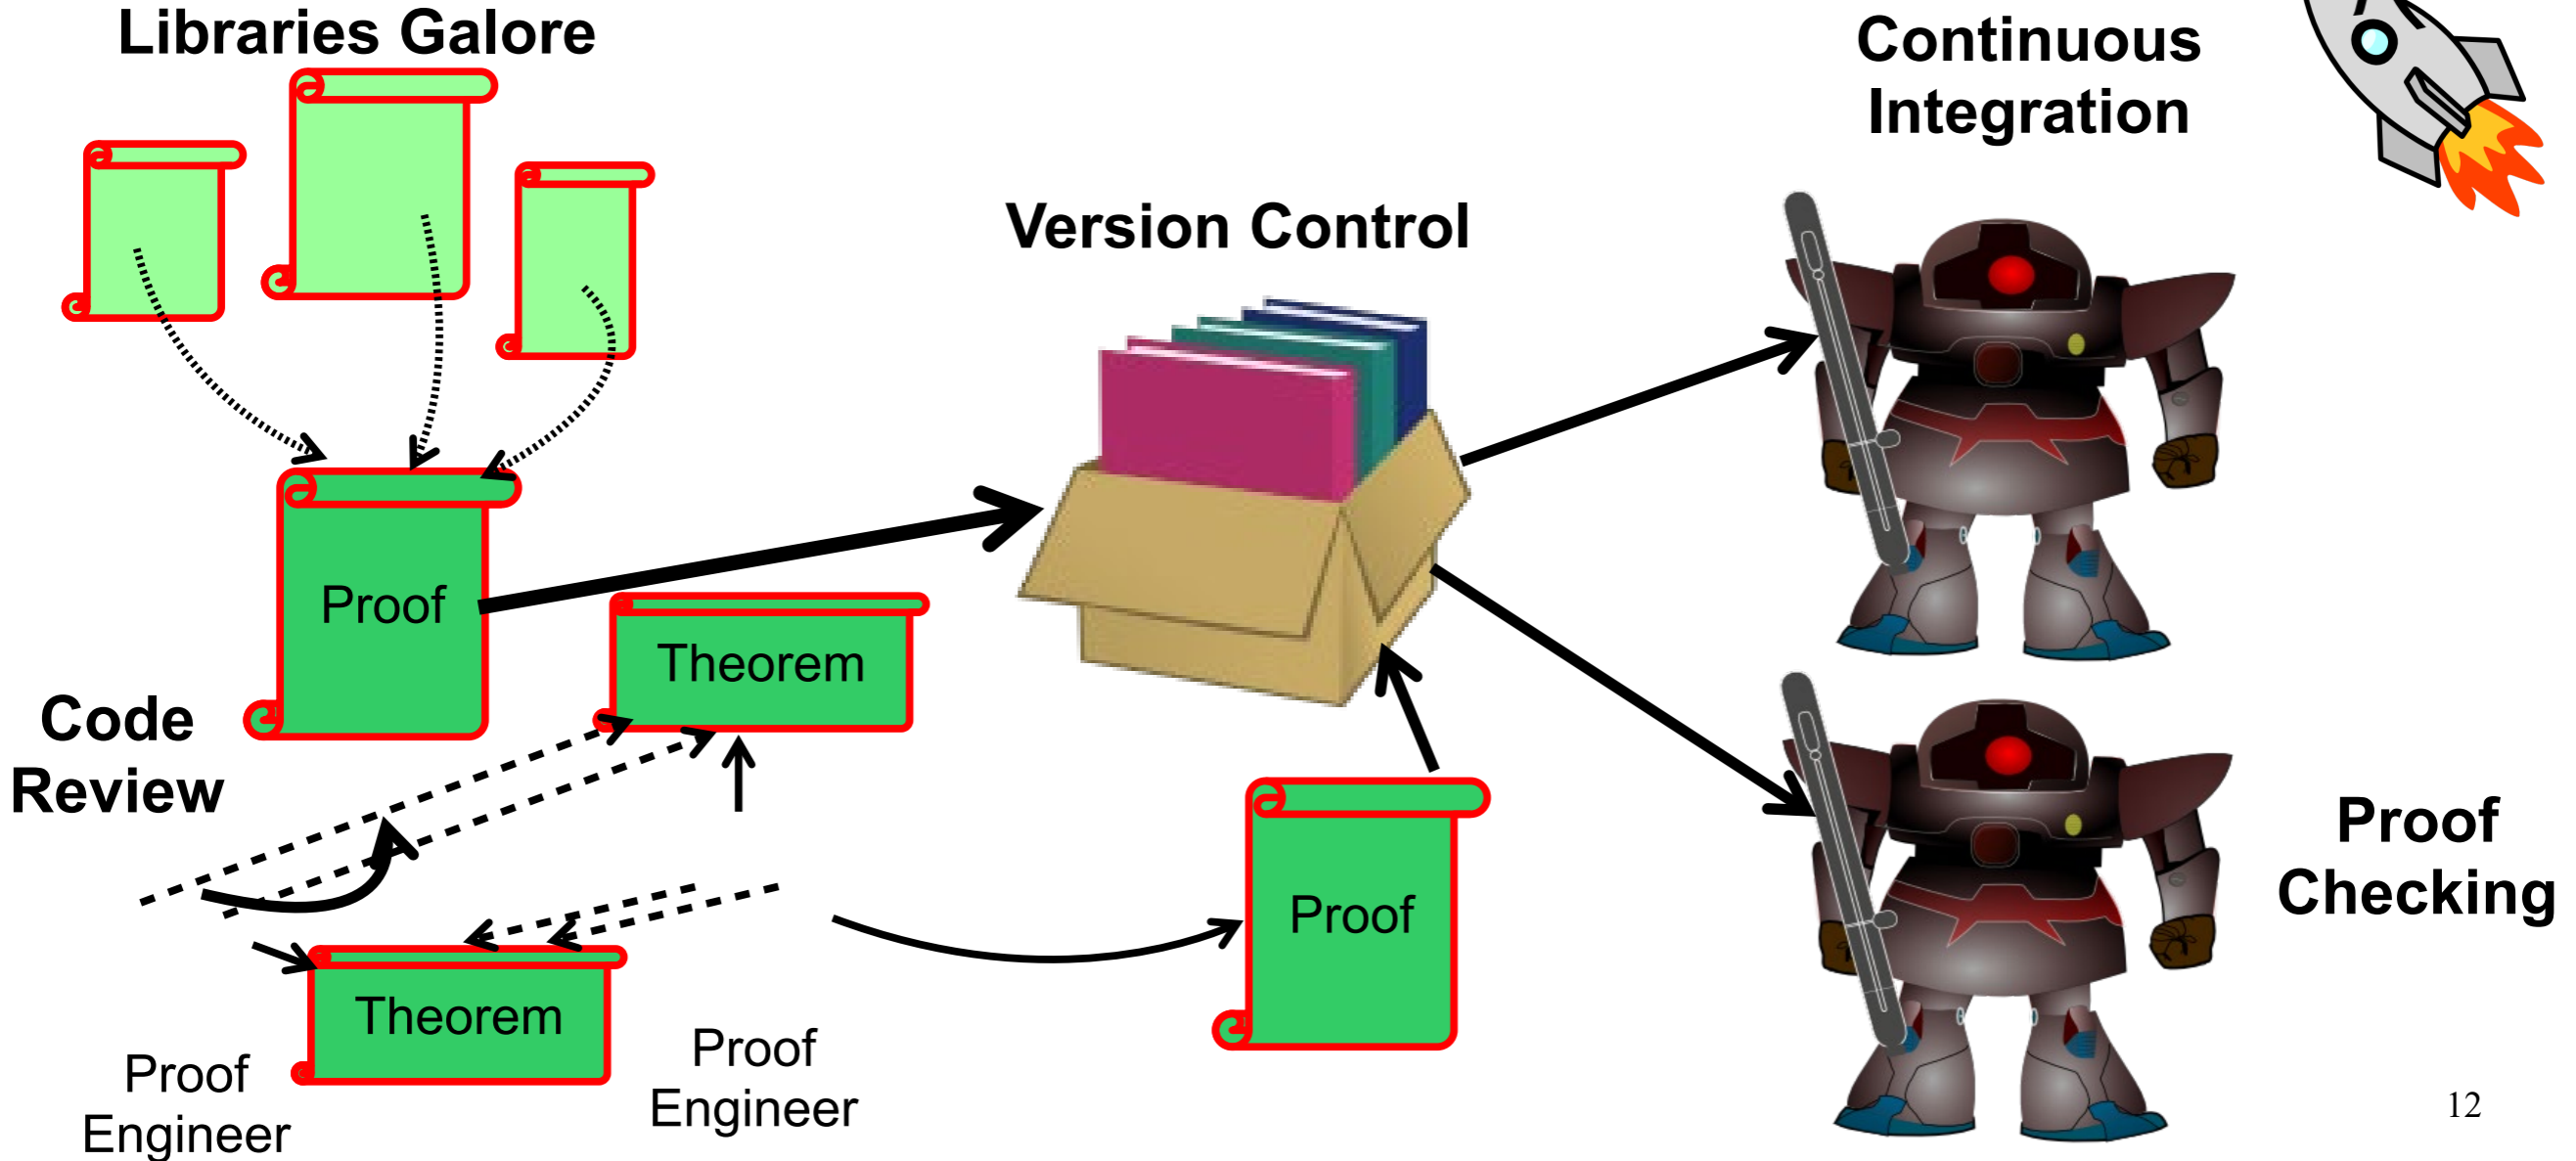

Q: Aren't These Proofs Too Boring for Mortals?

It is argued that formal verifications of programs, no matter how obtained, will not play the same key role in the development of computer science and software engineering as proofs do in mathematics. Furthermore the absence of continuity, the inevitability of change, and the complexity of specification of significantly many real programs make the formal verification process difficult to justify and manage.

– De Millo, Lipton, and Perlis,
“Social Processes and Proofs of
Theorems and Programs,” CACM, 1979




The Proof Workflow of the Future



Proof Assistant

noun: a software package essentially providing an
integrated development environment (IDE)
for stating and proving mathematical theorems
where writing proofs takes human effort
but **checking proofs is automatic**

The Most Popular Proof Assistants



Isabelle/HOL

One well-known application:



Verified microkernel OS

Coq

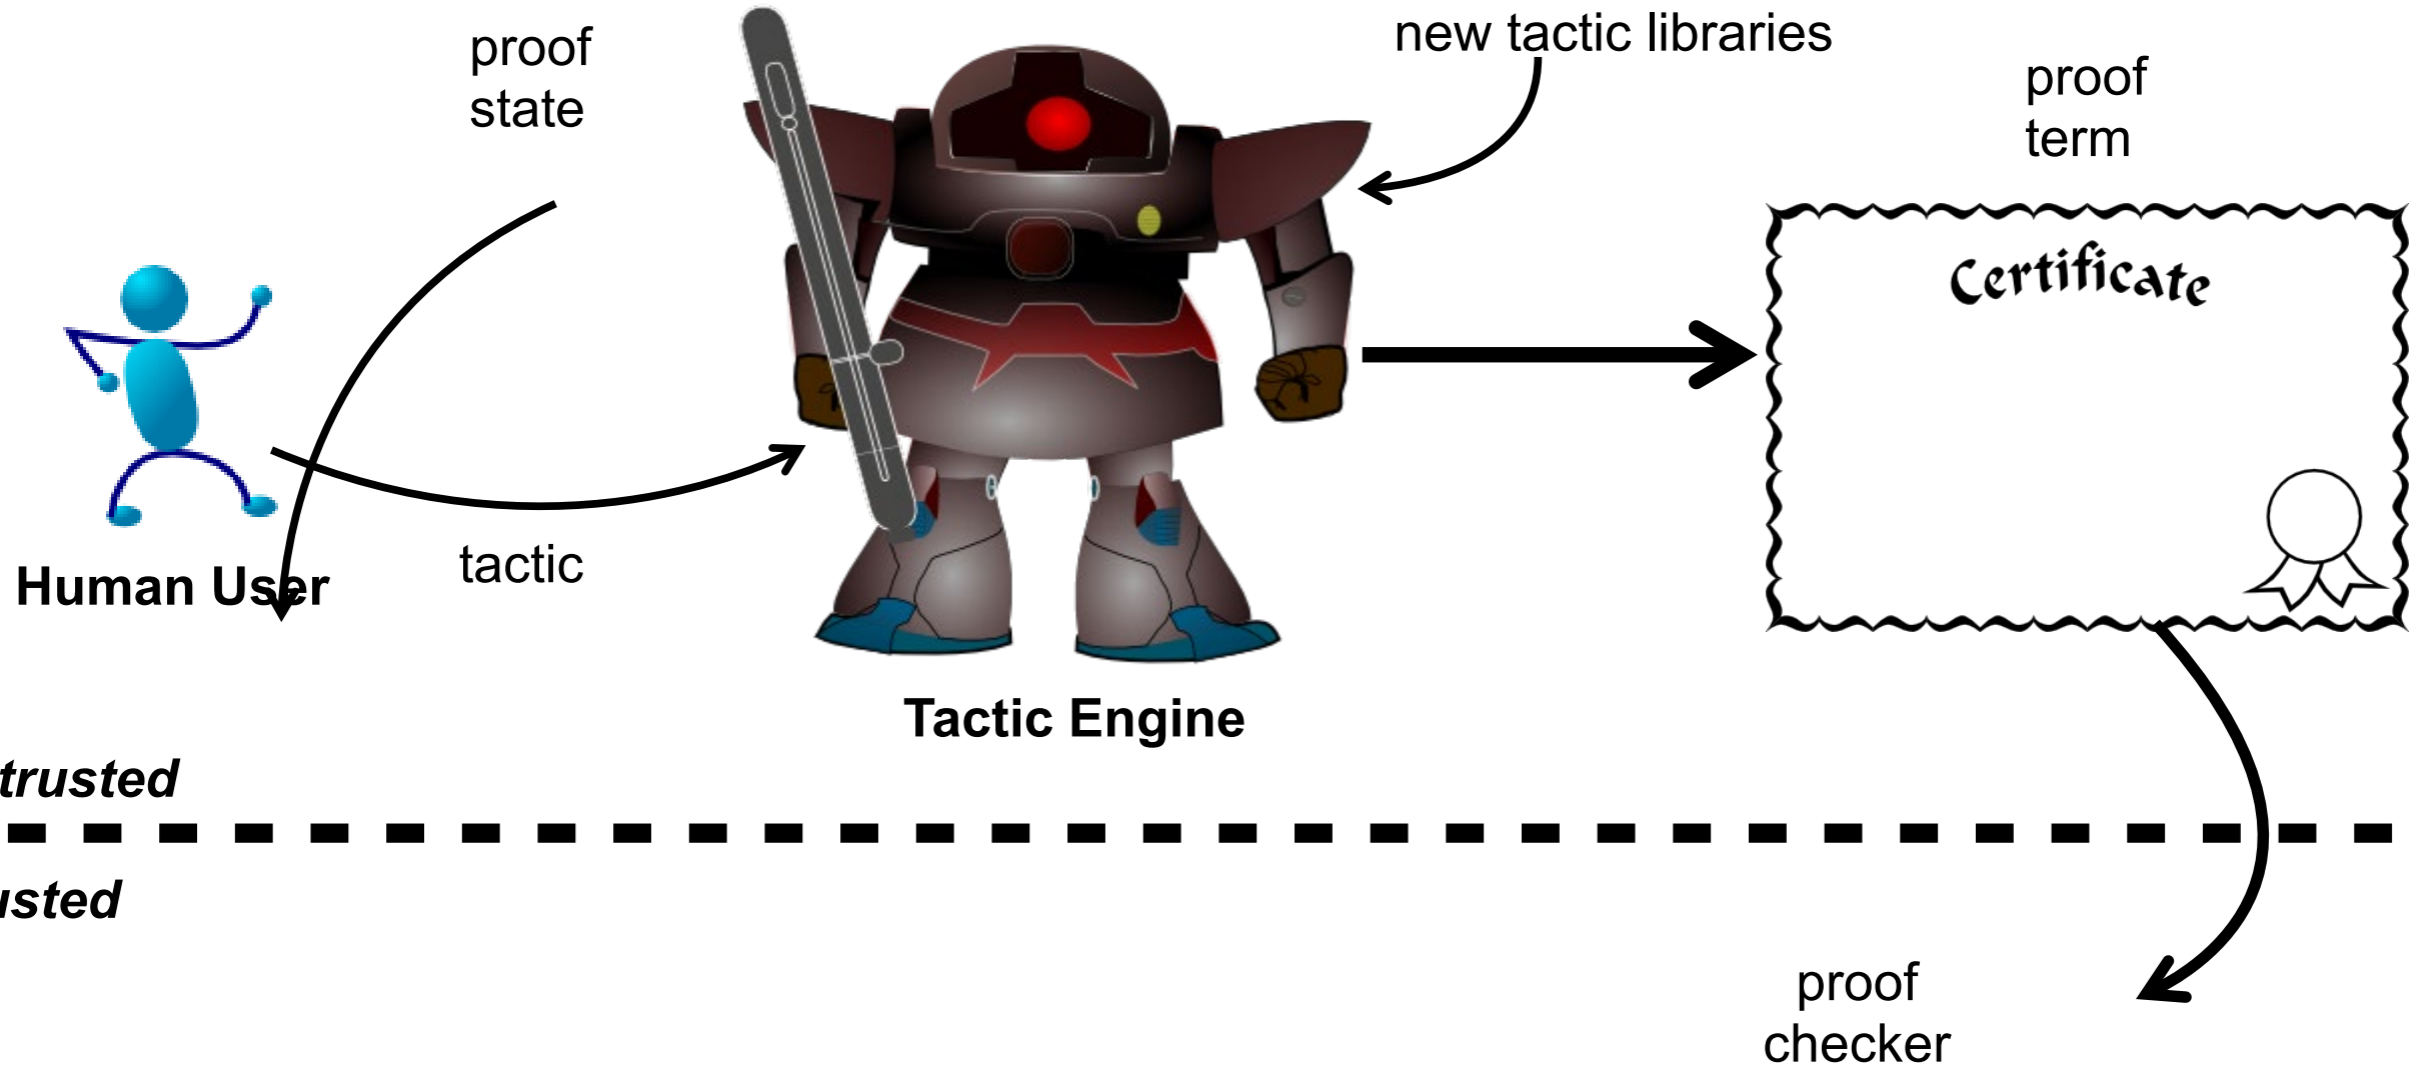
One well-known application:

CompCert



Verified C compiler

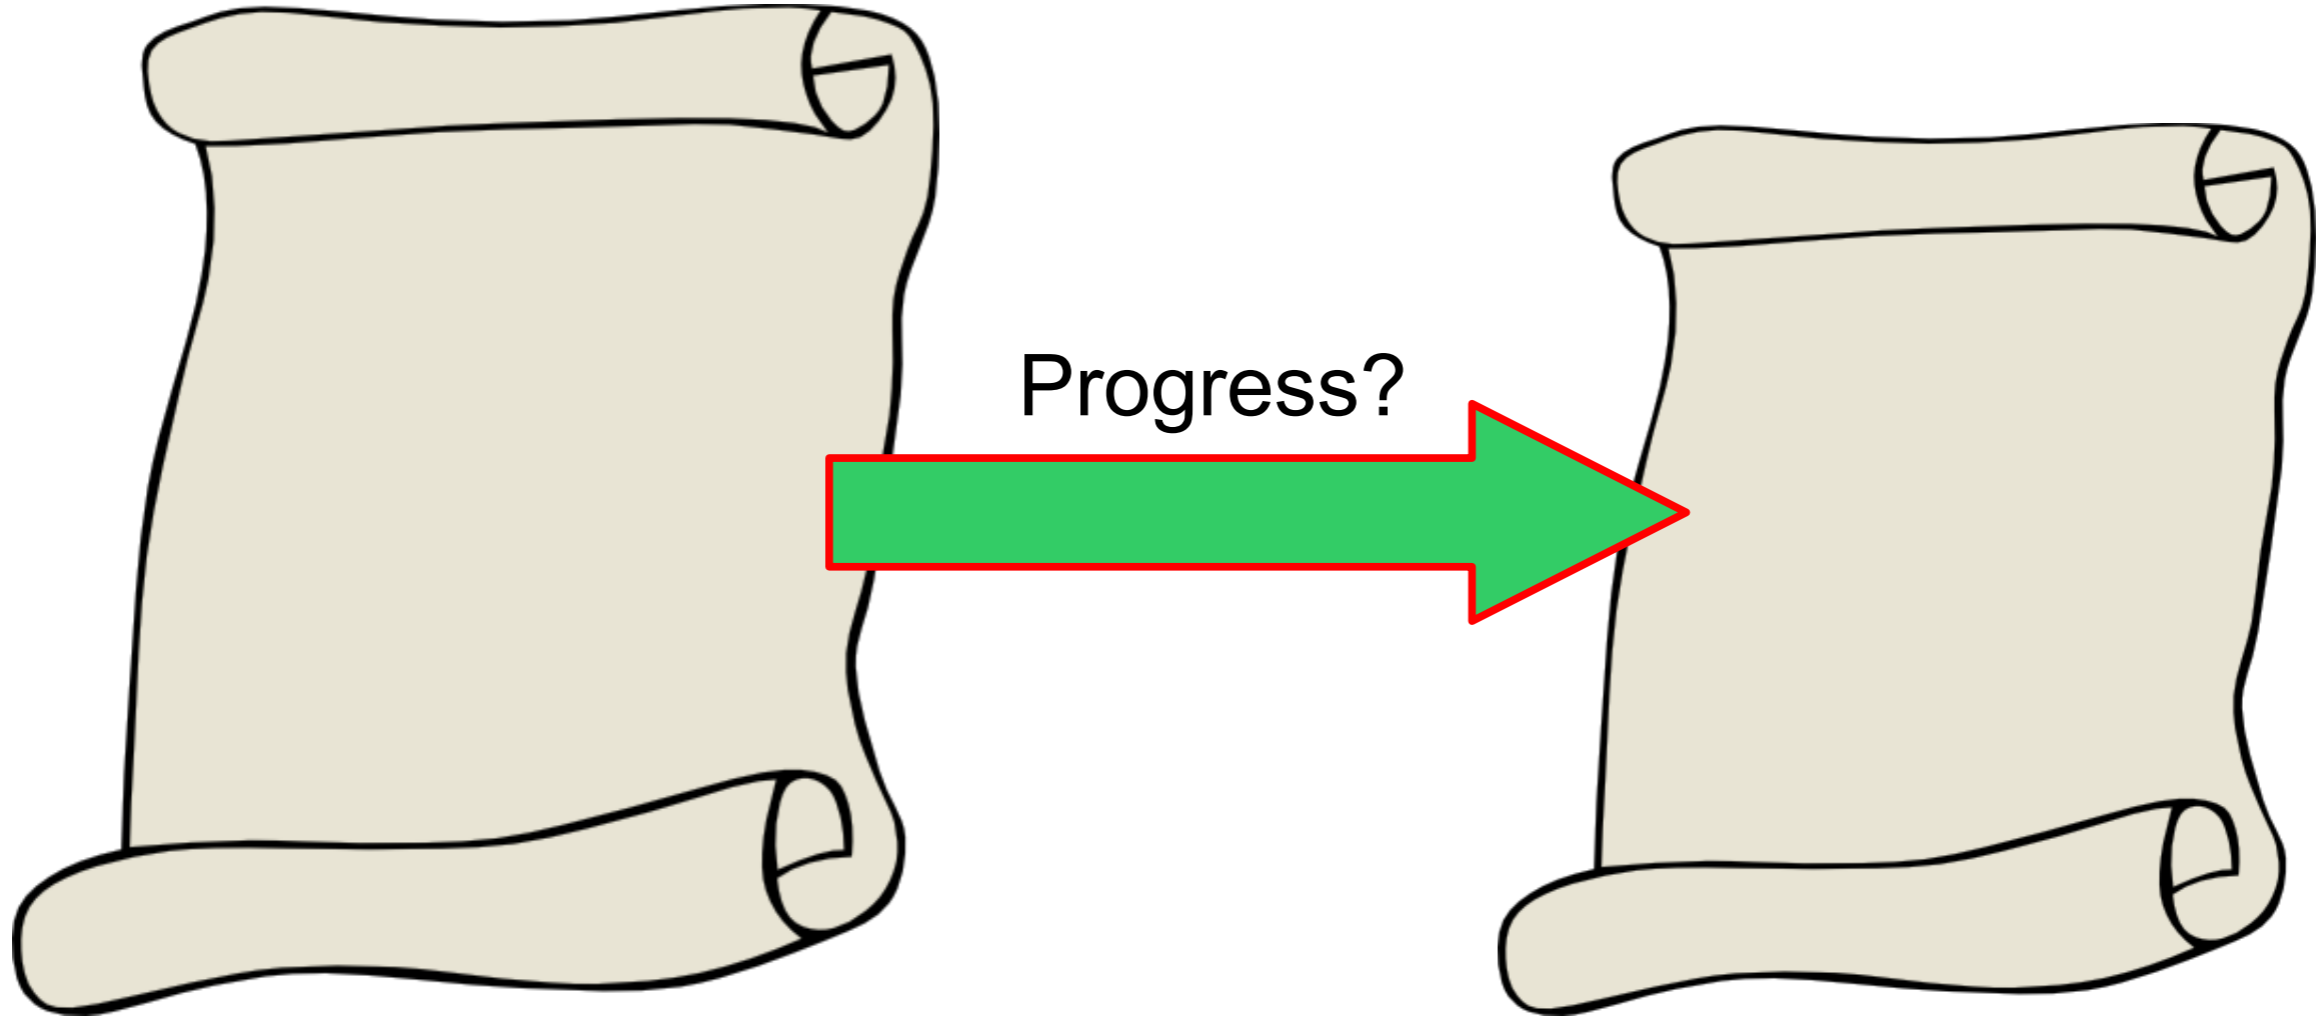
Why Should the [Machine|Human] Trust the [Human|Machine]?



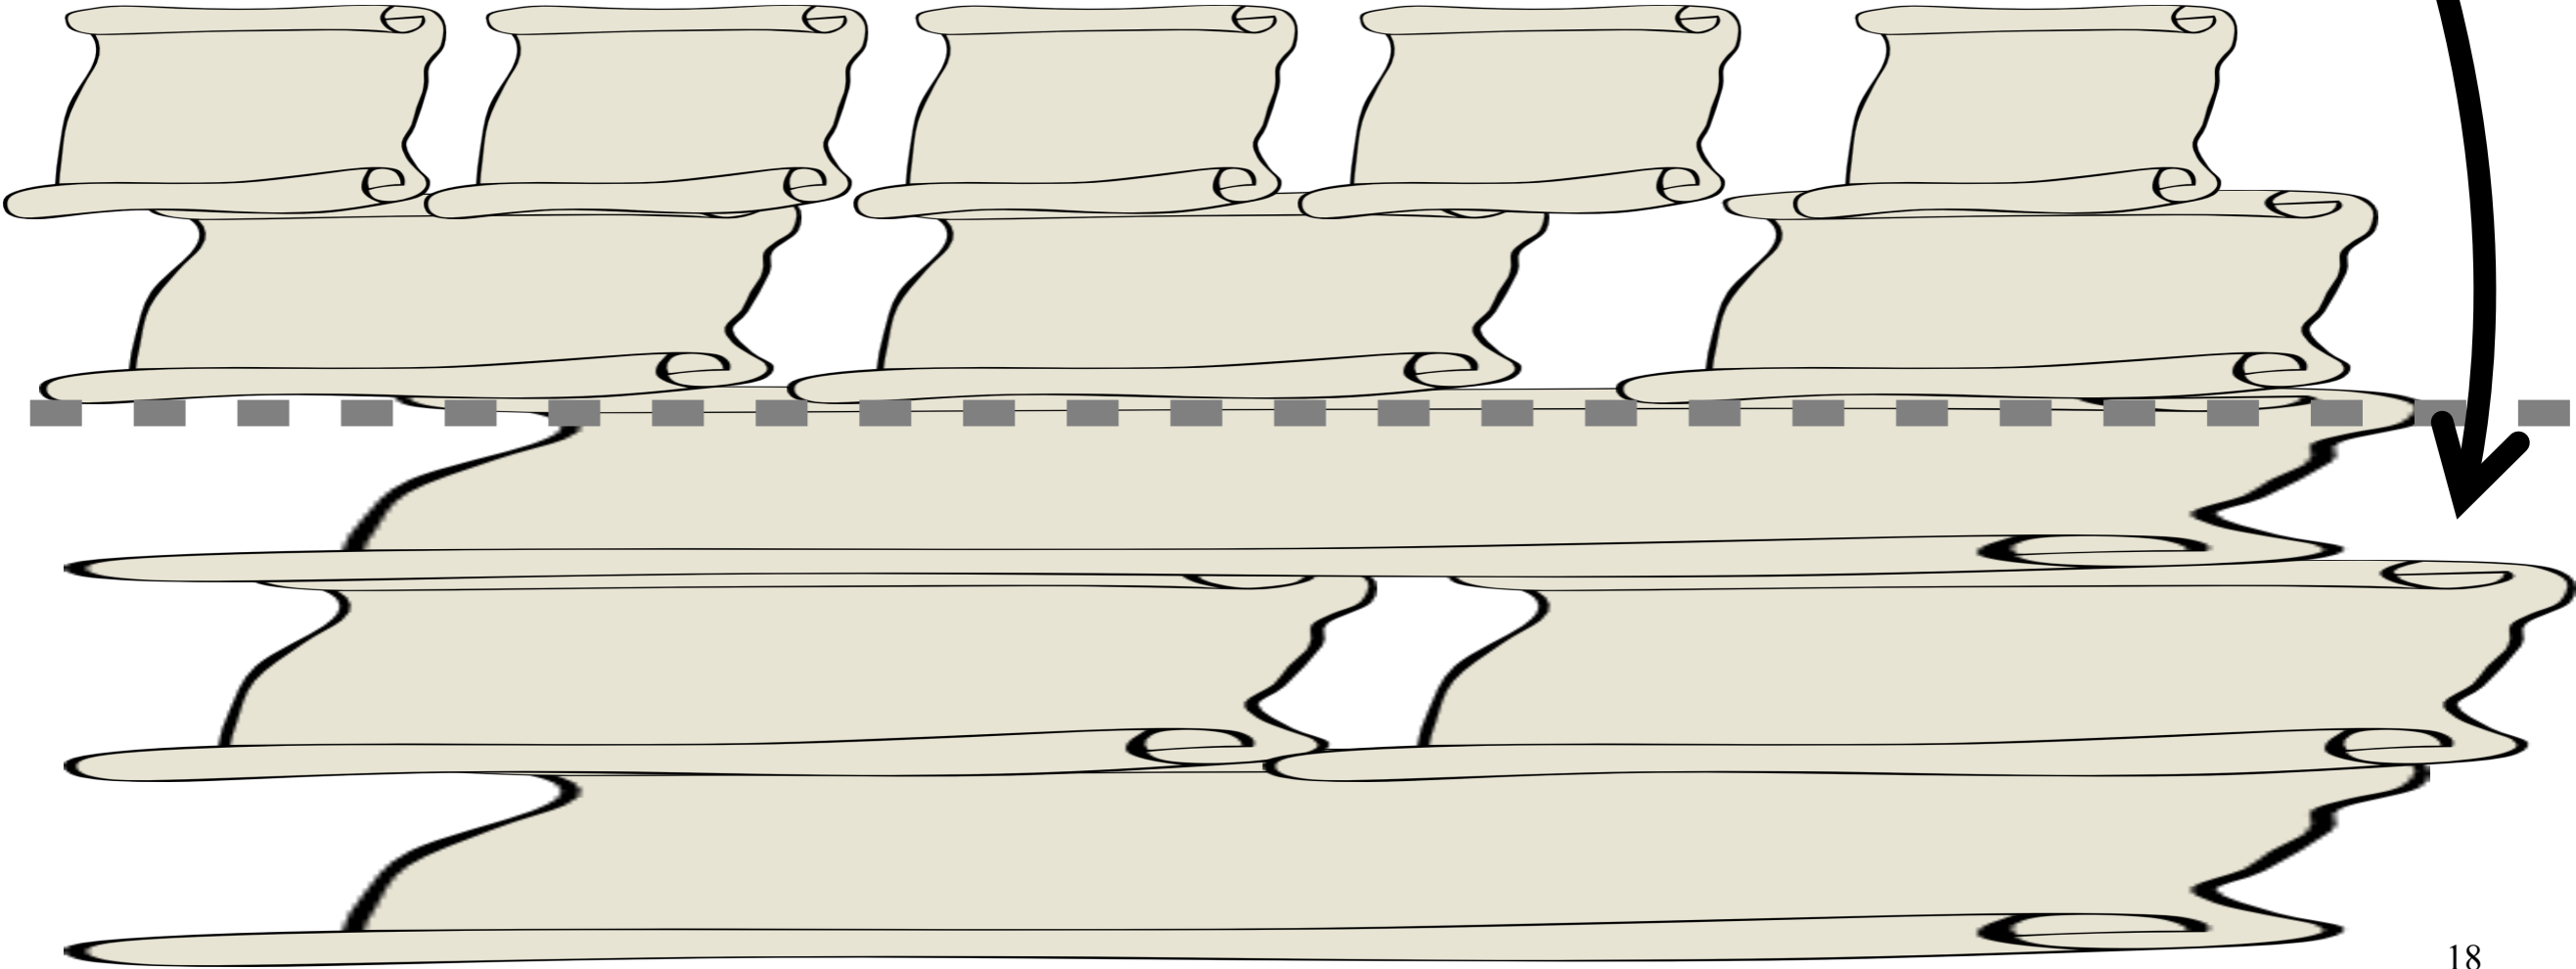
Demo

Some simple proofs in Coq

Q: Isn't It (About) As Hard to Get Specs Right?



A: Focus Spec-Writing on Systems Infrastructure



An Approximate Truth About Software

Spec

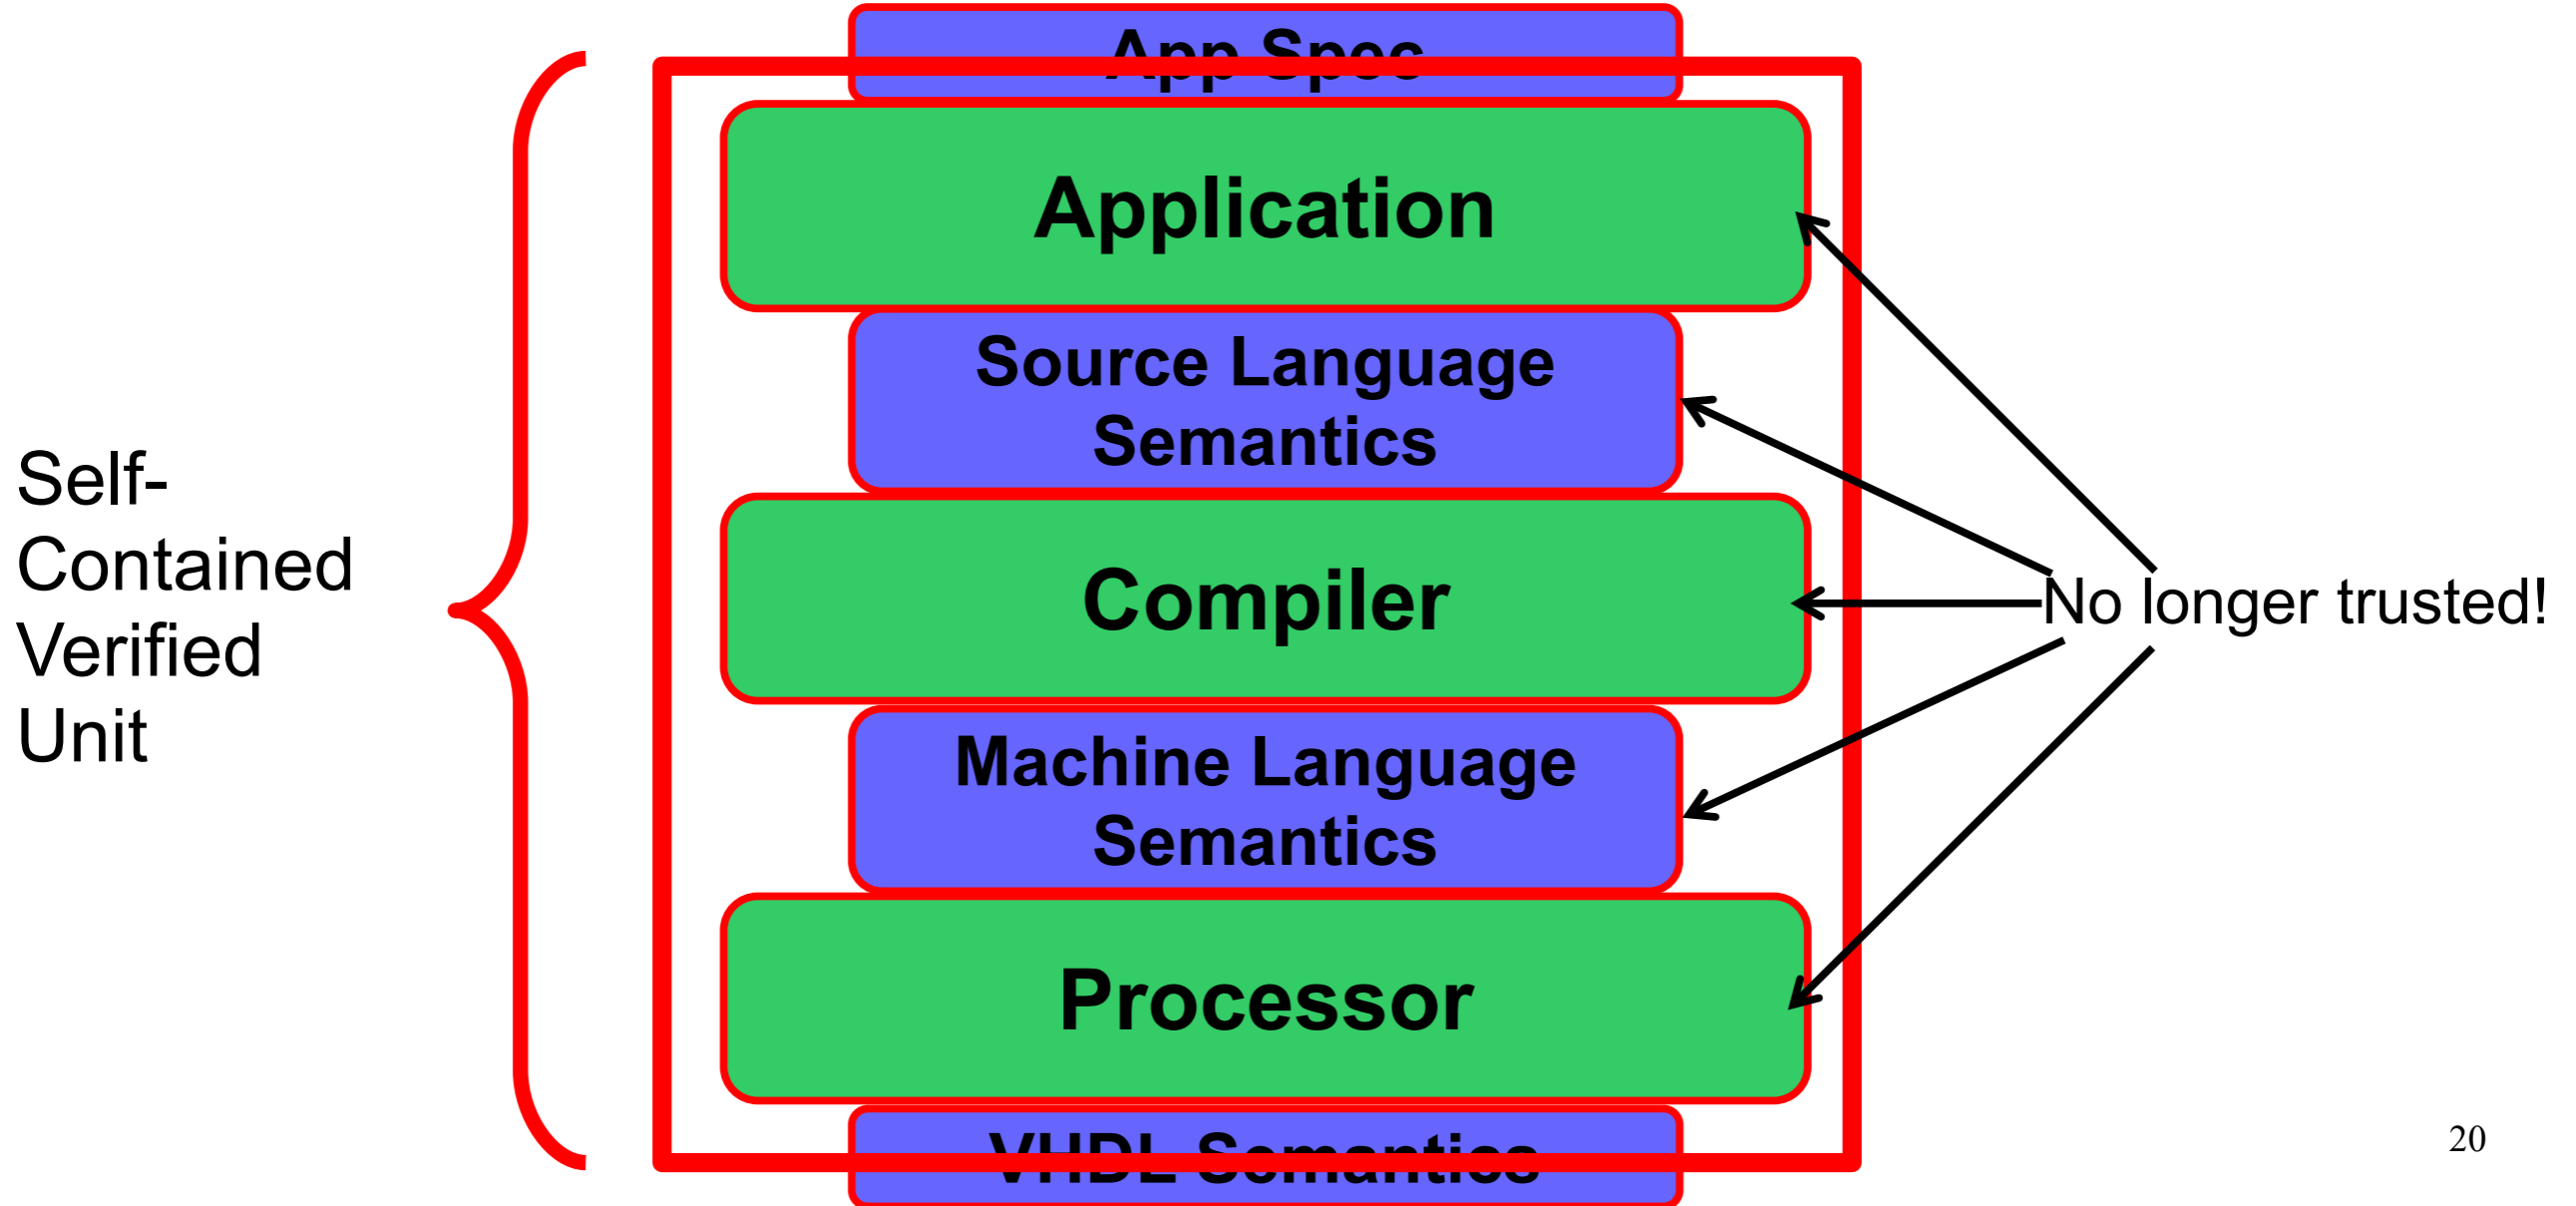
+

Optimizations

=

Implementation

Q: Aren't Those Specs *Still* Hard to Get Right?



Old vs. New

Old

System-integration tests
and unit tests,
since combined state space grows
exponentially as we compose pieces

Careful code review of all components,
since a corner-case bug in any of them
can wreck the whole system

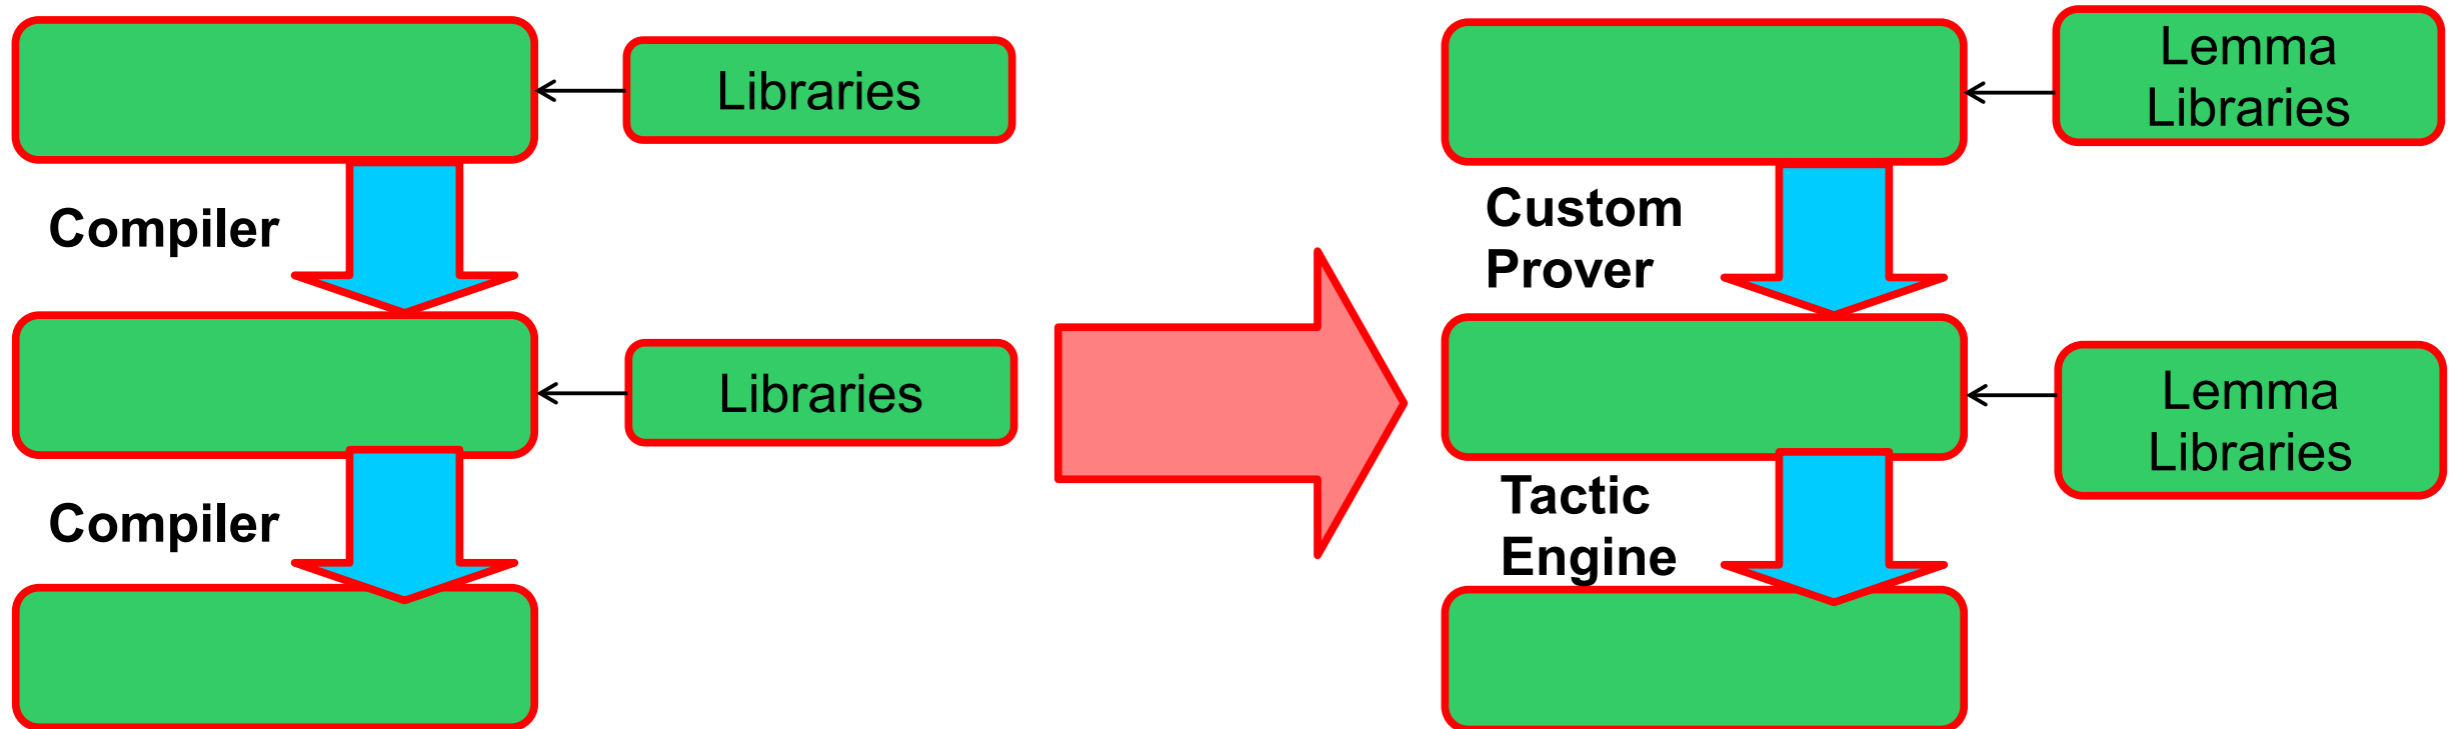
New

System-integration theorems imply
proper functioning of all
components.

Careful code review only of
externally facing specs

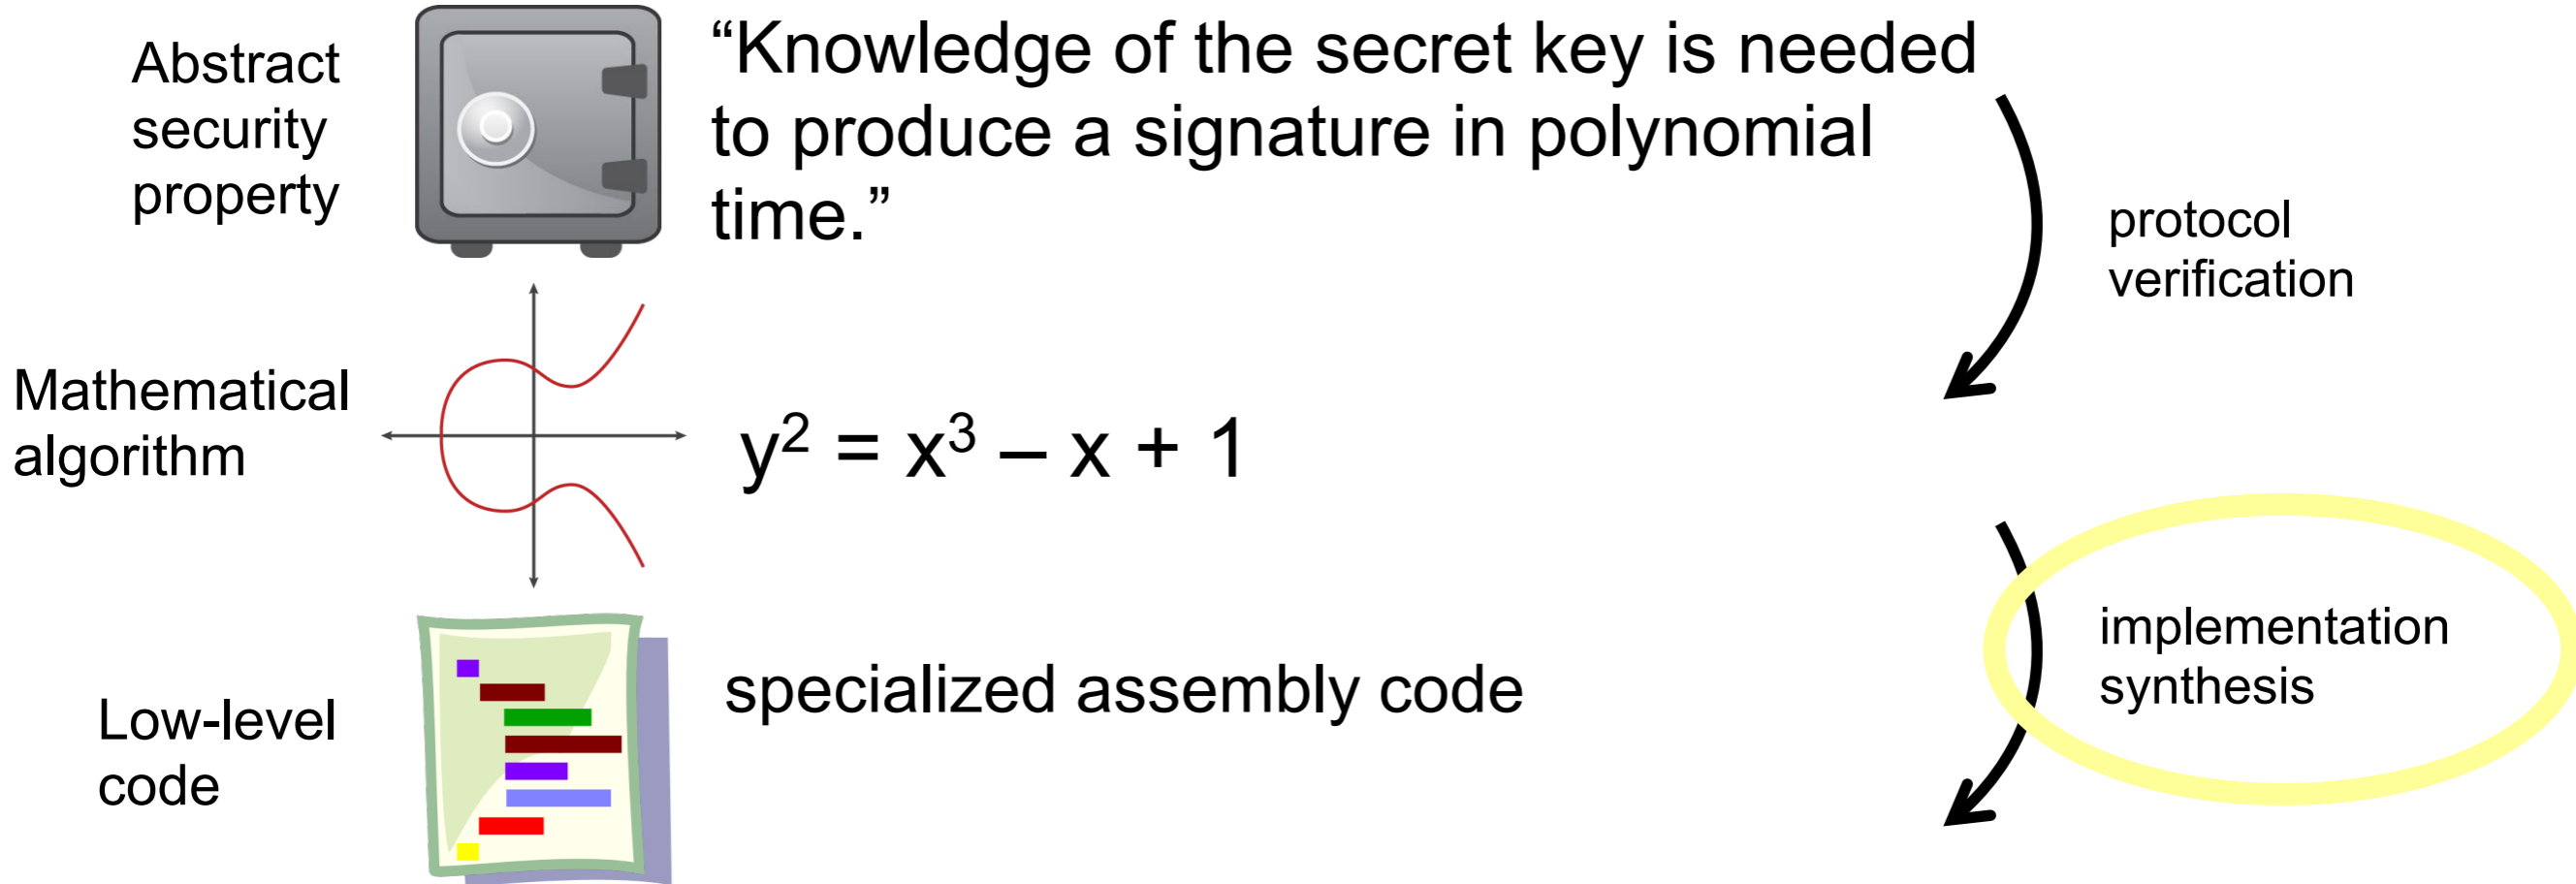
Q: Aren't the Proofs Huge and Unwieldy?

Well, aren't machine-code programs huge, too?




Motivation: correct-by-construction crypto

Correct-by-Construction Cryptography




Correct-by-Construction Cryptography

Mathematical algorithm




point = (x, y)

Optimized point format




point = (x, y, z, t)

High-level modular arithmetic



$x = x_0, x_1, \dots, x_n$
(mathematical integers)

Low-level code



specialized low-level code
(assumes fixed set of integer sizes)

classic verification
of functional programs

classic verification
of functional programs

compile-time code
specialization

compiler verification

Generated Code

Squaring a number (64-bit)

```
λ '(x7, x8, x6, x4, x2)%core,
uint64_t x9 = x2 * 0x2;
uint64_t x10 = x4 * 0x2;
uint64_t x11 = x6 * 0x2 * 0x13;
uint64_t x12 = x7 * 0x13;
uint64_t x13 = x12 * 0x2;
uint128_t x14 = (uint128_t) x2 * x2 + (uint128_t) x13 * x4 + (uint128_t) x11 * x8;
uint128_t x15 = (uint128_t) x9 * x4 + (uint128_t) x13 * x6 + (uint128_t) x8 * (x8 * 0x13);
uint128_t x16 = (uint128_t) x9 * x6 + (uint128_t) x4 * x4 + (uint128_t) x13 * x8;
uint128_t x17 = (uint128_t) x9 * x8 + (uint128_t) x10 * x6 + (uint128_t) x7 * x12;
uint128_t x18 = (uint128_t) x9 * x7 + (uint128_t) x10 * x8 + (uint128_t) x6 * x6;
uint64_t x19 = (uint64_t) (x14 >> 0x33);
uint64_t x20 = (uint64_t) x14 & 0x7fffffffffffff;
uint128_t x21 = x19 + x15;
uint64_t x22 = (uint64_t) (x21 >> 0x33);
uint64_t x23 = (uint64_t) x21 & 0x7fffffffffffff;
uint128_t x24 = x22 + x16;
uint64_t x25 = (uint64_t) (x24 >> 0x33);
uint64_t x26 = (uint64_t) x24 & 0x7fffffffffffff;
uint128_t x27 = x25 + x17;
uint64_t x28 = (uint64_t) (x27 >> 0x33);
uint64_t x29 = (uint64_t) x27 & 0x7fffffffffffff;
uint128_t x30 = x28 + x18;
uint64_t x31 = (uint64_t) (x30 >> 0x33);
uint64_t x32 = (uint64_t) x30 & 0x7fffffffffffff;
uint64_t x33 = x20 + 0x13 * x31;
uint64_t x34 = x33 >> 0x33;
uint64_t x35 = x33 & 0x7fffffffffffff;
uint64_t x36 = x34 + x23;
uint64_t x37 = x36 >> 0x33;
uint64_t x38 = x36 & 0x7fffffffffffff;
return (Return x32, Return x29, x37 + x26, Return x38, Return x35))
```

Squaring a number (32-bit)

```
λ '(x17, x18, x16, x14, x12, x10, x8, x6, x4, x2)%core,
uint64_t x19 = (uint64_t) x2 * x2;
uint64_t x20 = (uint64_t) (0x2 * x2) * x4;
uint64_t x21 = 0x2 * ((uint64_t) x4 * x4 + (uint64_t) x2 * x6);
uint64_t x22 = 0x2 * ((uint64_t) x4 * x6 + (uint64_t) x2 * x8);
uint64_t x23 = (uint64_t) x6 * x6 + (uint64_t) (0x4 * x4) * x8 + (uint64_t) (0x2 * x2) * x10;
uint64_t x24 = 0x2 * ((uint64_t) x6 * x8 + (uint64_t) x4 * x10 + (uint64_t) x2 * x12);
uint64_t x25 = 0x2 * ((uint64_t) x8 * x8 + (uint64_t) x6 * x10 + (uint64_t) x2 * x14 + (uint64_t) (0x2 * x4) * x12);
uint64_t x26 = 0x2 * ((uint64_t) x8 * x10 + (uint64_t) x6 * x12 + (uint64_t) x4 * x14 + (uint64_t) x2 * x16);
uint64_t x27 = (uint64_t) x10 * x10 + 0x2 * ((uint64_t) x6 * x14 + (uint64_t) x2 * x18 + 0x2 * ((uint64_t) x4 * x16 + (uint64_t) x8 * x12));
uint64_t x28 = 0x2 * ((uint64_t) x10 * x12 + (uint64_t) x8 * x14 + (uint64_t) x6 * x16 + (uint64_t) x4 * x18 + (uint64_t) x2 * x17);
uint64_t x29 = 0x2 * ((uint64_t) x12 * x12 + (uint64_t) x10 * x14 + (uint64_t) x8 * x16 + 0x2 * ((uint64_t) x6 * x18 + (uint64_t) x4 * x17));
uint64_t x30 = 0x2 * ((uint64_t) x12 * x14 + (uint64_t) x10 * x16 + (uint64_t) x8 * x18 + (uint64_t) x6 * x17);
uint64_t x31 = (uint64_t) x14 * x14 + 0x2 * ((uint64_t) x10 * x18 + 0x2 * ((uint64_t) x12 * x16 + (uint64_t) x8 * x17));
uint64_t x32 = 0x2 * ((uint64_t) x14 * x16 + (uint64_t) x12 * x18 + (uint64_t) x10 * x17);
uint64_t x33 = 0x2 * ((uint64_t) x16 * x16 + (uint64_t) x14 * x18 + (uint64_t) x12 * x17);
uint64_t x34 = 0x2 * ((uint64_t) x16 * x18 + (uint64_t) x14 * x17);
uint64_t x35 = (uint64_t) x18 * x18 + (uint64_t) (0x4 * x16) * x17;
uint64_t x36 = (uint64_t) (0x2 * x18) * x17;
uint64_t x37 = (uint64_t) (0x2 * x17) * x17;
uint64_t x38 = x27 * x37 << 0x4;
uint64_t x39 = x38 + x37 << 0x1;
uint64_t x40 = x39 * x37;
uint64_t x41 = x26 + x36 << 0x4;
uint64_t x42 = x41 + x36 << 0x1;
uint64_t x43 = x42 + x36;
uint64_t x44 = x25 + x35 << 0x4;
uint64_t x45 = x44 + x35 << 0x1;
uint64_t x46 = x45 + x35;
uint64_t x47 = x24 + x34 << 0x4;
uint64_t x48 = x47 + x34 << 0x1;
uint64_t x49 = x48 + x34;
uint64_t x50 = x23 + x33 << 0x4;
uint64_t x51 = x50 + x33 << 0x1;
uint64_t x52 = x51 + x33;
uint64_t x53 = x22 + x32 << 0x4;
uint64_t x54 = x53 + x32 << 0x1;
uint64_t x55 = x54 + x32;
uint64_t x56 = x21 + x31 << 0x4;
uint64_t x57 = x56 + x31 << 0x1;
uint64_t x58 = x57 + x31;
uint64_t x59 = x20 + x30 << 0x4;
uint64_t x60 = x59 + x30 << 0x1;
uint64_t x61 = x60 + x30;
uint64_t x62 = x19 + x29 << 0x4;
uint64_t x63 = x62 + x29 << 0x1;
uint64_t x64 = x63 + x29;
uint64_t x65 = x64 >> 0x1a;
uint32_t x66 = (uint32_t) x64 & 0x3fffffff;
uint64_t x67 = x65 + x61;
uint64_t x68 = x67 >> 0x19;
uint32_t x69 = (uint32_t) x67 & 0x1fffffff;
uint64_t x70 = x68 + x68;
uint64_t x71 = x70 >> 0x1a;
uint32_t x72 = (uint32_t) x70 & 0x3fffffff;
uint64_t x73 = x71 + x55;
uint64_t x74 = x73 >> 0x19;
uint32_t x75 = (uint32_t) x73 & 0x1fffffff;
uint64_t x76 = x74 + x52;
uint64_t x77 = x76 >> 0x1a;
uint32_t x78 = (uint32_t) x76 & 0x3fffffff;
uint64_t x79 = x77 + x49;
uint64_t x80 = x79 >> 0x19;
uint32_t x81 = (uint32_t) x79 & 0x1fffffff;
uint64_t x82 = x80 + x46;
uint32_t x83 = (uint32_t) (x82 >> 0x1a);
uint32_t x84 = (uint32_t) x82 & 0x3fffffff;
uint64_t x85 = x83 + x43;
uint32_t x86 = (uint32_t) (x85 >> 0x19);
uint32_t x87 = (uint32_t) x85 & 0x1fffffff;
uint64_t x88 = x86 + x40;
uint32_t x89 = (uint32_t) (x88 >> 0x1a);
uint32_t x90 = (uint32_t) x88 & 0x3fffffff;
uint64_t x91 = x89 + x28;
uint32_t x92 = (uint32_t) (x91 >> 0x19);
uint32_t x93 = (uint32_t) x91 & 0x1fffffff;
uint64_t x94 = x66 + (uint64_t) (0x13 * x92);
uint32_t x95 = (uint32_t) (x94 >> 0x1a);
uint32_t x96 = (uint32_t) x94 & 0x3fffffff;
uint32_t x97 = x95 + x69;
uint32_t x98 = x97 >> 0x19;
uint32_t x99 = x97 & 0x1fffffff;
return (Return x93, Return x90, Return x87, Return x84, Return x81, Return x78, Return x75, x98 + x72, Return x99, Return x96))
```

Surprising (?) Fact About Modular Arithmetic

Different prime moduli have dramatically different efficiency with best code on commodity processors.

$2^{255} - 19$ is a popular choice for relatively easy implementation.

General pattern: $2^k - c$, for $c \ll 2^k$. (Called *pseudo-Mersenne*.)

Example of a fast operation: *modular reduction*

$$\begin{aligned} t &= x + 2^k y \pmod{2^k - c} \quad \text{too big to fit below the modulus!} \\ &= x + (2^k - c + c)y \pmod{2^k - c} \\ &= x + \cancel{(2^k - c)y} + cy \pmod{2^k - c} \\ &= x + cy \pmod{2^k - c} \end{aligned}$$

Representing Numbers mod $2^{255} - 19$

t result of multiplying two numbers in the prime field, so **510 bits wide**

$$t = t_0 t_1 t_2 t_3 t_4 t_5 t_6 t_7$$
$$= (t_0 + 2^{64} t_1 + \dots) + 2^{256} (t_4 + 2^{64} t_5 + \dots)$$

each "digit" fits in 64-bit register

darn, that's 2^{256} , not 2^{255} , so we can't use that reduction trick!

However.... $51 \times 10 = 510$.

$$t = (t_0 + 2^{51} t_1 + \dots) + 2^{255} (t_5 + 2^{51} t_6 + \dots)$$

champion rep. on **64-bit processors**
(note: not using full bitwidth!)

Also.... $25.5 \times 2 = 51$.

$$t = s_0 + 2^{25.5} s_1 + 2^{2 \times 25.5} s_2 + 2^{3 \times 25.5} s_3 + \dots$$

$$t = s_0 + 2^{26} s_1 + 2^{51} s_2 + 2^{77} s_3 + \dots$$

champion rep. on **32-bit processors**
(note: nonuniform bitwidths!)

Demo

Invoking Fiat Cryptography

The Fiat Cryptography approach



The Basic Idea

Choice of base-system representation

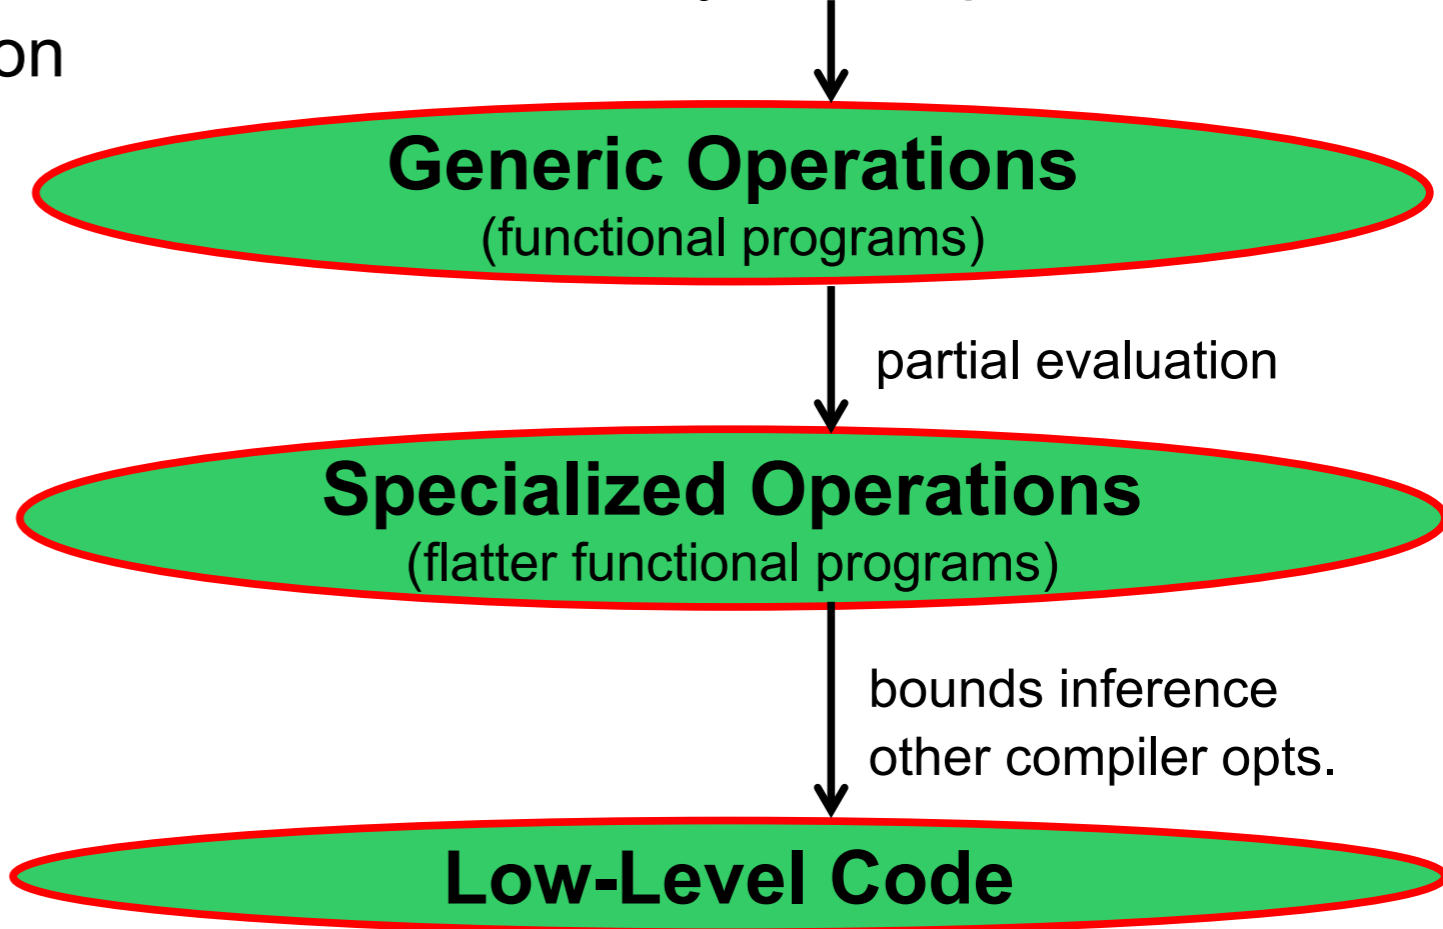


Our Library

proof

Fast C code

Choice of base-system representation



Generic Operations

(functional programs)

partial evaluation

Specialized Operations

(flatter functional programs)

bounds inference
other compiler opts.

Low-Level Code

Example: Multiplication (for modulus $2^{127} - 1$)

$$s = s_0 + 2^{43} s_1 + 2^{85} s_2$$

$$t = t_0 + 2^{43} t_1 + 2^{85} t_2$$

$$s \square t = u = u_0 u_1 u_2 u_3 u_4$$

$$s \square t = 1 \square s_0 t_0 + 2^{43} \square s_0 t_1 + 2^{85} \square s_0 t_2 + 2^{43} \square s_1 t_0 + 2^{86} \square s_1 t_1 + 2^{128} \square s_1 t_2 + 2^{85} \square s_2 t_0 + 2^{170} \square s_2 t_1 + 2^{170} \square s_2 t_2$$

$$u_0 = s_0 t_0$$

$$u_1 = s_0 t_1 + s_1 t_0$$

$$u_2 = s_0 t_2 + 2s_1 t_1 + s_2 t_0$$

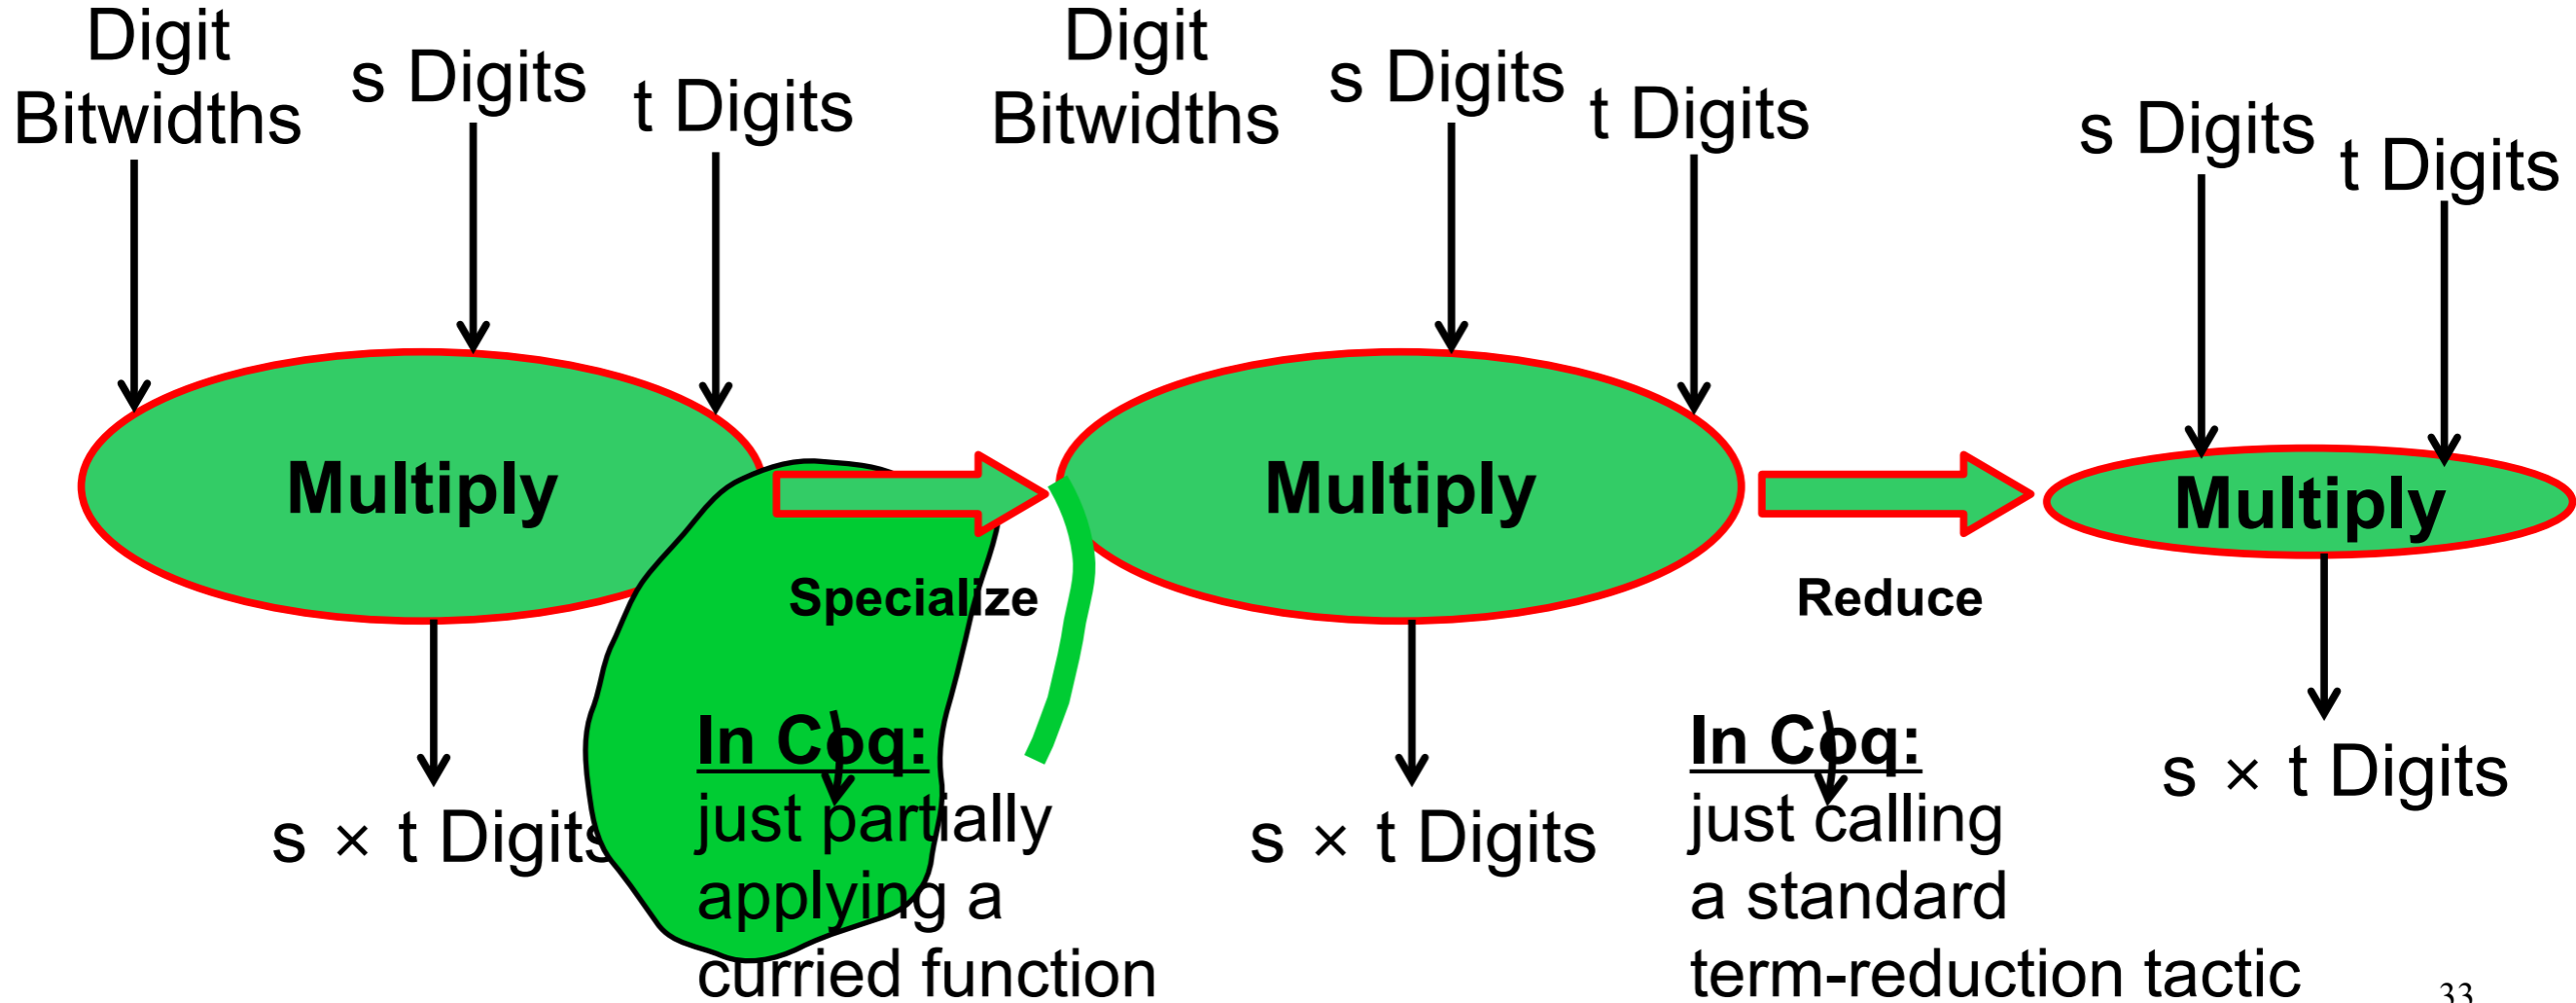
$$u_3 = 2s_1 t_2 + 2s_2 t_1$$

$$u_4 = s_2 t_2$$

$$u = u_0 + 2^{43} u_1 + 2^{85} u_2 + 2^{127} (u_3 + 2^{43} u_4)$$

$$= (u_0 + u_3) + 2^{43} (u_1 + u_4) + 2^{85} u_2$$

Time for Some Partial Evaluation



An Example

Definition $w (i:\text{nat}) : \mathbb{Z} := 2^{\lceil (25+1/2) * i \rceil}$.

Example $\text{base_25_5_mul } (f \ g:\text{tuple } \mathbb{Z} \ 10) :$

$\{ fg : \text{tuple } \mathbb{Z} \ 10 \mid$
 $(\text{eval } w \ fg) \bmod (2^{255}-19)$
 $= (\text{eval } w \ f * \text{eval } w \ g) \bmod (2^{255}-19) \}$.

$(f_0 * g_9 + f_1 * g_8 + f_2 * g_7 + f_3 * g_6 + f_4 * g_5 + f_5 * g_4 + f_6 * g_3 + f_7 * g_2 + f_8 * g_1 + f_9 * g_0,$
 $f_0 * g_8 + 2 * f_1 * g_7 + f_2 * g_6 + 2 * f_3 * g_5 + f_4 * g_4 + 2 * f_5 * g_3 + f_6 * g_2 + 2 * f_7 * g_1 + f_8 * g_0 + 38 * f_9 * g_9,$
 $f_0 * g_7 + f_1 * g_6 + f_2 * g_5 + f_3 * g_4 + f_4 * g_3 + f_5 * g_2 + f_6 * g_1 + f_7 * g_0 + 19 * f_8 * g_9 + 19 * f_9 * g_8,$
 $f_0 * g_6 + 2 * f_1 * g_5 + f_2 * g_4 + 2 * f_3 * g_3 + f_4 * g_2 + 2 * f_5 * g_1 + f_6 * g_0 + 38 * f_7 * g_9 + 19 * f_8 * g_8 + 38 * f_9 * g_7,$
 $f_0 * g_5 + f_1 * g_4 + f_2 * g_3 + f_3 * g_2 + f_4 * g_1 + f_5 * g_0 + 19 * f_6 * g_9 + 19 * f_7 * g_8 + 19 * f_8 * g_7 + 19 * f_9 * g_6,$
 $f_0 * g_4 + 2 * f_1 * g_3 + f_2 * g_2 + 2 * f_3 * g_1 + f_4 * g_0 + 38 * f_5 * g_9 + 19 * f_6 * g_8 + 38 * f_7 * g_7 + 19 * f_8 * g_6 + 38 * f_9 * g_5,$
 $f_0 * g_3 + f_1 * g_2 + f_2 * g_1 + f_3 * g_0 + 19 * f_4 * g_9 + 19 * f_5 * g_8 + 19 * f_6 * g_7 + 19 * f_7 * g_6 + 19 * f_8 * g_5 + 19 * f_9 * g_4,$
 $f_0 * g_2 + 2 * f_1 * g_1 + f_2 * g_0 + 38 * f_3 * g_9 + 19 * f_4 * g_8 + 38 * f_5 * g_7 + 19 * f_6 * g_6 + 38 * f_7 * g_5 + 19 * f_8 * g_4 + 38 * f_9 * g_3,$
 $f_0 * g_1 + f_1 * g_0 + 19 * f_2 * g_9 + 19 * f_3 * g_8 + 19 * f_4 * g_7 + 19 * f_5 * g_6 + 19 * f_6 * g_5 + 19 * f_7 * g_4 + 19 * f_8 * g_3 + 19 * f_9 * g_2,$
 $f_0 * g_0 + 38 * f_1 * g_9 + 19 * f_2 * g_8 + 38 * f_3 * g_7 + 19 * f_4 * g_6 + 38 * f_5 * g_5 + 19 * f_6 * g_4 + 38 * f_7 * g_3 + 19 * f_8 * g_2 + 38 * f_9 * g_1)$

Compiling to Low-Level Code

$$1 \times (1 \times 2^{52} + (1 \times x + 0)) + (1 \times (1 \times (-y) + 0) + 0)$$



reify to syntax tree



constant-fold

$$(2^{52} + x) - y$$



flatten

let c = $2^{52} + x$ in

let d = c - y in

d

Assume: $0 \leq x, y \leq 2^{51} + 2^{48}$

Deduce: $2^{52} \leq c \leq 2^{52} + 2^{51} + 2^{48}$

Deduce: $2^{51} - 2^{48} \leq d \leq 2^{52} + 2^{51} + 2^{48}$



infer bounds

```
uint64_t c = 252 +
```

```
x;
```

```
uint64_t d = c - y;
```

Implementation and Experiments

- ~38 kloc in full library (including significant parts that belong in stdlib)
- Very little code needed to instantiate to new prime moduli.
- In fact, we wrote a Python script (under 3000 lines) to generate parameters automatically from prime numbers, written suggestively, e.g. $2^{256} - 2^{224} + 2^{192} + 2^{96} - 1$.
- This script is outside the TCB, since any successful compilation is guaranteed to implement correct arithmetic.

Q: Where do we get a lot of reasonable moduli?

A: Scrape all prime numbers appearing in a popular mailing list.


We used the elliptic curves list at moderncrypto.org.

We found about 80 primes.


Only a few turned out to be terrible ideas posted by newbies.

Many-Primes Experiment

64-Bit Field Arithmetic Benchmarks



32-Bit Field Arithmetic Benchmarks



P256 Mixed Addition

Implementation	CPU cycles	μs at 2.6GHz
OpenSSL AMD64+ADX asm	544	.21
OpenSSL AMD64 asm	644	.25
<i>this work, icc</i>	1112	.43
<i>this work, gcc</i>	1808	.70
OpenSSL C	1968	.76


Towards correct-by-construction cryptographic appliances

The Verified IoT Lightbulb!

Ethernet card

FPGA running our formally verified processor and software stack

Power strip & lightbulb




The Verified IoT Lightbulb!

Ethernet card

FPGA running our formally verified processor and software stack


Power strip & lightbulb

Scope of formal proof:
essentially all code
(HW & SW) loaded on
this FPGA




Specification?

Consider all **traces** the system could generate:
00100, 11000, 00100, ...



Input Output



Recording pin values each cycle

Output pins: we as spec-writers may mandate what they are allowed to be!

Input pins: the environment may choose any values each cycle.

“Output pin controlling lightbulb is only on if the last valid Ethernet packet said so.”

Key Layers of End-to-End Proof



Controller Spec (Trace Predicate)

Controller SW



Programming Language Semantics

Verified Compiler



ISA Family Semantics



Verified Hardware

RTL Semantics

Disappearing Specs




Controller Spec (Trace Predicate)

Must get this spec right.



Everything this box hides is no longer trusted!

System as a Proved Black Box



Must get this one right, too.

RTL Semantics


Expanding Scope

Abstract security property




“Knowledge of the secret key is needed to produce a signature in polynomial time.”

Mathematical algorithm




$$y^2 = x^3 - x + 1$$

High-level modular arithmetic



$$X = X_0, X_1, \dots, X_n$$

Low-level code



specialized assembly code



Protocol verification, perhaps following past work by Appel & others, using our new higher-level notation for protocol programming

Synthesizing C code for more of a crypto library (beyond straightline code) with Rupicola, a proof-generating compiler

Genetic search for fast assembly code (collaboration with Prof. Yuval Yarom et al.), plus formally verified program-equivalence checker

Connect to verified HW & systems software

<https://github.com/mit-plv/flat>

<https://github.com/mit-plv/bedrock2>