

Fast verified
post-quantum software,
part 1: RAM subroutines

D. J. Bernstein

Performance pressure \Rightarrow
tons of new crypto software \Rightarrow
many mistakes passing tests \Rightarrow
frequent security disasters.

e.g. 2019.06 “Warning: Google
Researcher Drops Windows
10 Zero-Day Security Bomb” :
modular inverse.

e.g. 2019.09 “Produced signatures
were valid but leaked information
on the private key” : Falcon.

e.g. 2019.10 “Minerva attack can
recover private keys from smart
cards, cryptographic libraries” .

e.g. 2020.08 “A key-recovery
timing attack on . . . FrodoKEM” .

e.g. 2020.12 “It looks like the
FrodoKEM team also fixed the
timing oracle [GJN20] badly and
caused a more serious security
problem while trying to do that.”

ified
antum software,
RAM subroutines

ernstein

ance pressure ⇒

new crypto software ⇒

mistakes passing tests ⇒
security disasters.

9.06 “Warning: Google

ner Drops Windows

-Day Security Bomb”:

inverse.

1

e.g. 2019.09 “Produced signatures were valid but leaked information on the private key”: Falcon.

e.g. 2019.10 “Minerva attack can recover private keys from smart cards, cryptographic libraries”.

e.g. 2020.08 “A key-recovery timing attack on ... FrodoKEM”.

e.g. 2020.12 “It looks like the FrodoKEM team also fixed the timing oracle [GJN20] badly and caused a more serious security problem while trying to do that.”

2

Many fu

Keccak

“Keccak

>20 opt

of Kecca

Also, for

many fu

1

ware,
outines

ure ⇒

o software ⇒

ssing tests ⇒

disasters.

arning: Google

Windows

urity Bomb” :

e.g. 2019.09 “Produced signatures were valid but leaked information on the private key” : Falcon.

e.g. 2019.10 “Minerva attack can recover private keys from smart cards, cryptographic libraries” .

e.g. 2020.08 “A key-recovery timing attack on ... FrodoKEM” .

e.g. 2020.12 “It looks like the FrodoKEM team also fixed the timing oracle [GJN20] badly and caused a more serious security problem while trying to do that.”

2

Many functions ×

Keccak (SHA-3) t

“Keccak Code Pac

>20 optimized im

of Keccak: AVX2,

Also, for “parallel

many further impl

1

e.g. 2019.09 “Produced signatures were valid but leaked information on the private key”: Falcon.

e.g. 2019.10 “Minerva attack can recover private keys from smart cards, cryptographic libraries” .

e.g. 2020.08 “A key-recovery timing attack on ... FrodoKEM” .

e.g. 2020.12 “It looks like the FrodoKEM team also fixed the timing oracle [GJN20] badly and caused a more serious security problem while trying to do that.”

2

Many functions × many CP

Keccak (SHA-3) team main
“Keccak Code Package” wit
>20 optimized implementat
of Keccak: AVX2, NEON, e
Also, for “parallel Keccak” ,
many further implementatio

e.g. 2019.09 “Produced signatures were valid but leaked information on the private key”: Falcon.

e.g. 2019.10 “Minerva attack can recover private keys from smart cards, cryptographic libraries” .

e.g. 2020.08 “A key-recovery timing attack on ... FrodoKEM” .

e.g. 2020.12 “It looks like the FrodoKEM team also fixed the timing oracle [GJN20] badly and caused a more serious security problem while trying to do that.”

Many functions × many CPUs

Keccak (SHA-3) team maintains “Keccak Code Package” with >20 optimized implementations of Keccak: AVX2, NEON, etc. Also, for “parallel Keccak”, many further implementations.

e.g. 2019.09 “Produced signatures were valid but leaked information on the private key”: Falcon.

e.g. 2019.10 “Minerva attack can recover private keys from smart cards, cryptographic libraries” .

e.g. 2020.08 “A key-recovery timing attack on . . . FrodoKEM” .

e.g. 2020.12 “It looks like the FrodoKEM team also fixed the timing oracle [GJN20] badly and caused a more serious security problem while trying to do that.”

Many functions × many CPUs

Keccak (SHA-3) team maintains “Keccak Code Package” with >20 optimized implementations of Keccak: AVX2, NEON, etc. Also, for “parallel Keccak”, many further implementations.

Why not portable C code using “optimizing” compiler? Slower.

e.g. 2019.09 “Produced signatures were valid but leaked information on the private key”: Falcon.

e.g. 2019.10 “Minerva attack can recover private keys from smart cards, cryptographic libraries” .

e.g. 2020.08 “A key-recovery timing attack on ... FrodoKEM” .

e.g. 2020.12 “It looks like the FrodoKEM team also fixed the timing oracle [GJN20] badly and caused a more serious security problem while trying to do that.”

Many functions × many CPUs

Keccak (SHA-3) team maintains “Keccak Code Package” with >20 optimized implementations of Keccak: AVX2, NEON, etc. Also, for “parallel Keccak”, many further implementations.

Why not portable C code using “optimizing” compiler? Slower.

Post-quantum crypto is going down same path: AVX2, ARM Cortex-M4, Cortex-A7, Cortex-A53, Zen, AVX-512, RISC-V, ...

9.09 “Produced signatures
id but leaked information
private key”: Falcon.

9.10 “Minerva attack can
private keys from smart
cryptographic libraries”.

10.08 “A key-recovery
attack on ... FrodoKEM”.

10.12 “It looks like the
EM team also fixed the
oracle [GJN20] badly and
a more serious security
while trying to do that.”

2

Many functions × many CPUs

Keccak (SHA-3) team maintains
“Keccak Code Package” with
>20 optimized implementations
of Keccak: AVX2, NEON, etc.
Also, for “parallel Keccak”,
many further implementations.

Why not portable C code using
“optimizing” compiler? Slower.

Post-quantum crypto is going
down same path: AVX2, ARM
Cortex-M4, Cortex-A7, Cortex-
A53, Zen, AVX-512, RISC-V, ...

3

Some go

For some
and some

Without
can have
that the
does wh

2

duced signatures
ked information
": Falcon.

erva attack can
ys from smart
ic libraries".

ey-recovery

.. FrodoKEM".

ooks like the

also fixed the

[20] badly and

ious security

ng to do that."

Many functions × many CPUs

Keccak (SHA-3) team maintains
"Keccak Code Package" with
>20 optimized implementations
of Keccak: AVX2, NEON, etc.
Also, for "parallel Keccak",
many further implementations.

Why not portable C code using
"optimizing" compiler? Slower.

Post-quantum crypto is going
down same path: AVX2, ARM
Cortex-M4, Cortex-A7, Cortex-
A53, Zen, AVX-512, RISC-V, ...

3

Some good news

For some types of
and some types of

Without insane lev
can have an autom
that the optimized
does what the spe

2

Many functions × many CPUs

Keccak (SHA-3) team maintains “Keccak Code Package” with >20 optimized implementations of Keccak: AVX2, NEON, etc. Also, for “parallel Keccak”, many further implementations.

Why not portable C code using “optimizing” compiler? Slower.

Post-quantum crypto is going down same path: AVX2, ARM Cortex-M4, Cortex-A7, Cortex-A53, Zen, AVX-512, RISC-V, ...

3

Some good news

For some types of optimized and some types of specs:

Without insane levels of effort can have an automated guarantee that the optimized code does what the spec says.

Many functions × many CPUs

Keccak (SHA-3) team maintains “Keccak Code Package” with >20 optimized implementations of Keccak: AVX2, NEON, etc. Also, for “parallel Keccak”, many further implementations.

Why not portable C code using “optimizing” compiler? Slower.

Post-quantum crypto is going down same path: AVX2, ARM Cortex-M4, Cortex-A7, Cortex-A53, Zen, AVX-512, RISC-V, ...

Some good news

For some types of optimized code, and some types of specs:

Without insane levels of effort, can have an automated guarantee that the optimized code does what the spec says.

Many functions × many CPUs

Keccak (SHA-3) team maintains “Keccak Code Package” with >20 optimized implementations of Keccak: AVX2, NEON, etc. Also, for “parallel Keccak”, many further implementations.

Why not portable C code using “optimizing” compiler? Slower.

Post-quantum crypto is going down same path: AVX2, ARM Cortex-M4, Cortex-A7, Cortex-A53, Zen, AVX-512, RISC-V, ...

Some good news

For some types of optimized code, and some types of specs:

Without insane levels of effort, can have an automated guarantee that the optimized code does what the spec says.

Security reviewer still has to check whether the spec is secure *and* has to check for bugs in the verification tools—but saves tons of time in checking code optimized for each CPU.

nctions × many CPUs

(SHA-3) team maintains
“Code Package” with
optimized implementations
for: AVX2, NEON, etc.
“parallel Keccak”,
further implementations.

portable C code using
“gcc” compiler? Slower.

quantum crypto is going
same path: AVX2, ARM
M4, Cortex-A7, Cortex-
M4, AVX-512, RISC-V, ...

3

Some good news

For some types of optimized code,
and some types of specs:

Without insane levels of effort,
can have an automated guarantee
that the optimized code
does what the spec says.

Security reviewer still has to
check whether the spec is secure
and has to check for
bugs in the verification tools—
but saves tons of time in checking
code optimized for each CPU.

4

What ex

Starting
algorithm
Why do
written i

3

many CPUs

Team maintains
“package” with
implementations
NEON, etc.
“Keccak”,
implementations.

C code using
compiler? Slower.

Opto is going
AVX2, ARM
Cortex-A7, Cortex-
A52, RISC-V, ...

Some good news

For some types of optimized code,
and some types of specs:

Without insane levels of effort,
can have an automated guarantee
that the optimized code
does what the spec says.

Security reviewer still has to
check whether the spec is secure
and has to check for
bugs in the verification tools—
but saves tons of time in checking
code optimized for each CPU.

4

What exactly is “t

Starting in 1960, C
algorithms—written
Why do we tolerate
written in English

3

Some good news

For some types of optimized code, and some types of specs:

Without insane levels of effort, can have an automated guarantee that the optimized code does what the spec says.

Security reviewer still has to check whether the spec is secure *and* has to check for bugs in the verification tools—but saves tons of time in checking code optimized for each CPU.

4

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL. Why do we tolerate algorithms written in English “pseudocode”?

Some good news

For some types of optimized code, and some types of specs:

Without insane levels of effort, can have an automated guarantee that the optimized code does what the spec says.

Security reviewer still has to check whether the spec is secure *and* has to check for bugs in the verification tools—but saves tons of time in checking code optimized for each CPU.

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

Some good news

For some types of optimized code, and some types of specs:

Without insane levels of effort, can have an automated guarantee that the optimized code does what the spec says.

Security reviewer still has to check whether the spec is secure *and* has to check for bugs in the verification tools—but saves tons of time in checking code optimized for each CPU.

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

Some good news

For some types of optimized code, and some types of specs:

Without insane levels of effort, can have an automated guarantee that the optimized code does what the spec says.

Security reviewer still has to check whether the spec is secure *and* has to check for bugs in the verification tools—but but saves tons of time in checking code optimized for each CPU.

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being (1) easy to read, (2) executable.

Verify spec = ref = avx2 = ...

Security reviewers focus on spec.

Good news

the types of optimized code,
the types of specs:

insane levels of effort,
an automated guarantee
optimized code
at the spec says.

reviewer still has to
whether the spec is secure
to check for

the verification tools—
tons of time in checking
optimized for each CPU.

4

What exactly is “the spec”?

Starting in 1960, CACM published
algorithms—written in ALGOL.
Why do we tolerate algorithms
written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being
(1) easy to read, (2) executable.

Verify spec = ref = avx2 = ...

Security reviewers focus on spec.

5

Case stu

Many al
CPU RA
secret ac
Can we

4

optimized code,
specs:

levels of effort,
guarantee

code

c says.

still has to

spec is secure

for

ation tools—

time in checking

each CPU.

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being
(1) easy to read, (2) executable.

Verify spec = ref = avx2 = ...

Security reviewers focus on spec.

5

Case study: RAM

Many algorithms r
CPU RAM instruc

secret addresses th

Can we eliminate

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being (1) easy to read, (2) executable.

Verify $\text{spec} = \text{ref} = \text{avx2} = \dots$.

Security reviewers focus on spec.

Case study: RAM subroutines

Many algorithms rely on RAM. CPU RAM instructions leak

secret addresses through timing

Can we eliminate timing leaks?

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being (1) easy to read, (2) executable.

Verify spec = ref = avx2 = ...

Security reviewers focus on spec.

Case study: RAM subroutines

Many algorithms rely on RAM. CPU RAM instructions leak secret addresses through timing. Can we eliminate timing leaks?

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being
(1) easy to read, (2) executable.

Verify spec = ref = avx2 = ...

Security reviewers focus on spec.

Case study: RAM subroutines

Many algorithms rely on RAM.
CPU RAM instructions leak
secret addresses through timing.
Can we eliminate timing leaks?

Yes! Replace CPU RAM insns
with software to simulate RAM.

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being (1) easy to read, (2) executable.

Verify spec = ref = avx2 = ...

Security reviewers focus on spec.

Case study: RAM subroutines

Many algorithms rely on RAM. CPU RAM instructions leak secret addresses through timing. Can we eliminate timing leaks?

Yes! Replace CPU RAM insns with software to simulate RAM.

Speedup #1: Use sorting to efficiently simulate *parallel* RAM.

What exactly is “the spec”?

Starting in 1960, CACM published algorithms—written in ALGOL.

Why do we tolerate algorithms written in English “pseudocode”?

“Easier to read than ref”:

that’s because ref

- was forced to be in C,
- often tries to be constant time,
- sometimes tries to be fast.

No conflict between spec being (1) easy to read, (2) executable.

Verify spec = ref = avx2 = ...

Security reviewers focus on spec.

Case study: RAM subroutines

Many algorithms rely on RAM. CPU RAM instructions leak secret addresses through timing. Can we eliminate timing leaks?

Yes! Replace CPU RAM insns with software to simulate RAM.

Speedup #1: Use sorting to efficiently simulate *parallel* RAM.

Speedup #2: Sometimes same permutation is applied to many inputs. Precompute “control bits” for permutation.

Exactly is “the spec”?

in 1960, CACM published
 programs—written in ALGOL.

we tolerate algorithms
 in English “pseudocode”?

to read than ref”:

because ref

forced to be in C,

tries to be constant time,
 sometimes tries to be fast.

conflict between spec being

to read, (2) executable.

spec = ref = avx2 = ...

reviewers focus on spec.

Case study: RAM subroutines

Many algorithms rely on RAM.
 CPU RAM instructions leak
 secret addresses through timing.
 Can we eliminate timing leaks?

Yes! Replace CPU RAM insns
 with software to simulate RAM.

Speedup #1: Use sorting to
 efficiently simulate *parallel* RAM.

Speedup #2: Sometimes
 same permutation is applied
 to many inputs. Precompute
 “control bits” for permutation.

2018 Be
 for sortin
 Verified

2020 Be
 for const
 HOL Lig

Coming
 of the p

This soft
 inside cu
 for Class
 permuta
 NTRU P

5

the spec”?

CACM published
in ALGOL.

te algorithms

“pseudocode”?

an ref”:

f

e in C,

constant time,
to be fast.

en spec being

(2) executable.

$\text{f} = \text{avx2} = \dots$

focus on spec.

Case study: RAM subroutines

Many algorithms rely on RAM.

CPU RAM instructions leak

secret addresses through timing.

Can we eliminate timing leaks?

Yes! Replace CPU RAM insns

with software to simulate RAM.

Speedup #1: Use sorting to
efficiently simulate *parallel* RAM.

Speedup #2: Sometimes
same permutation is applied
to many inputs. Precompute
“control bits” for permutation.

6

2018 Bernstein: sp

for sorting integer

Verified constant-t

2020 Bernstein: sp

for constant-time

HOL Light proof o

Coming soon: veri

of the permutation

This software is al

inside current softw

for Classic McElie

permutations), NT

NTRU Prime (sort

5

Case study: RAM subroutines

Many algorithms rely on RAM.
CPU RAM instructions leak
secret addresses through timing.
Can we eliminate timing leaks?

Yes! Replace CPU RAM insns
with software to simulate RAM.

Speedup #1: Use sorting to
efficiently simulate *parallel* RAM.

Speedup #2: Sometimes
same permutation is applied
to many inputs. Precompute
“control bits” for permutation.

6

2018 Bernstein: speed record
for sorting integer arrays.

Verified constant-time software

2020 Bernstein: speed record
for constant-time permutation

HOL Light proof of algorithm

Coming soon: verification
of the permutation software

This software is already used
inside current software releases
for Classic McEliece (sorting
permutations), NTRU (sorting
NTRU Prime (sorting)).

Case study: RAM subroutines

Many algorithms rely on RAM.
CPU RAM instructions leak
secret addresses through timing.
Can we eliminate timing leaks?

Yes! Replace CPU RAM insns
with software to simulate RAM.

Speedup #1: Use sorting to
efficiently simulate *parallel* RAM.

Speedup #2: Sometimes
same permutation is applied
to many inputs. Precompute
“control bits” for permutation.

2018 Bernstein: speed records
for sorting integer arrays.

Verified constant-time software.

2020 Bernstein: speed records
for constant-time permutations.

HOL Light proof of algorithm.

Coming soon: verification
of the permutation software.

This software is already used
inside current software releases
for Classic McEliece (sorting and
permutations), NTRU (sorting),
NTRU Prime (sorting).

Why: RAM subroutines

algorithms rely on RAM.

RAM instructions leak

addresses through timing.

eliminate timing leaks?

replace CPU RAM insns

software to simulate RAM.

o #1: Use sorting to

efficiently simulate *parallel* RAM.

o #2: Sometimes

permutation is applied

to inputs. Precompute

“bits” for permutation.

6

2018 Bernstein: speed records
for sorting integer arrays.

Verified constant-time software.

2020 Bernstein: speed records
for constant-time permutations.

HOL Light proof of algorithm.

Coming soon: verification
of the permutation software.

This software is already used
inside current software releases
for Classic McEliece (sorting and
permutations), NTRU (sorting),
NTRU Prime (sorting).

7

The con

Imagine

automat

fast bina

“Compil

prove th

always w

If all of -

6

subroutines

ely on RAM.

tions leak

rough timing.

timing leaks?

RAM insns

simulate RAM.

sorting to

parallel RAM.

netimes

is applied

recompute

permutation.

2018 Bernstein: speed records
for sorting integer arrays.

Verified constant-time software.

2020 Bernstein: speed records
for constant-time permutations.

HOL Light proof of algorithm.

Coming soon: verification
of the permutation software.

This software is already used
inside current software releases
for Classic McEliece (sorting and
permutations), NTRU (sorting),
NTRU Prime (sorting).

7

The conventional

Imagine an optimi
automatically conv

fast binary for whi

“Compiler verificat

prove that the com

always works corre

If all of this is don

6

2018 Bernstein: speed records
for sorting integer arrays.

Verified constant-time software.

2020 Bernstein: speed records
for constant-time permutations.

HOL Light proof of algorithm.

Coming soon: verification
of the permutation software.

This software is already used
inside current software releases
for Classic McEliece (sorting and
permutations), NTRU (sorting),
NTRU Prime (sorting).

7

The conventional path

Imagine an optimizing compiler
automatically converting source code
to fast binary for whichever CPU

“Compiler verification”:
prove that the compiler
always works correctly.

If all of this is done, great!

2018 Bernstein: speed records
for sorting integer arrays.
Verified constant-time software.

2020 Bernstein: speed records
for constant-time permutations.
HOL Light proof of algorithm.

Coming soon: verification
of the permutation software.

This software is already used
inside current software releases
for Classic McEliece (sorting and
permutations), NTRU (sorting),
NTRU Prime (sorting).

The conventional path

Imagine an optimizing compiler
automatically converting spec →
fast binary for whichever CPU.

“Compiler verification”:
prove that the compiler
always works correctly.

If all of this is done, great!

2018 Bernstein: speed records
for sorting integer arrays.
Verified constant-time software.

2020 Bernstein: speed records
for constant-time permutations.
HOL Light proof of algorithm.

Coming soon: verification
of the permutation software.

This software is already used
inside current software releases
for Classic McEliece (sorting and
permutations), NTRU (sorting),
NTRU Prime (sorting).

The conventional path

Imagine an optimizing compiler
automatically converting spec →
fast binary for whichever CPU.

“Compiler verification”:
prove that the compiler
always works correctly.

If all of this is done, great!

Reality: Again, look at Keccak.

Speedups >
automated speedups >
verified automated speedups.

rnstein: speed records
ng integer arrays.

constant-time software.

rnstein: speed records
stant-time permutations.
ght proof of algorithm.

soon: verification
ermutation software.

ftware is already used
urrent software releases
sic McEliece (sorting and
tions), NTRU (sorting),
Prime (sorting).

The conventional path

Imagine an optimizing compiler
automatically converting spec \rightarrow
fast binary for whichever CPU.

“Compiler verification”:
prove that the compiler
always works correctly.

If all of this is done, great!

Reality: Again, look at Keccak.

Speedups $>$
automated speedups $>$
verified automated speedups.

Verifying

Optimiza
spec \rightarrow
opt4 \rightarrow
Some m
CPUs sh

7

The conventional path

Imagine an optimizing compiler automatically converting spec → fast binary for whichever CPU.

“Compiler verification” : prove that the compiler always works correctly.

If all of this is done, great!

Reality: Again, look at Keccak.

Speedups >
automated speedups >
verified automated speedups.

8

Verifying fast software

Optimization experience
spec → opt → opt1
opt2 → opt3 → opt4 → opt5 → ...
Some manual steps
CPUs share some

7

The conventional path

Imagine an optimizing compiler automatically converting spec \rightarrow fast binary for whichever CPU.

“Compiler verification” :
prove that the compiler always works correctly.

If all of this is done, great!

Reality: Again, look at Keccak.

Speedups $>$
automated speedups $>$
verified automated speedups.

8

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx

Some manual steps, some to
CPUs share some steps.

The conventional path

Imagine an optimizing compiler automatically converting spec \rightarrow fast binary for whichever CPU.

“Compiler verification” :
prove that the compiler always works correctly.

If all of this is done, great!

Reality: Again, look at Keccak.

Speedups $>$
automated speedups $>$
verified automated speedups.

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Some manual steps, some tools.

CPUs share some steps.

The conventional path

Imagine an optimizing compiler automatically converting spec \rightarrow fast binary for whichever CPU.

“Compiler verification”:
prove that the compiler always works correctly.

If all of this is done, great!

Reality: Again, look at Keccak.

Speedups $>$
automated speedups $>$
verified automated speedups.

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Some manual steps, some tools.
CPUs share some steps.

“Translation validation”:
verify equivalence of
tool output to tool input.
Doesn't require verifying
that the tool *always* works.

“Transformation verification”:
verify equivalence of
manual output to manual input.

Conventional path

an optimizing compiler
typically converting spec \rightarrow
binary for whichever CPU.

“no verification”:
that the compiler
works correctly.

if this is done, great!

Again, look at Keccak.

10x >

10x speedups >

10x automated speedups.

8

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow ... \rightarrow avx2.

Some manual steps, some tools.
CPUs share some steps.

“Translation validation”:

verify equivalence of
tool output to tool input.

Doesn't require verifying
that the tool *always* works.

“Transformation verification”:

verify equivalence of
manual output to manual input.

9

Allowing

For verification
spec \leftrightarrow
verif3

Don't try to
match the

spec \rightarrow
opt4 \rightarrow

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Some manual steps, some tools.
CPUs share some steps.

“Translation validation”:

verify equivalence of
tool output to tool input.

Doesn't require verifying
that the tool *always* works.

“Transformation verification”:

verify equivalence of
manual output to manual input.

Allowing new verif

For verification, su

spec \leftrightarrow verif \leftrightarrow
verif3 \leftrightarrow \dots \leftrightarrow a

Don't try to force
match the develop

spec \rightarrow opt \rightarrow op
opt4 \rightarrow opt5 \rightarrow .

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Some manual steps, some tools.

CPUs share some steps.

“Translation validation”:

verify equivalence of
tool output to tool input.

Doesn't require verifying
that the tool *always* works.

“Transformation verification”:

verify equivalence of
manual output to manual input.

Allowing new verification ch

For verification, suffices to b

spec \leftrightarrow verif \leftrightarrow verif2 \leftrightarrow
verif3 \leftrightarrow \dots \leftrightarrow avx2.

Don't try to force this chain

match the development path

spec \rightarrow opt \rightarrow opt2 \rightarrow opt
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Some manual steps, some tools.
CPUs share some steps.

“Translation validation”:

verify equivalence of
tool output to tool input.

Doesn't require verifying
that the tool *always* works.

“Transformation verification”:

verify equivalence of
manual output to manual input.

Allowing new verification chains

For verification, suffices to build
spec \leftrightarrow verif \leftrightarrow verif2 \leftrightarrow
verif3 \leftrightarrow \dots \leftrightarrow avx2.

Don't try to force this chain to
match the development path

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Verifying fast software

Optimization experts:

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Some manual steps, some tools.
CPUs share some steps.

“Translation validation”:

verify equivalence of
tool output to tool input.

Doesn't require verifying
that the tool *always* works.

“Transformation verification”:

verify equivalence of
manual output to manual input.

Allowing new verification chains

For verification, suffices to build
spec \leftrightarrow verif \leftrightarrow verif2 \leftrightarrow
verif3 \leftrightarrow \dots \leftrightarrow avx2.

Don't try to force this chain to
match the development path

spec \rightarrow opt \rightarrow opt2 \rightarrow opt3 \rightarrow
opt4 \rightarrow opt5 \rightarrow \dots \rightarrow avx2.

Separation promotes independent
speedups in (1) the development
process and (2) the verification
process: e.g., vectorization is
often challenging for development
but trivial for verification.