

# Cryptographic Primitives Implemented Efficiently, Correctly and Securely

Nik Swamy, OPLSS 2021

Thanks to Jonathan Protzenko and Chris Hawblitzel  
for these slides. Errors are mine

Crypto code is hard to get right

But is critical for system security

# Many bugs in Curve25519 implementations

(C and assembly)

agl / curve25519-donna

<> Code

Issues 2

Pull requests 7

Projects 0

Wiki

Insights

## Correct bounds in 32-bit code.

The 32-bit code was illustrative of the tricks used in the original curve25519 paper rather than rigorous. However, it has proven quite popular.

This change fixes an issue that Robert Ransom found where outputs between  $2^{255}-19$  and  $2^{255}-1$  weren't correctly reduced in fcontract. This appears to leak a small fraction of a bit of security of private keys.

Additionally, the code has been cleaned up to reflect the real-world needs. The ref10 code also exists for 32-bit, generic C but is somewhat slower and objections around the lack of qasm availability have been raised.

master 1.3

Curve25519-donna



agl committed on Jun 9, 2014

1 parent

Ed25519 amd64 bug

gistfile1.md

Raw

Watch

NaCl (asm)

While visiting 30c3, I attended the [You-broke-the-Internet workshop on NaCl](#).

One thing mentioned in the talk was that auditing crypto code is a lot of work, and that this is one of the reasons why Ed25519 isn't included in NaCl yet (they promised a version including it for 2014). The speakers mentioned a bug in the amd64 assembly implementation of Ed25519 as an example of a bug that can only be found by auditing, not by randomized tests. This bug is caused by a carry being added in the wrong place, but since that carry is usually zero, the bug is hard to find (occurs with probability  $2^{-60}$  or so).

The [TweetNaCl](#) paper briefly mentions this bug as well:

Partial audits have revealed a bug in this software ( `r1 += 0 + carry` should be `r2 += 0 + carry` in `amd64-64-24k` ) that would not be caught by random tests; this illustrates the importance of audits.

Searching for this string in the SUPERCOP source code turns up four matches:

```
crypto_scalarmult\curve25519\amd64-64\fe25519_mul.s
crypto_scalarmult\curve25519\amd64-64\fe25519_square.s
crypto_sign\ed25519\amd64-64-24k\fe25519_mul.s
crypto_sign\ed25519\amd64-64-24k\fe25519_square.s
```

So it appears like the `amd64-64` implementation of both Curve25519 and Ed25519 is affected.

It seems difficult to exploit this when used for key generation or signing since the attacker cannot influence the data. Key-exchange and signature verification might be a problem.

TweetNaCl

```
sv pack25519(u8 *o)
{
    int i,j,b;
    gf m,t;
    FOR(i,16) t[i]=n
    car25519(t);
    car25519(t);
    car25519(t);
    FOR(j,2) {
        m[0]=t[0]-0xff
        for(i=1;i<15;i
        m[i]=t[i]-0x
        m[i-1]&=0xff
    }
    m[15]=t[15]-0x
    b=(m[15]>>16)&
    m[15]&=0xffff;
    sel25519(t,m,1-b);
}
FOR(i,16) {
    o[2*i]=t[i]&0xff;
    o[2*i+1]=t[i]>>8;
}
}
```

This bug is triggered when the last limb `n[15]` of the input argument `n` of this function is greater or equal than `0xffff`. In these cases the result of the scalar multiplication is not reduced as expected resulting in a wrong packed value. This code can be fixed simply by replacing `m[15]&=0xffff;` by `m[14]&=0xffff;`.

# 3 Bugs in OpenSSL implementation of Poly1305

OpenSSL Security Advisory [10 Nov 2016]

[openssl-dev] [openssl.org #4439] poly1305-x86.pl produces incorrect output

“These produce wrong results. The first example does so only on 32 bit, the other three also on 64 bit.”

“I believe this affects both the SSE2 and AVX2 code. It does seem to be dependent on this input pattern.”

“I'm probably going to write something to generate random inputs and stress all your other poly1305 code paths against a reference implementation.”

the other three also on 64 bit.

recommend doing the same in your own test harness, to make sure there aren't others of these bugs lurking around.

# Implementation bug in AES-GCM

## **The fragility of AES-GCM authentication algorithm**

Shay Gueron<sup>1,2</sup>, Vlad Krasnov<sup>2</sup>

<sup>1</sup> Department of Mathematics, University of Haifa, Israel

<sup>2</sup> Intel Corporation, Israel Development Center, Haifa, Israel

March 15, 2013

**Abstract.** A new implementation of the GHASH function has been recently committed to a Git version of OpenSSL, to speed up AES-GCM. We identified a bug in that implementation, and made sure it was quickly fixed before trickling into an official OpenSSL trunk. Here, we use this (already fixed) bug

# AES-GCM

Evaluate polynomials in this field to get an **authentication code!** (see also: Poly1305)

## GHASH (AES-GCM):

- $p = 2^{128} \ (q = 2, n = 128)$
- $P = x^{128} + x^7 + x^2 + x + 1$

**“the math”**



# Writing the actual code

A long way from the math

“the reality”

```
444 .Lgcm_dec_body:
445
446 $code.=<<__ ;
447     vzeroupper
448
449     vmovdqu    ($ivp), $T1          # input counter value
450     add        \($-128,%rsp
451     mov        12($ivp), $counter
452     lea        .Lbswap_mask(%rip), $const
453     lea        -0x80($key), $in0    # borrow $in0
454     mov        \($0xf80,$end0      # borrow $end0
455     vmovdqu    ($Xip), $Xi         # load Xi
456     and        \($-128,%rsp        # ensure stack alignment
457     vmovdqu    ($const), $Ii       # borrow $Ii for .Lbswap_mask
458     lea        0x80($key), $key     # size optimization
459     lea        0x20+0x20($Xip), $Xip # size optimization
460     mov        0xf0-0x80($key), $rounds
461     vpshufb    $Ii, $Xi, $Xi
462
463     and        $end0, $in0
464     and        %rsp, $end0
465     sub        $in0, $end0
466     jc         .Ldec_no_key_aliasing
467     cmp        \($768,$end0
468     jnc        .Ldec_no_key_aliasing
469     sub        $end0, %rsp          # avoid aliasing with key
470 .Ldec_no_key_aliasing:
471
472     vmovdqu    0x50($inp), $Z3      # I[5]
473     lea        ($inp), $in0
474     vmovdqu    0x40($inp), $Z0
475     lea        -0xc0($inp,$len), $end0
476     vmovdqu    0x30($inp), $Z1
477     shr        \($4,$len
478     xor        $ret, $ret
479     vmovdqu    0x20($inp), $Z2
480     vpshufb    $Ii, $Z3, $Z3        # passed to _aesni_ctr32_ghash_6x
481     vmovdqu    0x10($inp), $T2
482     vpshufb    $Ii, $Z0, $Z0
483     vmovdqu    ($inp), $Hkey
484     vpshufb    $Ii, $Z1, $Z1
485     vmovdqu    $Z0, 0x30(%rsp)
486     vpshufb    $Ii, $Z2, $Z2
487     vmovdqu    $Z1, 0x40(%rsp)
488     vpshufb    $Ii, $T2, $T2
489     vmovdqu    $Z2, 0x50(%rsp)
490     vpshufb    $Ii, $Hkey, $Hkey
491     vmovdqu    $T2, 0x60(%rsp)
492     vmovdqu    $Hkey, 0x70(%rsp)
493
494     call       _aesni_ctr32_ghash_6x
495
496     vmovups    $inout0, -0x60($out) # save output
497     vmovups    $inout1, -0x50($out)
```



$$GF(2^{128}) = GF(2)[X]/(x^{128} + x^7 + x^2 + x + 1)$$

refines

**Algorithm 1** Multiplication in  $GF(2^{128})$ .  
 $Z \in GF(2^{128})$ .

---

```

 $Z \leftarrow 0, V \leftarrow X$ 
for  $i = 0$  to 127 do
  if  $Y_i = 1$  then
     $Z \leftarrow Z \oplus V$ 
  end if
  if  $V_{127} = 0$  then
     $V \leftarrow \text{rightshift}(V)$ 
  else
     $V \leftarrow \text{rightshift}(V) \oplus R$ 
  end if
end for
return  $Z$ 

```

---

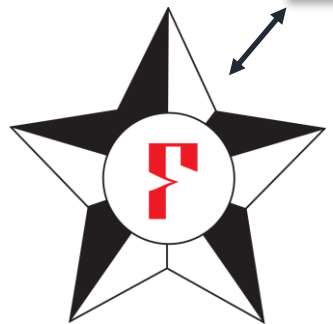
refines

```

vmovdqu    ($ivp), $T1
add        \$. -128, %rsp
mov        12($ivp), $counter
lea        .Lbswap_mask(%rip), $cons
lea        -0x80($key), $in0
mov        \$.0xf80, $end0
vmovdqu    ($Xip), $Xi
and        \$. -128, %rsp
vmovdqu    ($const), $Ii
lea        0x80($key), $key
lea        0x20+0x20($Xip), $Xip
mov        0xf0-0x80($key), $rounds
vpslufb    $Ii, $Xi, $Xi

```

Z3



Specification  
(*"the mathematical truth"*)

Pseudo-code  
(*"implementation blueprint"*)

proof

proof

proof

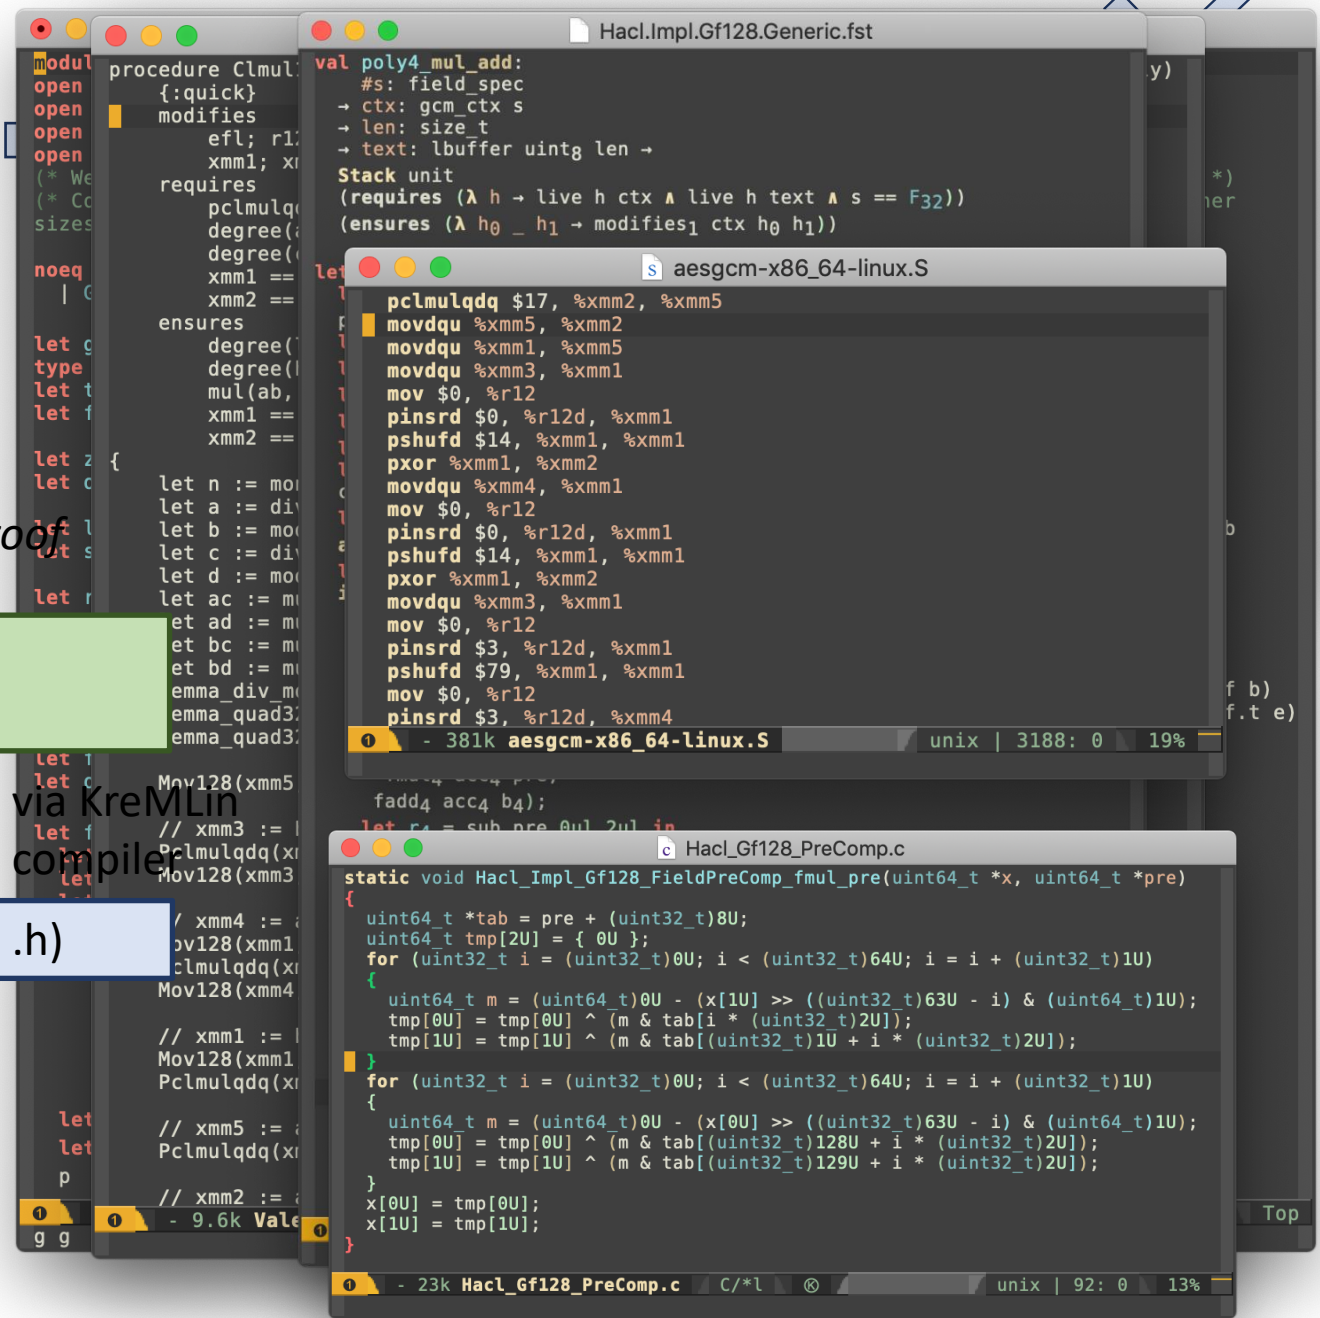
**Vale**

Vale/F\*  
(*"assembly-like"*)

Low\*  
(*"C-like"*)

via Vale  
printer  
Assembly (.asm)

via KreMLin  
compiler  
C code (.c, .h)



# What do we verify?

## Safety

Memory- and type-safety. Mitigates buffer overruns, dangling pointers, code injections. No undefined behavior.

## Functional correctness

Our fast implementations behave precisely as our simpler specifications.

## Secrecy

Access to secrets, including crypto keys and private app data is restricted according to design.

Our specifications and implementations are written together, in one language (F\*)

Drift between spec and implementation cannot happen.

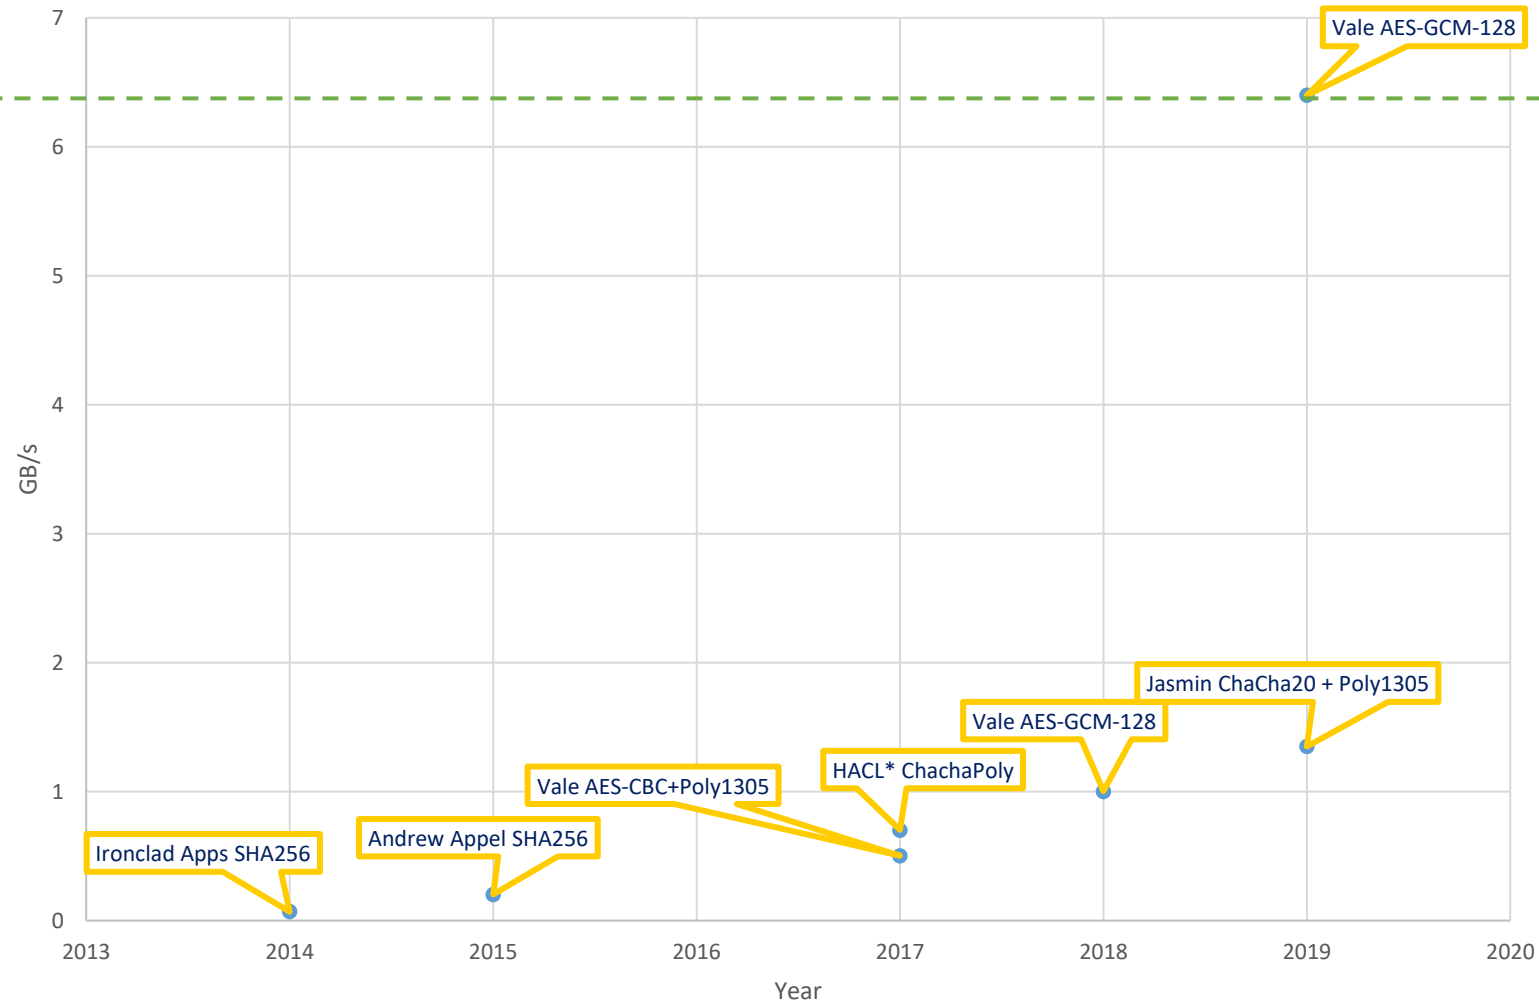
Each application can do custom proofs beyond functional correctness and safety:

- non malleability (parsers)
- crypto games (TLS)
- security reduction (Merkle Trees)
- etc. etc.

# Verified Assembly Language in Vale / F\*

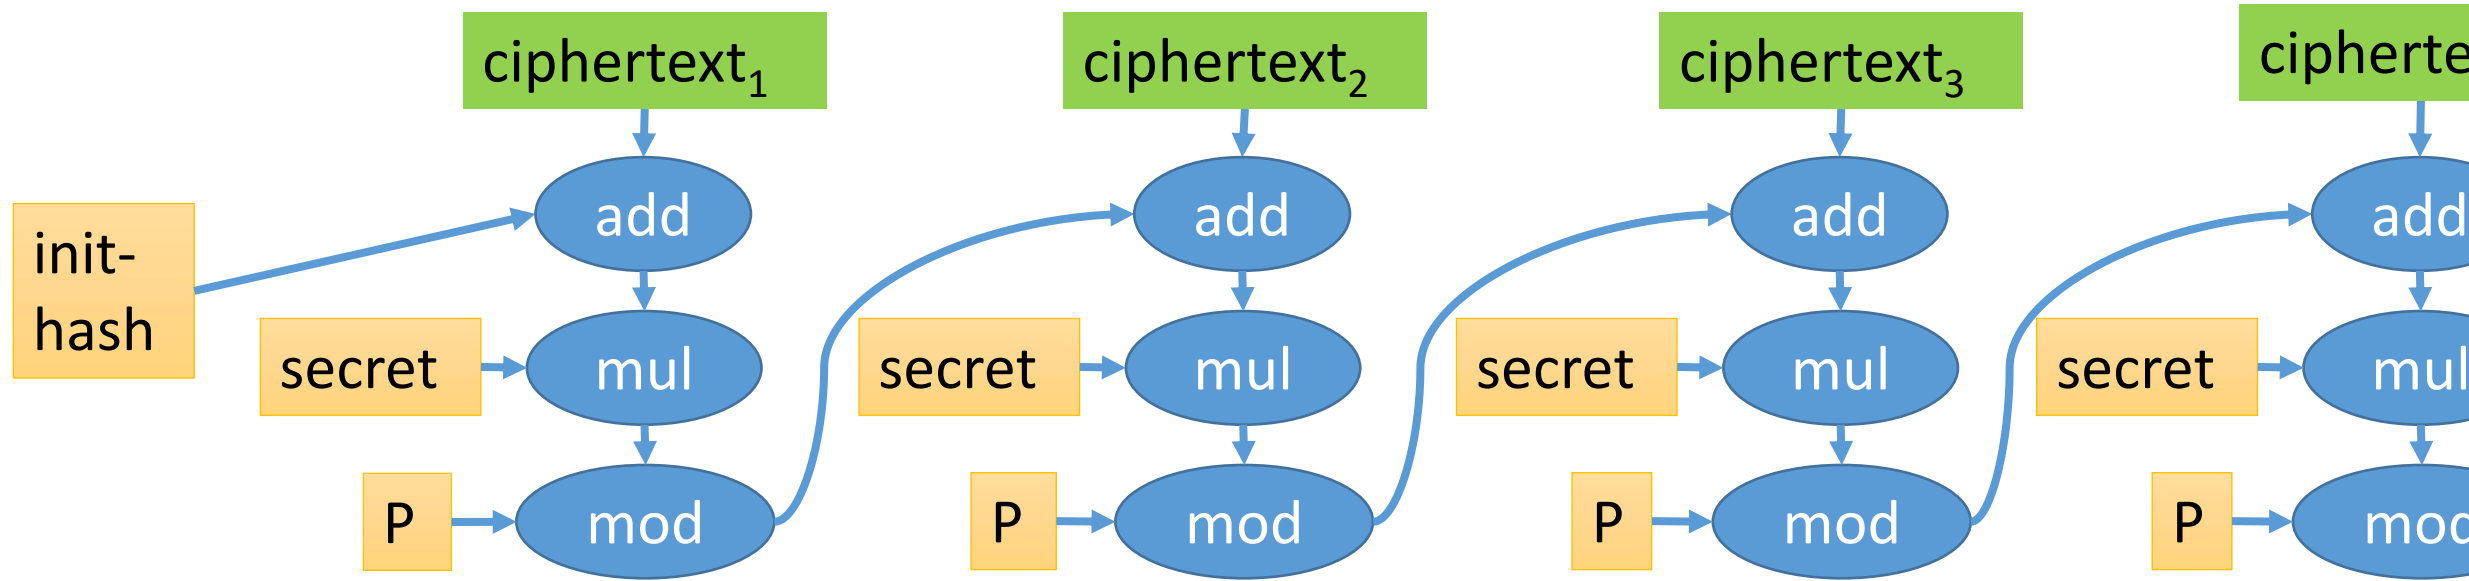
# We have a *fast* verified AES-GCM

Performance of various verified symmetric crypto / hash implementations



Fastest  
OpenSSL  
assembly  
code

# Optimizing AES-GCM



Important optimizations:

- delay mod operations
- parallelize add/mul operations
- math+bitwise tricks for mod
- careful instruction scheduling

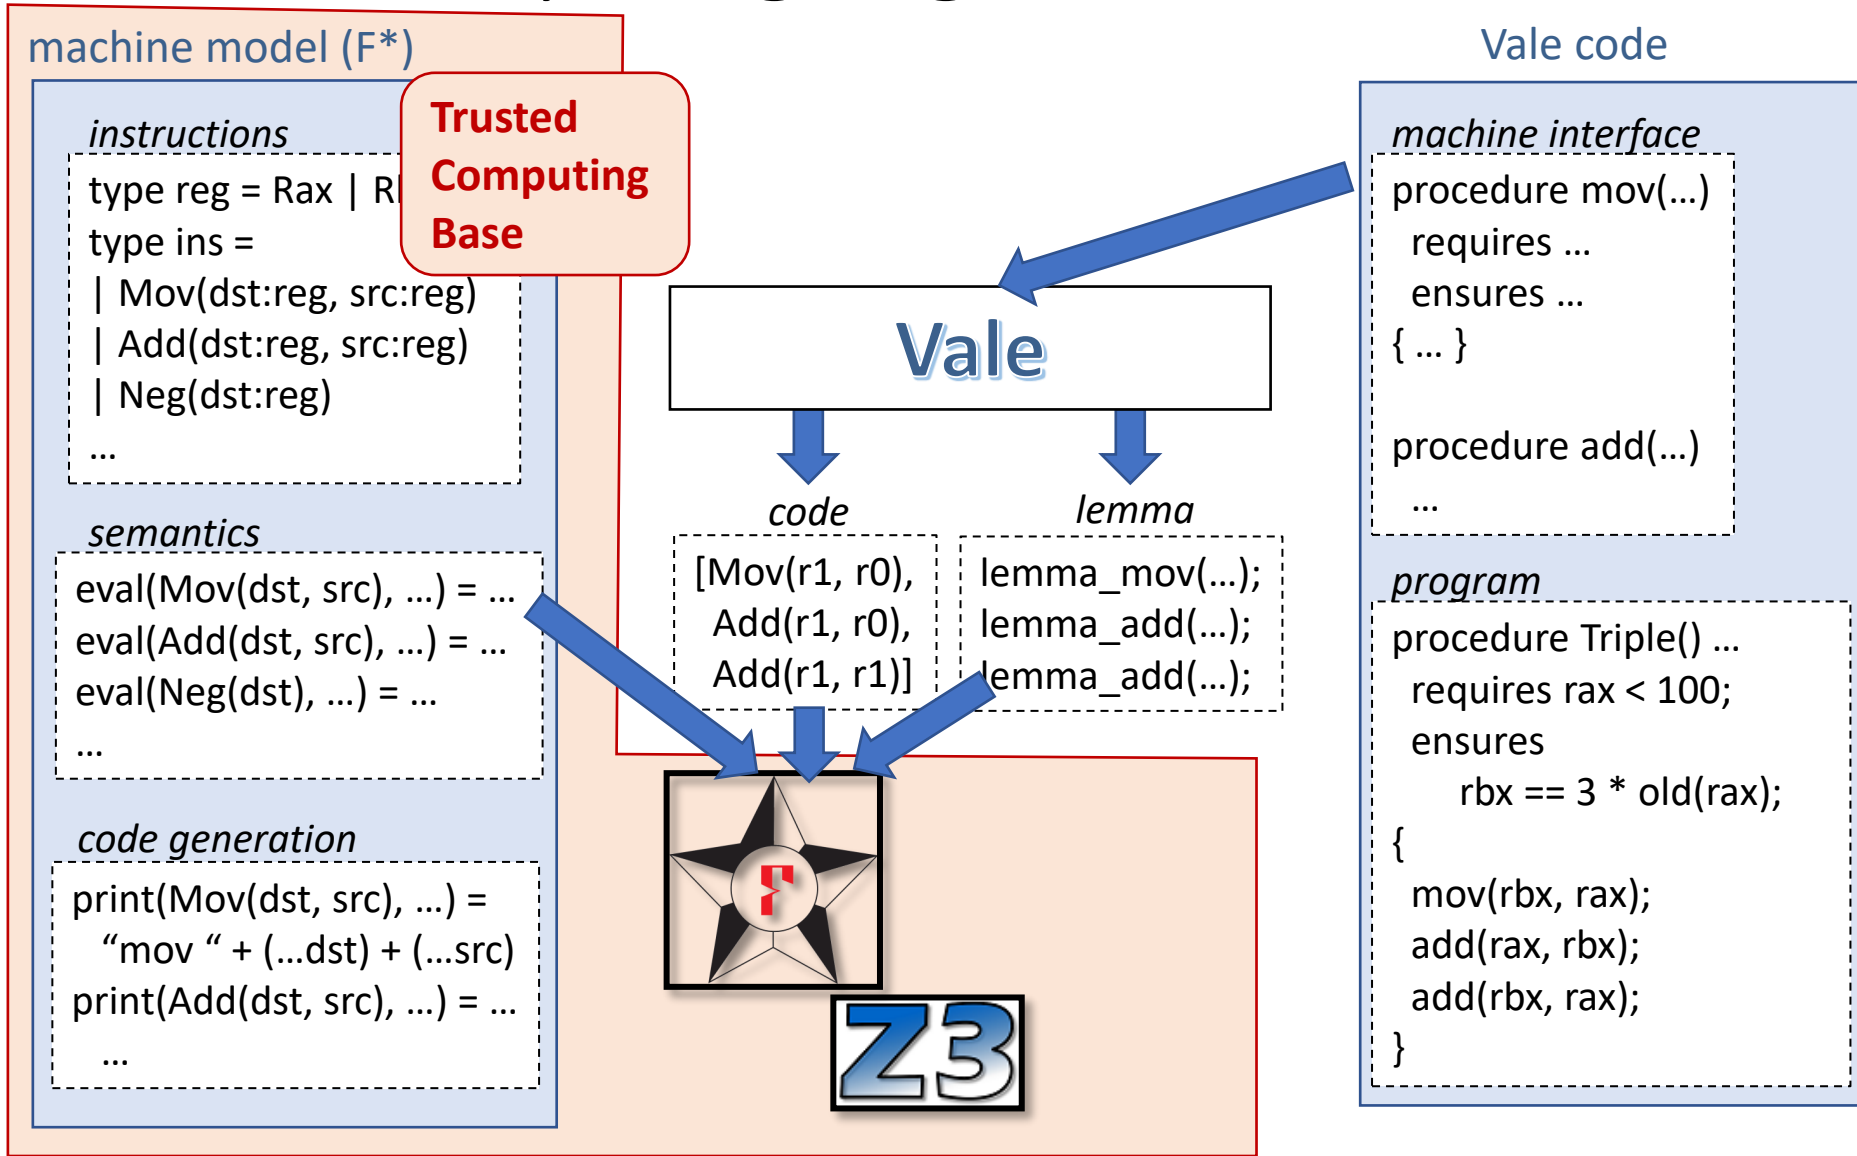
$$(((\text{init} + c_1) * s \% P + c_2) * s \% P + c_3) * s \% P$$

$$\rightarrow (((\text{init} + c_1) * s + c_2) * s + c_3) * s \% P$$

$$\rightarrow ((\text{init} + c_1) * s^3 + c_2 * s^2 + c_3 * s^1) \% P$$

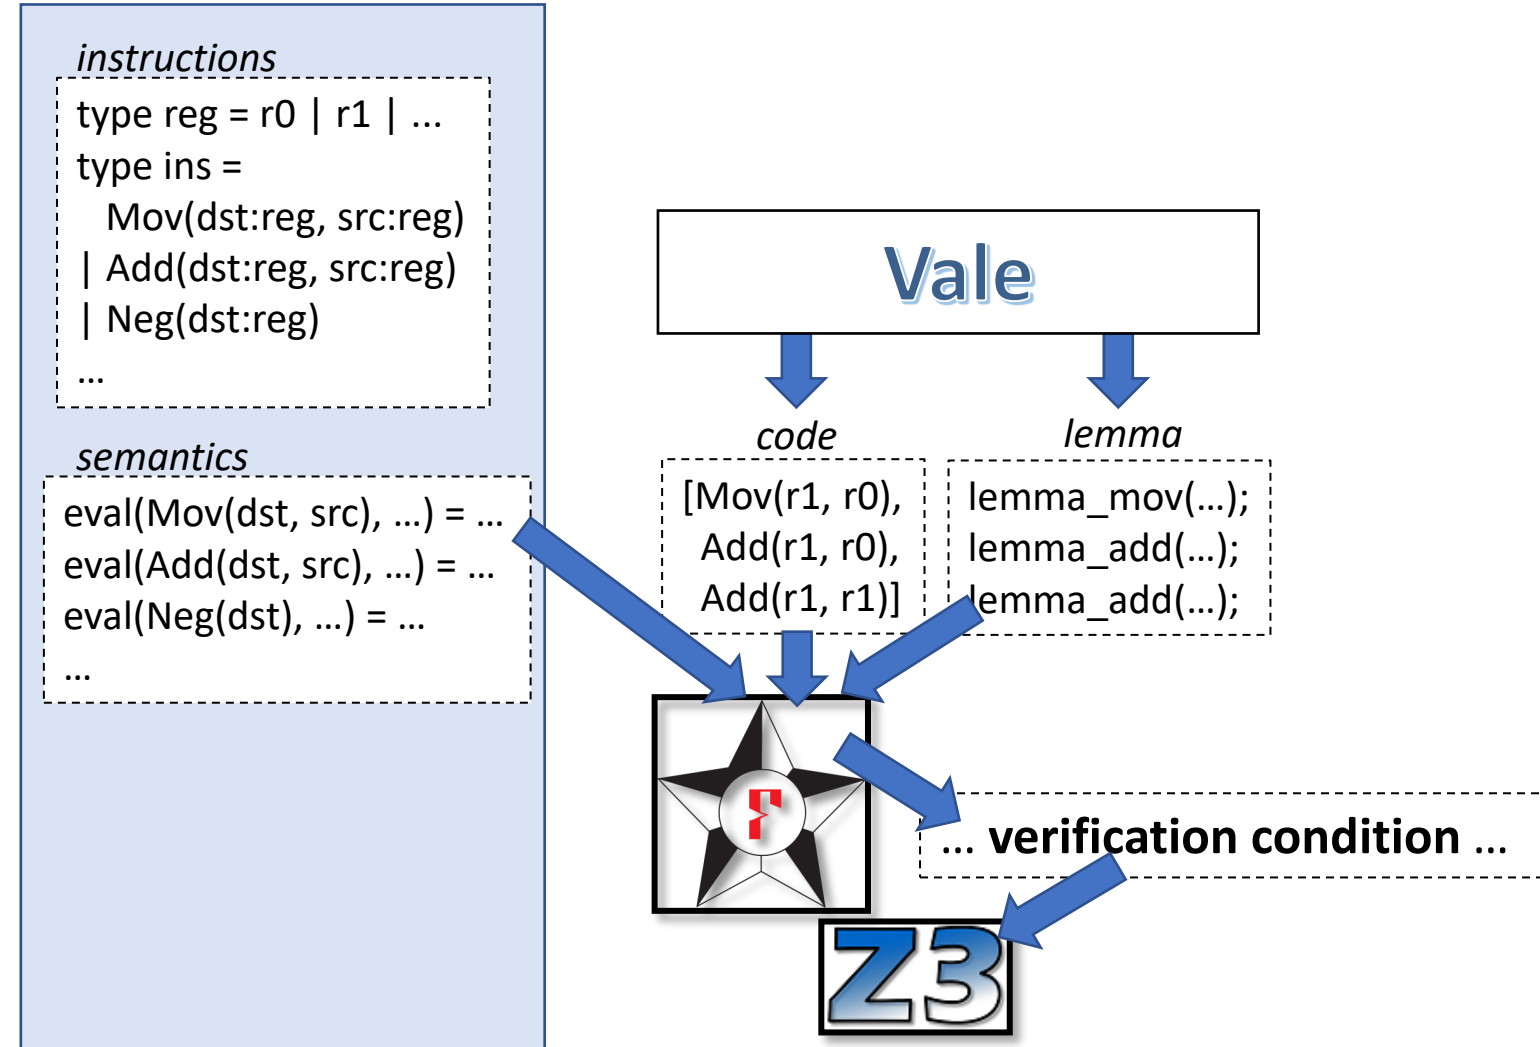
$$\rightarrow ((\text{init} + c_1) * (s^3 \% P) + c_2 * (s^2 \% P) + c_3 * s^1) \% P$$

# Vale: extensible, automated assembly language verification



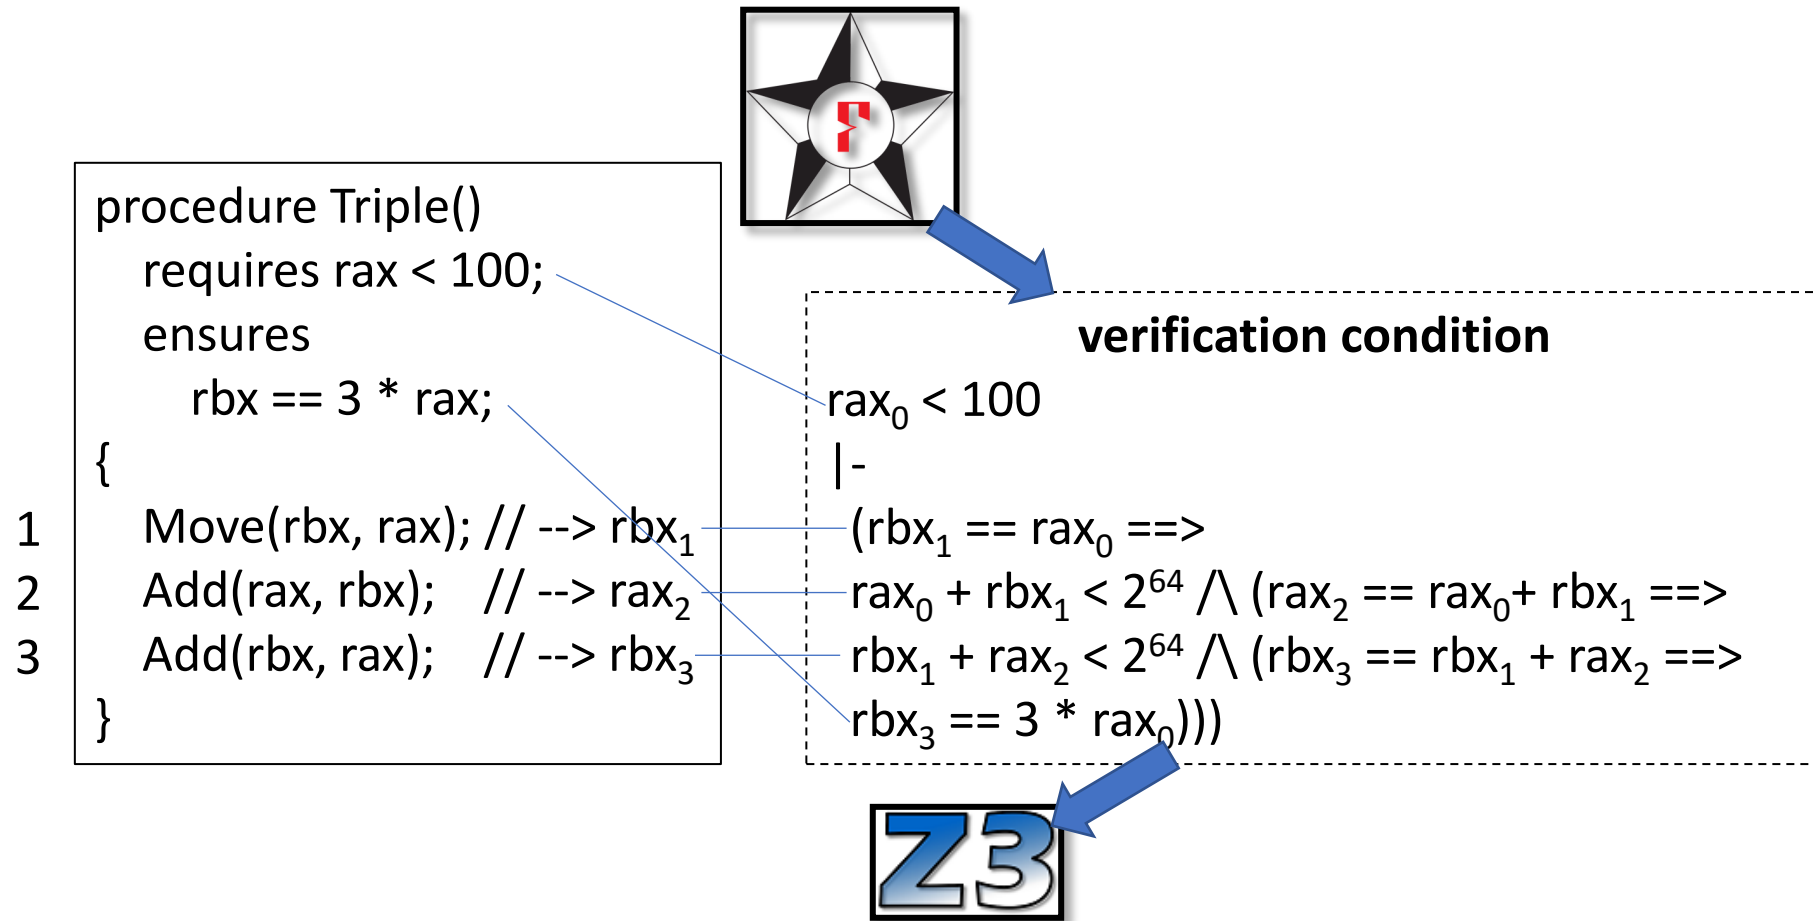
# Vale: extensible, automated assembly language verification

machine model (F\*)





# Verification condition



# Demo

- Verification condition generation for Vale

# Ugh! Default SMT query looks awful!

## verification condition we want:

.....  $(\text{rax}_2 == \text{rax}_0 + \text{rbx}_1 ==>$   
 $\text{rbx}_1 + \text{rax}_2 < 2^{64}$  .....

## verification condition we get:

...  
(forall (ghost\_result\_0:(state \* fuel)).  
 (let (s3, fc3) = ghost\_result\_0 in  
 eval\_code (Ins (Add64 (OReg (Rax)) (OReg (Rbx)))) fc3 s2 == Some s3 /\  
 eval\_operand (OReg Rax) s3 == eval\_operand (OReg Rax) s2 + eval\_operand (OReg Rbx) s2 /\  
 s3 == update\_state (OReg Rax).r s3 s2) ==>  
 lemma\_Add s2 (OReg Rax) (OReg Rbx) == ghost\_result\_0 ==>  
 (forall (s3:state) (fc3:fuel). lemma\_Add s2 (OReg Rax) (OReg Rbx) == Mktuple2 s3 fc3 ==>  
 Cons? codes\_Tuple.tl /\  
 (forall (any\_result0:list code). codes\_Tuple.tl == any\_result0 ==>  
 (forall (any\_result1:list code). codes\_Tuple.tl.tl == any\_result1 ==>  
 OReg? (OReg Rbx) /\ eval\_operand (OReg Rbx) s3 + eval\_operand (OReg Rax) s3 < 2<sup>64</sup>  
 )  
 )  
 )  
...)

# Let's write our own VC generator!

- ??? Maybe like this: ???

I'm lonely  
and sad.



Our own Vale  
VC generator

procedure Triple() ...  
...

**verification condition we want:**

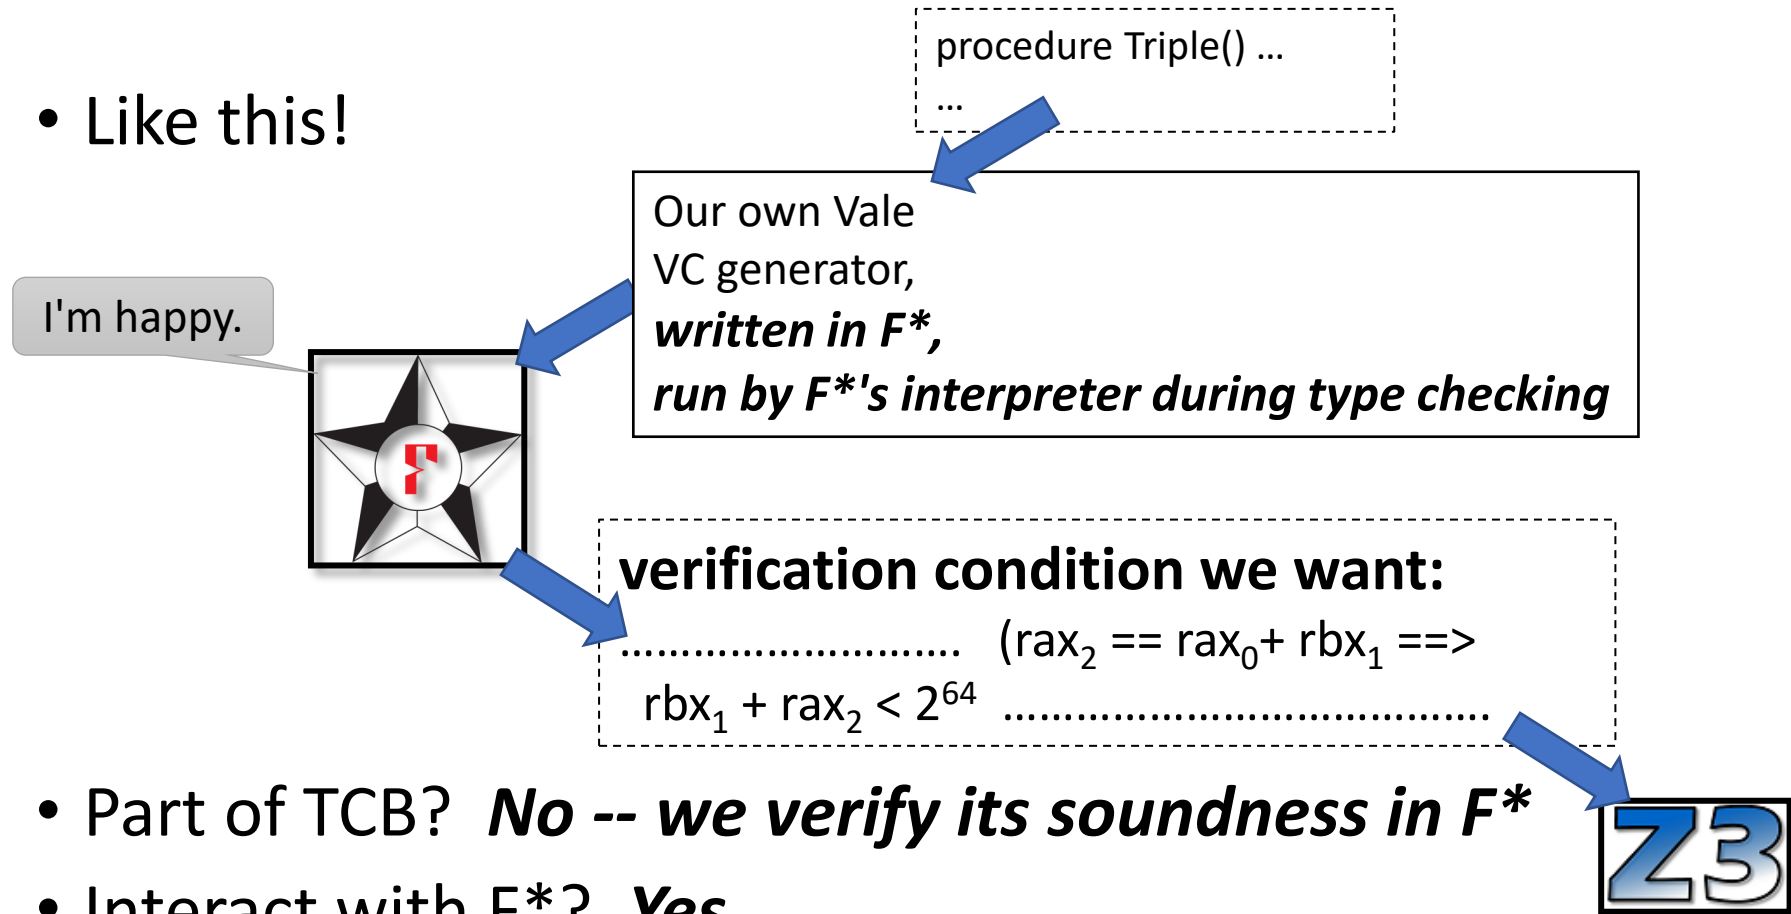
..... (rax<sub>2</sub> == rax<sub>0</sub> + rbx<sub>1</sub> ==>  
rbx<sub>1</sub> + rax<sub>2</sub> < 2<sup>64</sup> .....



- But won't it be part of TCB?
- And how do we interact with F\*?
- Can we reuse F\* features and libraries?

# Let's write our own VC generator!

- Like this!



- Part of TCB? **No -- we verify its soundness in  $F^*$**
- Interact with  $F^*$ ? **Yes**
- Reuse  $F^*$  features and libraries? **Yes**

# Let's write our own VC generator!

