

Error-prone cryptographic designs

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

*“The poor user is
given enough rope with which
to hang himself—something
a standard should not do.”*

—1992 Rivest,
commenting on nonce generation
inside Digital Signature Algorithm
(1991 proposal by NIST,
1992 credited to NSA,
1994 standardized by NIST)

Crypto horror story #1

2010 Bushing–Marcan–Segher–
Sven “failOverflow” demolition
of Sony PS3 security system:
Sony had ignored requirement
to generate new random nonce
for each ECDSA signature.
⇒ Sony’s signatures leaked
Sony’s secret code-signing key.

Error-prone cryptographic designs

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

*“The poor user is
given enough rope with which
to hang himself—something
a standard should not do.”*

—1992 Rivest,
commenting on nonce generation
inside Digital Signature Algorithm
(1991 proposal by NIST,
1992 credited to NSA,
1994 standardized by NIST)

Crypto horror story #1

2010 Bushing–Marcan–Segher–
Sven “failOverflow” demolition
of Sony PS3 security system:
Sony had ignored requirement
to generate new random nonce
for each ECDSA signature.

⇒ Sony’s signatures leaked
Sony’s secret code-signing key.

Traditional response: Blame Sony.
Blame the crypto implementor.

Error-prone cryptographic designs

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

*“The poor user is
given enough rope with which
to hang himself—something
a standard should not do.”*

—1992 Rivest,
commenting on nonce generation
inside Digital Signature Algorithm
(1991 proposal by NIST,
1992 credited to NSA,
1994 standardized by NIST)

Crypto horror story #1

2010 Bushing–Marcan–Segher–
Sven “failOverflow” demolition
of Sony PS3 security system:
Sony had ignored requirement
to generate new random nonce
for each ECDSA signature.

⇒ Sony’s signatures leaked
Sony’s secret code-signing key.

Traditional response: Blame Sony.
Blame the crypto implementor.

Rivest’s response: Blame DSA.
Blame the crypto designer.

Error-prone cryptographic designs

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

*“The poor user is
given enough rope with which
to hang himself—something
a standard should not do.”*

—1992 Rivest,
commenting on nonce generation
inside Digital Signature Algorithm
(1991 proposal by NIST,
1992 credited to NSA,
1994 standardized by NIST)

Crypto horror story #1

2010 Bushing–Marcan–Segher–
Sven “failOverflow” demolition
of Sony PS3 security system:
Sony had ignored requirement
to generate new random nonce
for each ECDSA signature.

⇒ Sony’s signatures leaked
Sony’s secret code-signing key.

Traditional response: Blame Sony.
Blame the crypto implementor.

Rivest’s response: Blame DSA.
Blame the crypto designer.

Change DSA to avoid this pitfall!

one cryptographic designs

. Bernstein

ty of Illinois at Chicago &

che Universiteit Eindhoven

oor user is

ough rope with which

himself—something

ard should not do.”

Rivest,

ting on nonce generation

igital Signature Algorithm

proposal by NIST,

edited to NSA,

andardized by NIST)

Crypto horror story #1

2010 Bushing–Marcan–Segher–

Sven “failOverflow” demolition

of Sony PS3 security system:

Sony had ignored requirement

to generate new random nonce

for each ECDSA signature.

⇒ Sony’s signatures leaked

Sony’s secret code-signing key.

Traditional response: Blame Sony.

Blame the crypto implementor.

Rivest’s response: Blame DSA.

Blame the crypto designer.

Change DSA to avoid this pitfall!

Crypto h

2005 Os

65ms to

used for

Attack p

but with

Almost a

use fast

Kernel’s

influence

influenci

influenci

of the at

65ms to

ographic designs

n
is at Chicago &
siteit Eindhoven

*e with which
something
not do."*

once generation
ature Algorithm
NIST,
ISA,
by NIST)

Crypto horror story #1

2010 Bushing–Marcan–Segher–
Sven “failOverflow” demolition
of Sony PS3 security system:
Sony had ignored requirement
to generate new random nonce
for each ECDSA signature.

⇒ Sony’s signatures leaked

Sony’s secret code-signing key.

Traditional response: Blame Sony.

Blame the crypto implementor.

Rivest’s response: Blame DSA.

Blame the crypto designer.

Change DSA to avoid this pitfall!

Crypto horror story

2005 Osvik–Shamir
65ms to steal Linux
used for hard-disk
Attack process on
but without privile

Almost all AES im
use fast lookup ta
Kernel’s secret AE
influences table-lo
influencing CPU c
influencing measur
of the attack proc
65ms to compute

esigns

ago &
hoven

ch

ation
orithm

Crypto horror story #1

2010 Bushing–Marcan–Segher–
Sven “failOverflow” demolition
of Sony PS3 security system:

Sony had ignored requirement
to generate new random nonce
for each ECDSA signature.

⇒ Sony’s signatures leaked

Sony’s secret code-signing key.

Traditional response: Blame Sony.

Blame the crypto implementor.

Rivest’s response: Blame DSA.

Blame the crypto designer.

Change DSA to avoid this pitfall!

Crypto horror story #2

2005 Osvik–Shamir–Tromer:
65ms to steal Linux AES key
used for hard-disk encryption
Attack process on same CPU
but without privileges.

Almost all AES implementat
use fast lookup tables.

Kernel’s secret AES key
influences table-load address
influencing CPU cache state
influencing measurable timing
of the attack process.

65ms to compute influence

Crypto horror story #1

2010 Bushing–Marcan–Segher–Sven “fail0verflow” demolition of Sony PS3 security system: Sony had ignored requirement to generate new random nonce for each ECDSA signature.
⇒ Sony’s signatures leaked Sony’s secret code-signing key.

Traditional response: Blame Sony.
Blame the crypto implementor.

Rivest’s response: Blame DSA.
Blame the crypto designer.
Change DSA to avoid this pitfall!

Crypto horror story #2

2005 Osvik–Shamir–Tromer: 65ms to steal Linux AES key used for hard-disk encryption. Attack process on same CPU but without privileges.

Almost all AES implementations use fast lookup tables. Kernel’s secret AES key influences table-load addresses, influencing CPU cache state, influencing measurable timings of the attack process. 65ms to compute influence⁻¹.

Horror story #1

“Flashing–Marcan–Segher–
MailOverflow” demolition

PS3 security system:

did ignored requirement

to generate new random nonce

for ECDSA signature.

Attacker's signatures leaked

secret code-signing key.

Official response: Blame Sony.

Blame the crypto implementor.

Alternative response: Blame DSA.

Blame the crypto designer.

Switch to DSA to avoid this pitfall!

Crypto horror story #2

2005 Osvik–Shamir–Tromer:

65ms to steal Linux AES key

used for hard-disk encryption.

Attack process on same CPU

but without privileges.

Almost all AES implementations
use fast lookup tables.

Kernel's secret AES key

influences table-load addresses,

influencing CPU cache state,

influencing measurable timings

of the attack process.

65ms to compute influence⁻¹.

2012 Mc

Shacham

data-cache

x86 proc

somehow

physical

memory

program

Story #1

Arkan–Segher–

“demolition

city system:

requirement

random nonce

signature.

keys leaked

re-signing key.

Case: Blame Sony.

to implementor.

Blame DSA.

to designer.

Avoid this pitfall!

Crypto horror story #2

2005 Osvik–Shamir–Tromer:

65ms to steal Linux AES key

used for hard-disk encryption.

Attack process on same CPU

but without privileges.

Almost all AES implementations

use fast lookup tables.

Kernel’s secret AES key

influences table-load addresses,

influencing CPU cache state,

influencing measurable timings

of the attack process.

65ms to compute influence⁻¹.

2012 Mowery–Kee

Shacham: “We po

data-cache timing

x86 processors tha

somehow subvert

physical indexing,

memory requireme

programs is doome

Crypto horror story #2

2005 Osvik–Shamir–Tromer:
65ms to steal Linux AES key
used for hard-disk encryption.
Attack process on same CPU
but without privileges.

Almost all AES implementations
use fast lookup tables.

Kernel's secret AES key
influences table-load addresses,
influencing CPU cache state,
influencing measurable timings
of the attack process.

65ms to compute influence⁻¹.

2012 Mowery–Keelveedhi–
Shacham: “We posit that a
data-cache timing attack ag
x86 processors that does not
somehow subvert the prefet
physical indexing, and massi
memory requirements of mo
programs is doomed to fail.”

Crypto horror story #2

2005 Osvik–Shamir–Tromer:
65ms to steal Linux AES key
used for hard-disk encryption.
Attack process on same CPU
but without privileges.

Almost all AES implementations
use fast lookup tables.

Kernel's secret AES key
influences table-load addresses,
influencing CPU cache state,
influencing measurable timings
of the attack process.

65ms to compute influence⁻¹.

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

Crypto horror story #2

2005 Osvik–Shamir–Tromer:
65ms to steal Linux AES key
used for hard-disk encryption.
Attack process on same CPU
but without privileges.

Almost all AES implementations
use fast lookup tables.

Kernel's secret AES key
influences table-load addresses,
influencing CPU cache state,
influencing measurable timings
of the attack process.

65ms to compute influence⁻¹.

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

2014 Irazoqui–Inci–Eisenbarth–
Sunar “Wait a minute! A fast,
Cross-VM attack on AES”
recovers “the AES keys
of OpenSSL 1.0.1 running inside
the victim VM” in 60 seconds
despite VMware virtualization.

Horror story #2

Avik–Shamir–Tromer:
steal Linux AES key
hard-disk encryption.
process on same CPU
out privileges.

all AES implementations
lookup tables.

secret AES key
es table-load addresses,
ng CPU cache state,
ng measurable timings
ttack process.
compute influence⁻¹.

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

2014 Irazoqui–Inci–Eisenbarth–
Sunar “Wait a minute! A fast,
Cross-VM attack on AES”
recovers “the AES keys
of OpenSSL 1.0.1 running inside
the victim VM” in 60 seconds
despite VMware virtualization.

After ma
on imple
today we
plagued
Warning

y #2

ir–Tromer:
ix AES key
encryption.
same CPU
ages.

plementations
bles.

S key
ad addresses,
ache state,
rable timings
ess.

influence⁻¹.

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

2014 Irazoqui–Inci–Eisenbarth–
Sunar “Wait a minute! A fast,
Cross-VM attack on AES”
recovers “the AES keys
of OpenSSL 1.0.1 running inside
the victim VM” in 60 seconds
despite VMware virtualization.

After many, many,
on implementation
today we still have
plagued with AES
Warning: more pa

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

2014 Irazoqui–Inci–Eisenbarth–
Sunar “Wait a minute! A fast,
Cross-VM attack on AES”
recovers “the AES keys
of OpenSSL 1.0.1 running inside
the victim VM” in 60 seconds
despite VMware virtualization.

After many, many, many papers
on implementations and attacks
today we still have an ecosystem
plagued with AES vulnerabilities.
Warning: more papers \neq security.

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

2014 Irazoqui–Inci–Eisenbarth–
Sunar “Wait a minute! A fast,
Cross-VM attack on AES”
recovers “the AES keys
of OpenSSL 1.0.1 running inside
the victim VM” in 60 seconds
despite VMware virtualization.

After many, many, many papers
on implementations and attacks,
today we still have an ecosystem
plagued with AES vulnerabilities.
Warning: more papers \neq security.

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

2014 Irazoqui–Inci–Eisenbarth–
Sunar “Wait a minute! A fast,
Cross-VM attack on AES”
recovers “the AES keys
of OpenSSL 1.0.1 running inside
the victim VM” in 60 seconds
despite VMware virtualization.

After many, many, many papers
on implementations and attacks,
today we still have an ecosystem
plagued with AES vulnerabilities.
Warning: more papers \neq security.

AES has a serious conflict
between security, simplicity, speed.
It’s tough to achieve security
while insisting on the AES design
—i.e., blaming the implementor.

2012 Mowery–Keelveedhi–
Shacham: “We posit that any
data-cache timing attack against
x86 processors that does not
somehow subvert the prefetcher,
physical indexing, and massive
memory requirements of modern
programs is doomed to fail.”

2014 Irazoqui–Inci–Eisenbarth–
Sunar “Wait a minute! A fast,
Cross-VM attack on AES”
recovers “the AES keys
of OpenSSL 1.0.1 running inside
the victim VM” in 60 seconds
despite VMware virtualization.

After many, many, many papers
on implementations and attacks,
today we still have an ecosystem
plagued with AES vulnerabilities.
Warning: more papers \neq security.

AES has a serious conflict
between security, simplicity, speed.
It’s tough to achieve security
while insisting on the AES design
—i.e., blaming the implementor.

Allowing the *design* to vary
makes security much easier.
Next-generation ciphers are
naturally constant-time and fast.

Howery–Keelveedhi–
n: “We posit that any
the timing attack against
processors that does not
v subvert the prefetcher,
indexing, and massive
requirements of modern
s is doomed to fail.”

zoqui–Inci–Eisenbarth–
Wait a minute! A fast,
M attack on AES”

“the AES keys
SSL 1.0.1 running inside
m VM” in 60 seconds
VMware virtualization.

After many, many, many papers
on implementations and attacks,
today we still have an ecosystem
plagued with AES vulnerabilities.
Warning: more papers \neq security.

AES has a serious conflict
between security, simplicity, speed.
It’s tough to achieve security
while insisting on the AES design
—i.e., blaming the implementor.

Allowing the *design* to vary
makes security much easier.
Next-generation ciphers are
naturally constant-time and fast.

The big

Cry
des

Cry
impleme

elvedhi—
osit that any
attack against
t does not
the prefetcher,
and massive
ents of modern
ed to fail.”

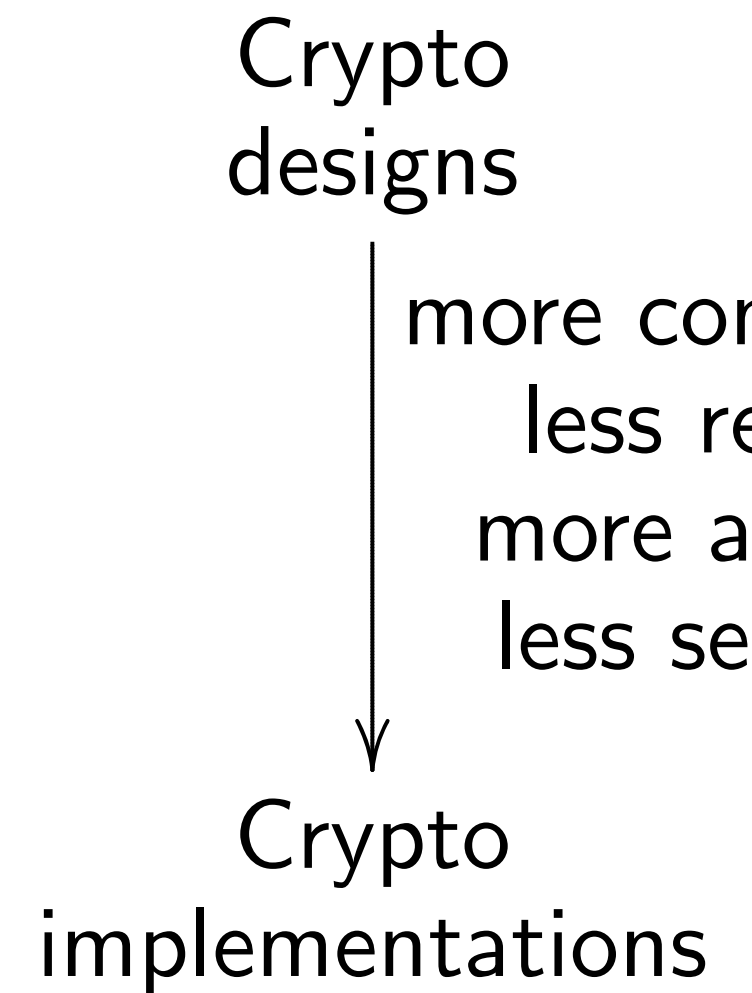
—Eisenbarth—
nute! A fast,
on AES”
keys
running inside
60 seconds
irtualization.

After many, many, many papers
on implementations and attacks,
today we still have an ecosystem
plagued with AES vulnerabilities.
Warning: more papers \neq security.

AES has a serious conflict
between security, simplicity, speed.
It’s tough to achieve security
while insisting on the AES design
—i.e., blaming the implementor.

Allowing the *design* to vary
makes security much easier.
Next-generation ciphers are
naturally constant-time and fast.

The big picture

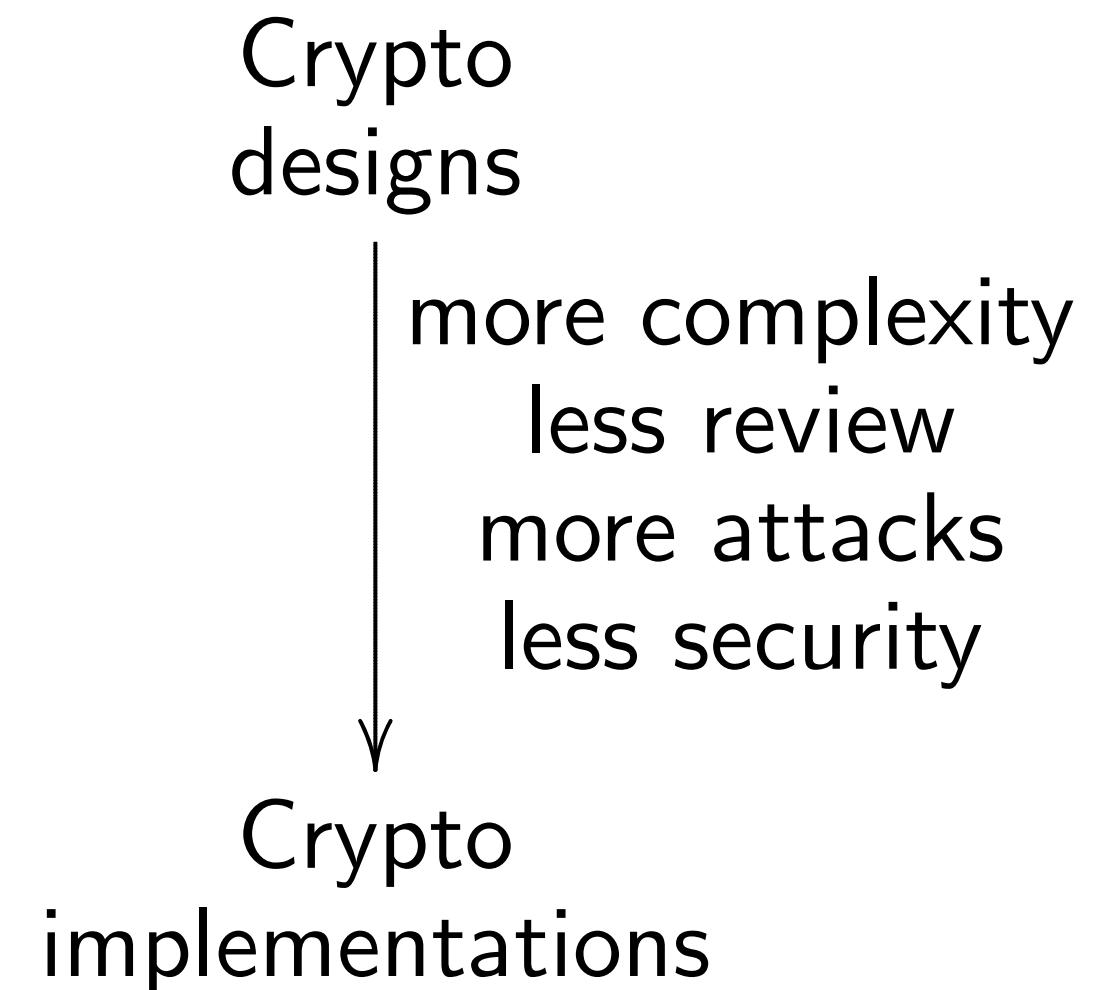


After many, many, many papers on implementations and attacks, today we still have an ecosystem plagued with AES vulnerabilities. Warning: more papers \neq security.

AES has a serious conflict between security, simplicity, speed. It's tough to achieve security while insisting on the AES design —i.e., blaming the implementor.

Allowing the *design* to vary makes security much easier. Next-generation ciphers are naturally constant-time and fast.

The big picture

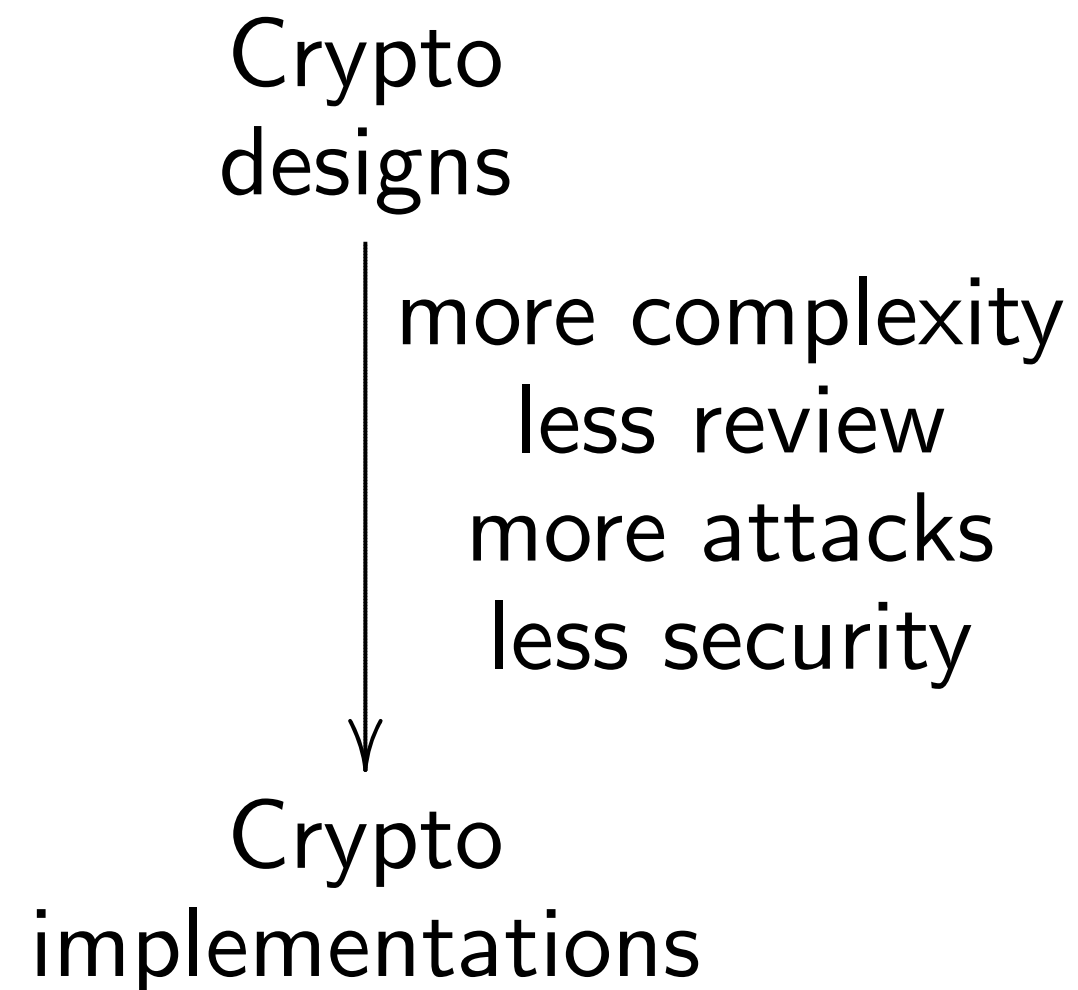


After many, many, many papers on implementations and attacks, today we still have an ecosystem plagued with AES vulnerabilities. Warning: more papers \neq security.

AES has a serious conflict between security, simplicity, speed. It's tough to achieve security while insisting on the AES design —i.e., blaming the implementor.

Allowing the *design* to vary makes security much easier. Next-generation ciphers are naturally constant-time and fast.

The big picture

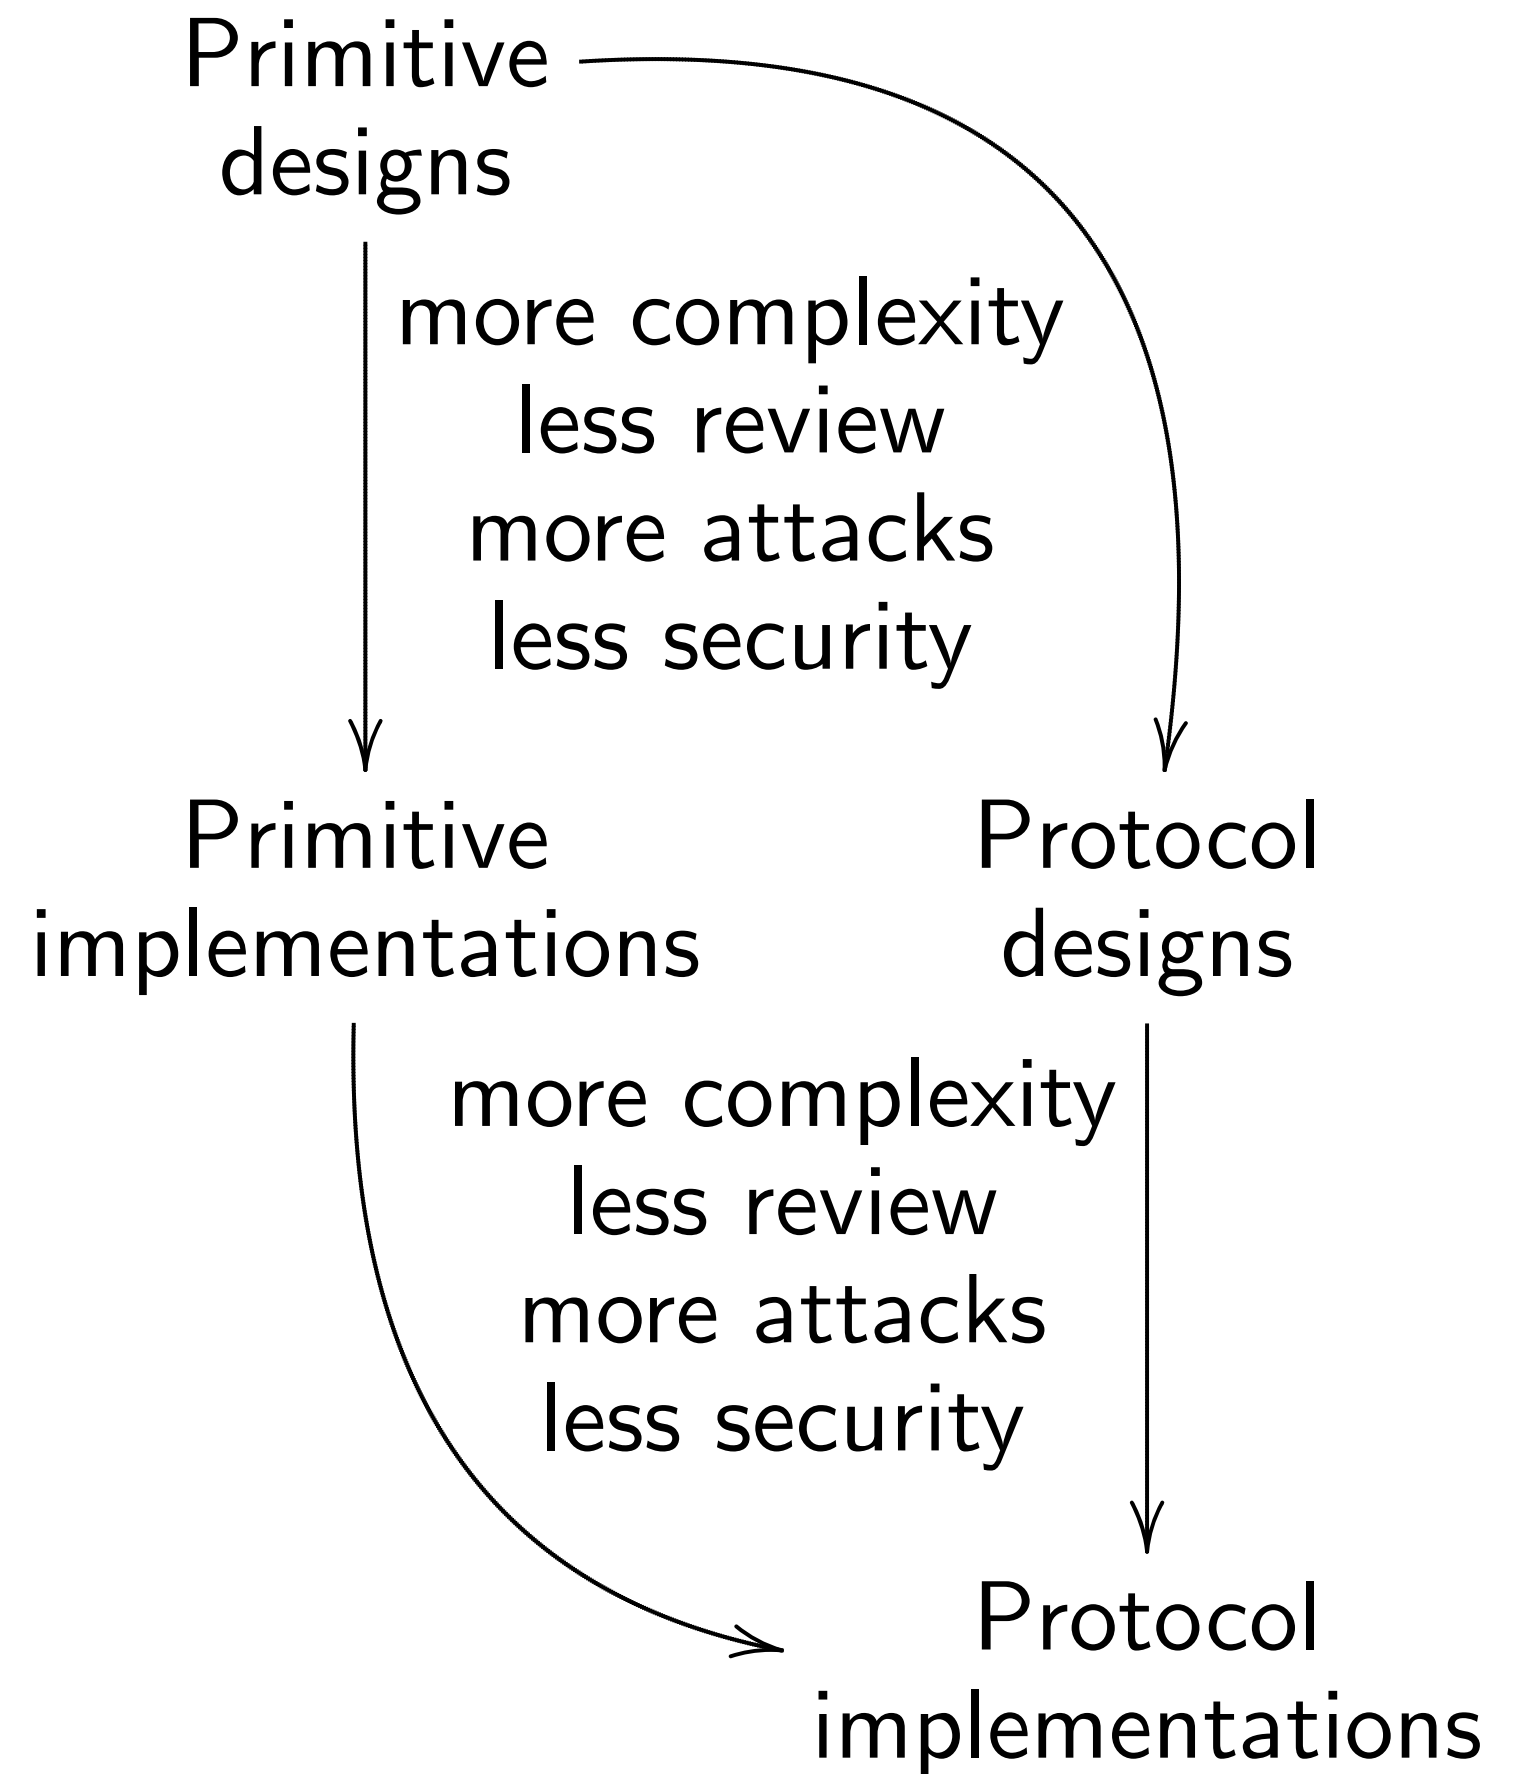


After many, many, many papers on implementations and attacks, today we still have an ecosystem plagued with AES vulnerabilities. Warning: more papers \neq security.

AES has a serious conflict between security, simplicity, speed. It's tough to achieve security while insisting on the AES design —i.e., blaming the implementor.

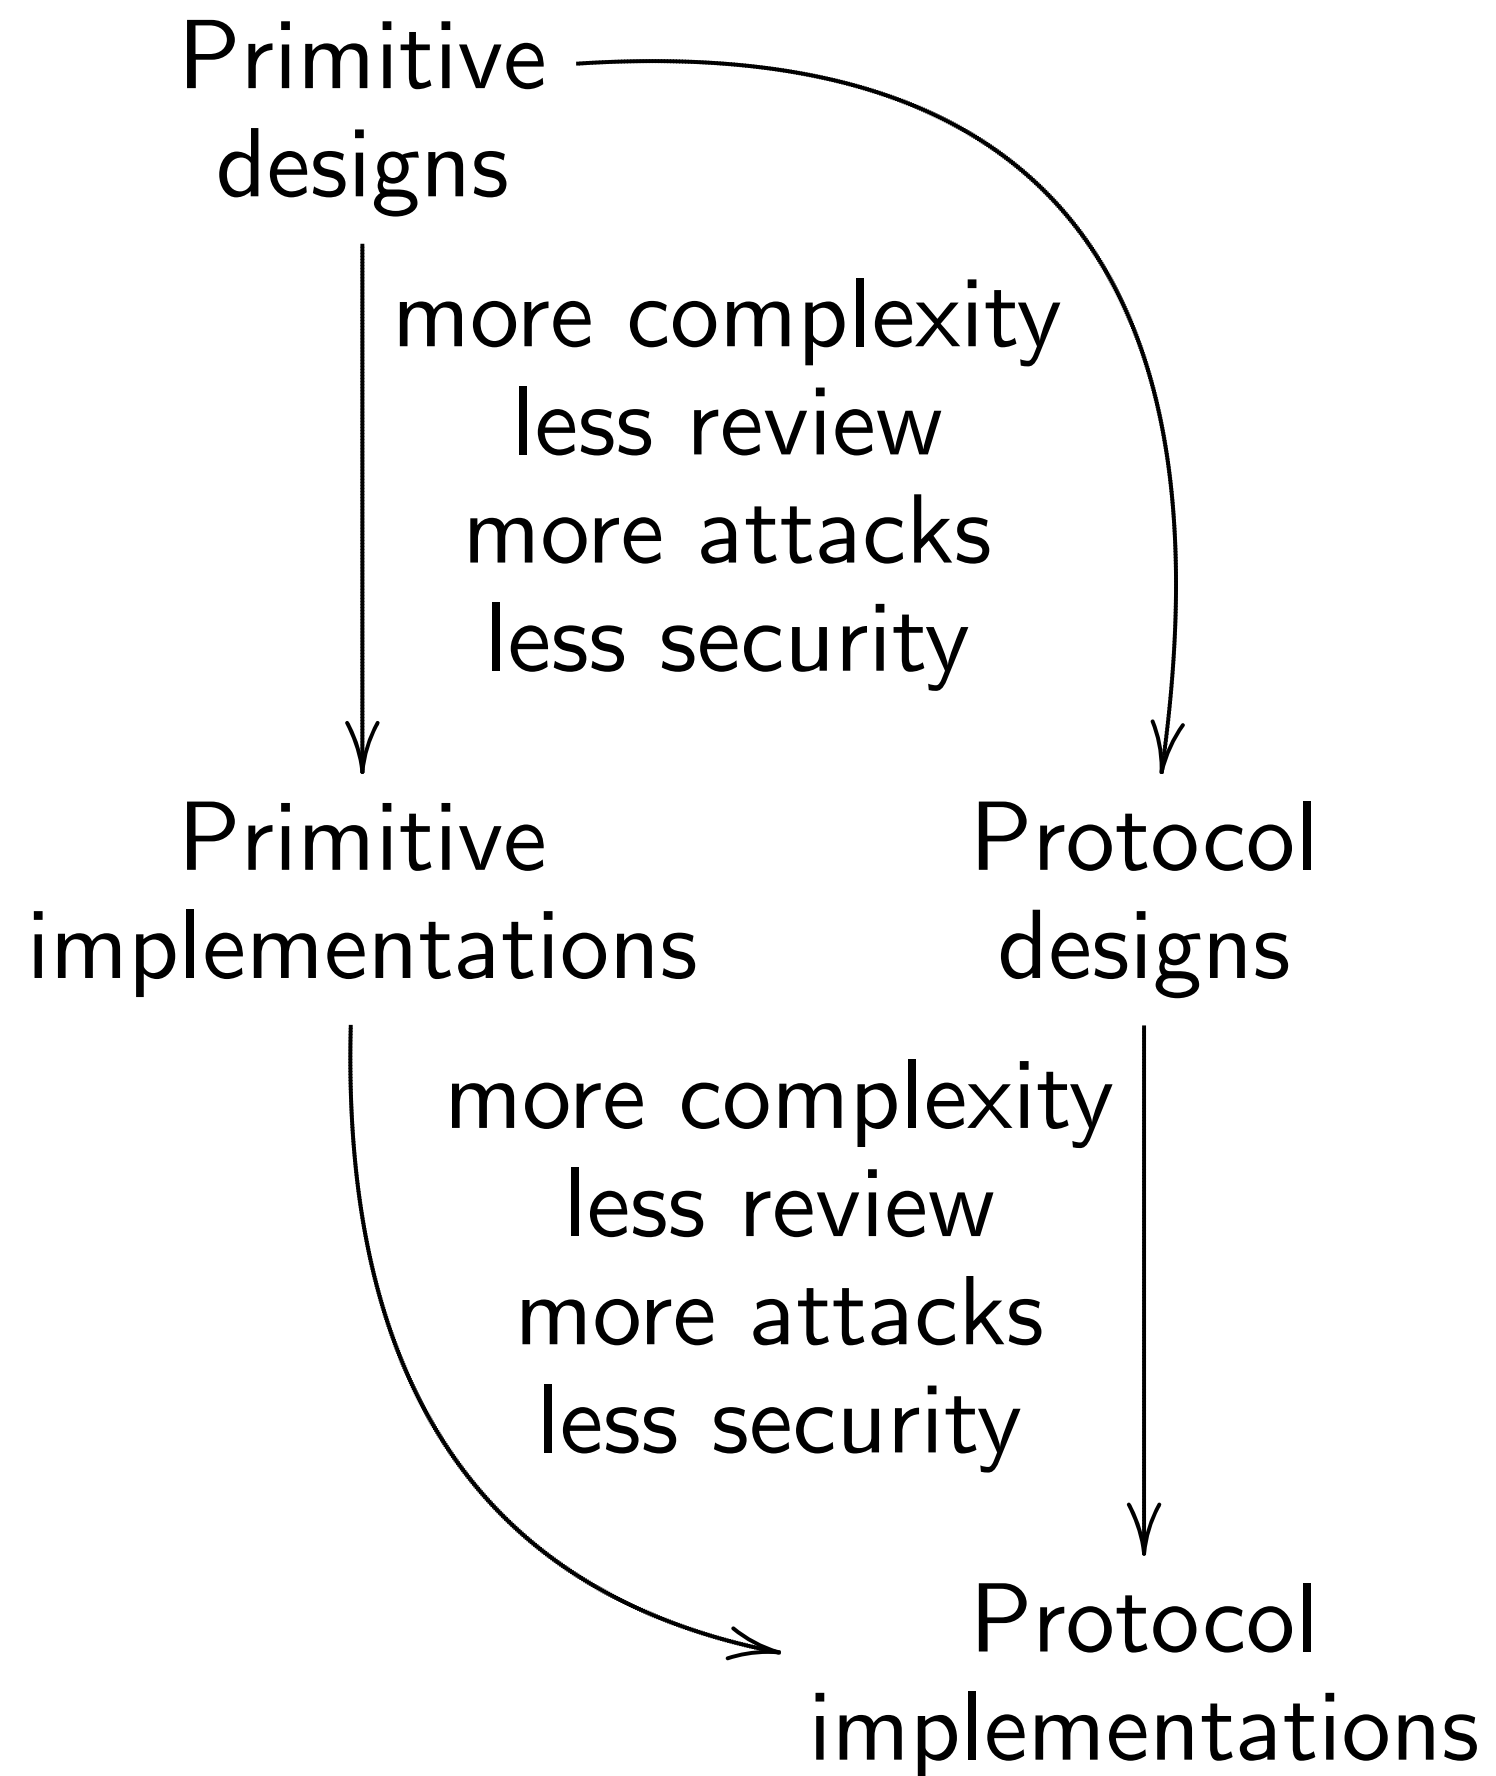
Allowing the *design* to vary makes security much easier. Next-generation ciphers are naturally constant-time and fast.

The big picture



any, many, many papers
implementations and attacks,
we still have an ecosystem
with AES vulnerabilities.
: more papers \neq security.
s a serious conflict
security, simplicity, speed.
gh to achieve security
sisting on the AES design
blaming the implementor.
g the *design* to vary
security much easier.
eneration ciphers are
y constant-time and fast.

The big picture



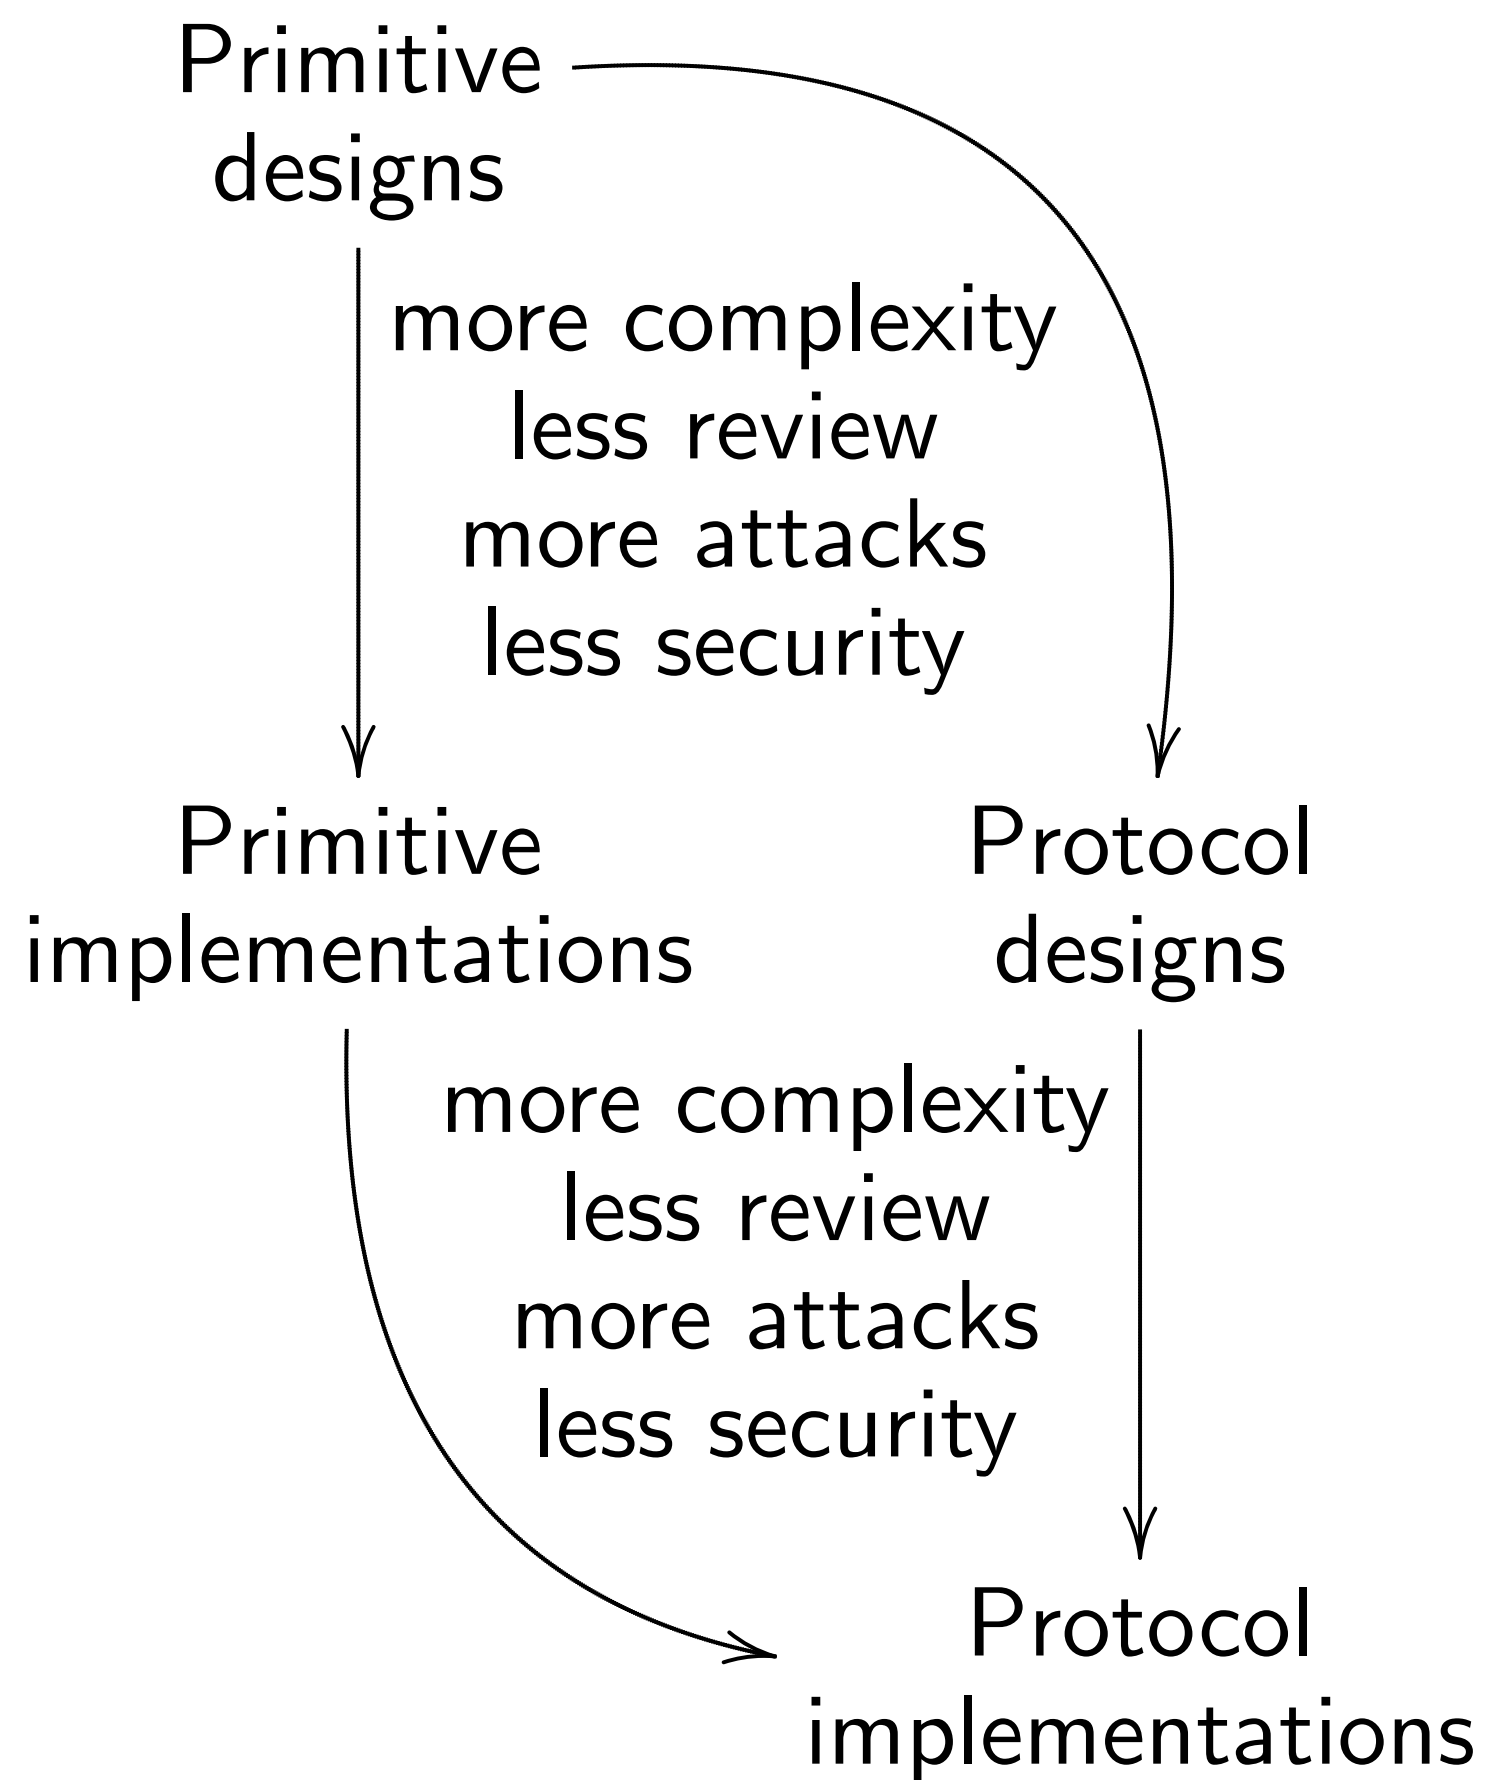
Public re
towards
ignoring
protocol
There's
of, e.g.,
than of

many papers
ns and attacks,
e an ecosystem
vulnerabilities.
pers \neq security.

conflict
simplicity, speed.
ve security
the AES design
e implementor.

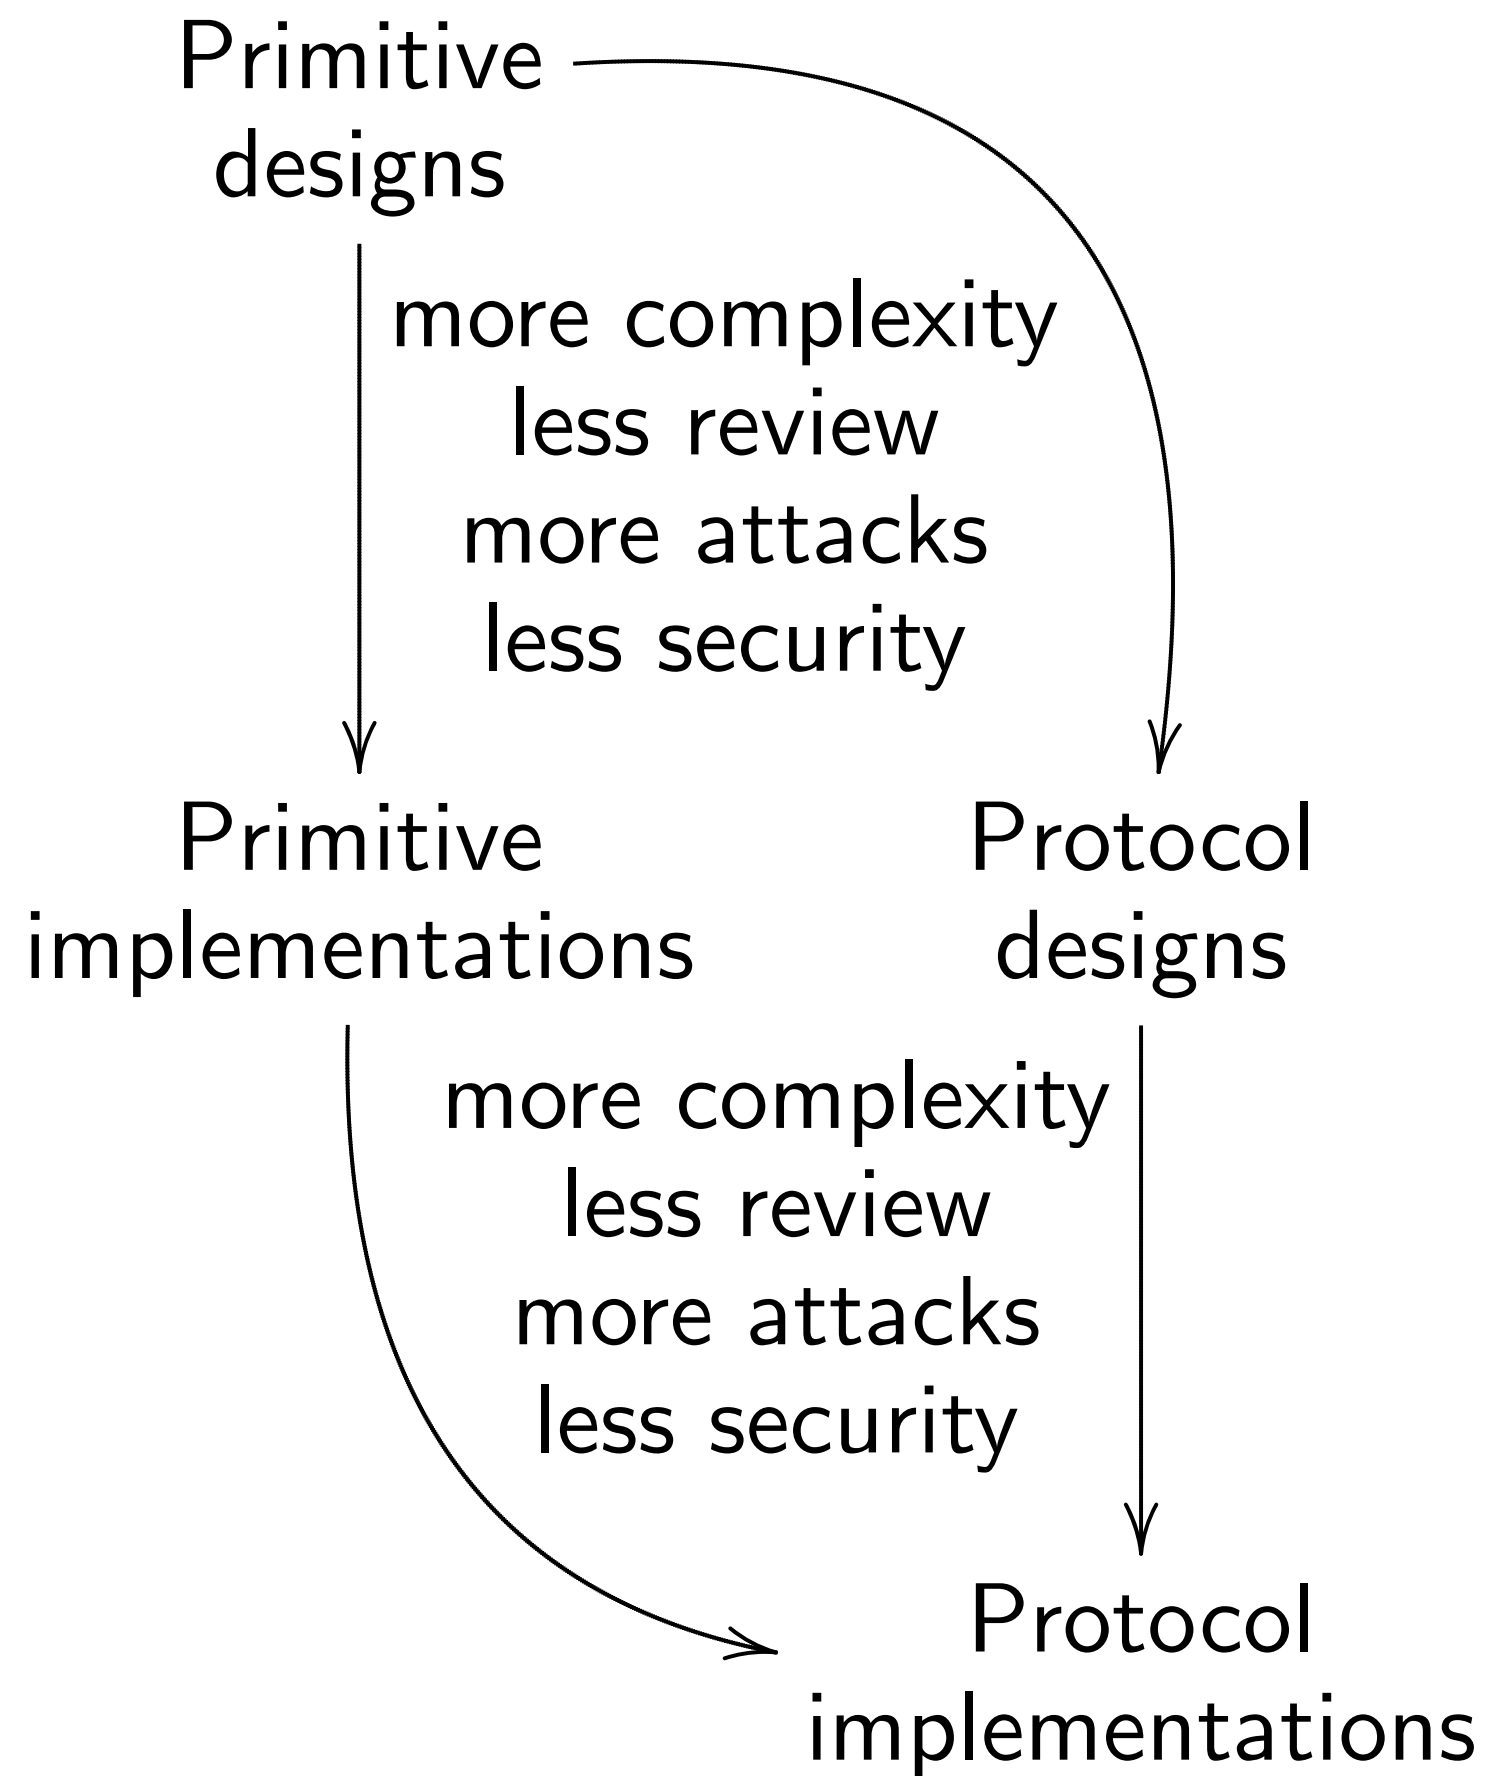
n to vary
ch easier.
phers are
-time and fast.

The big picture



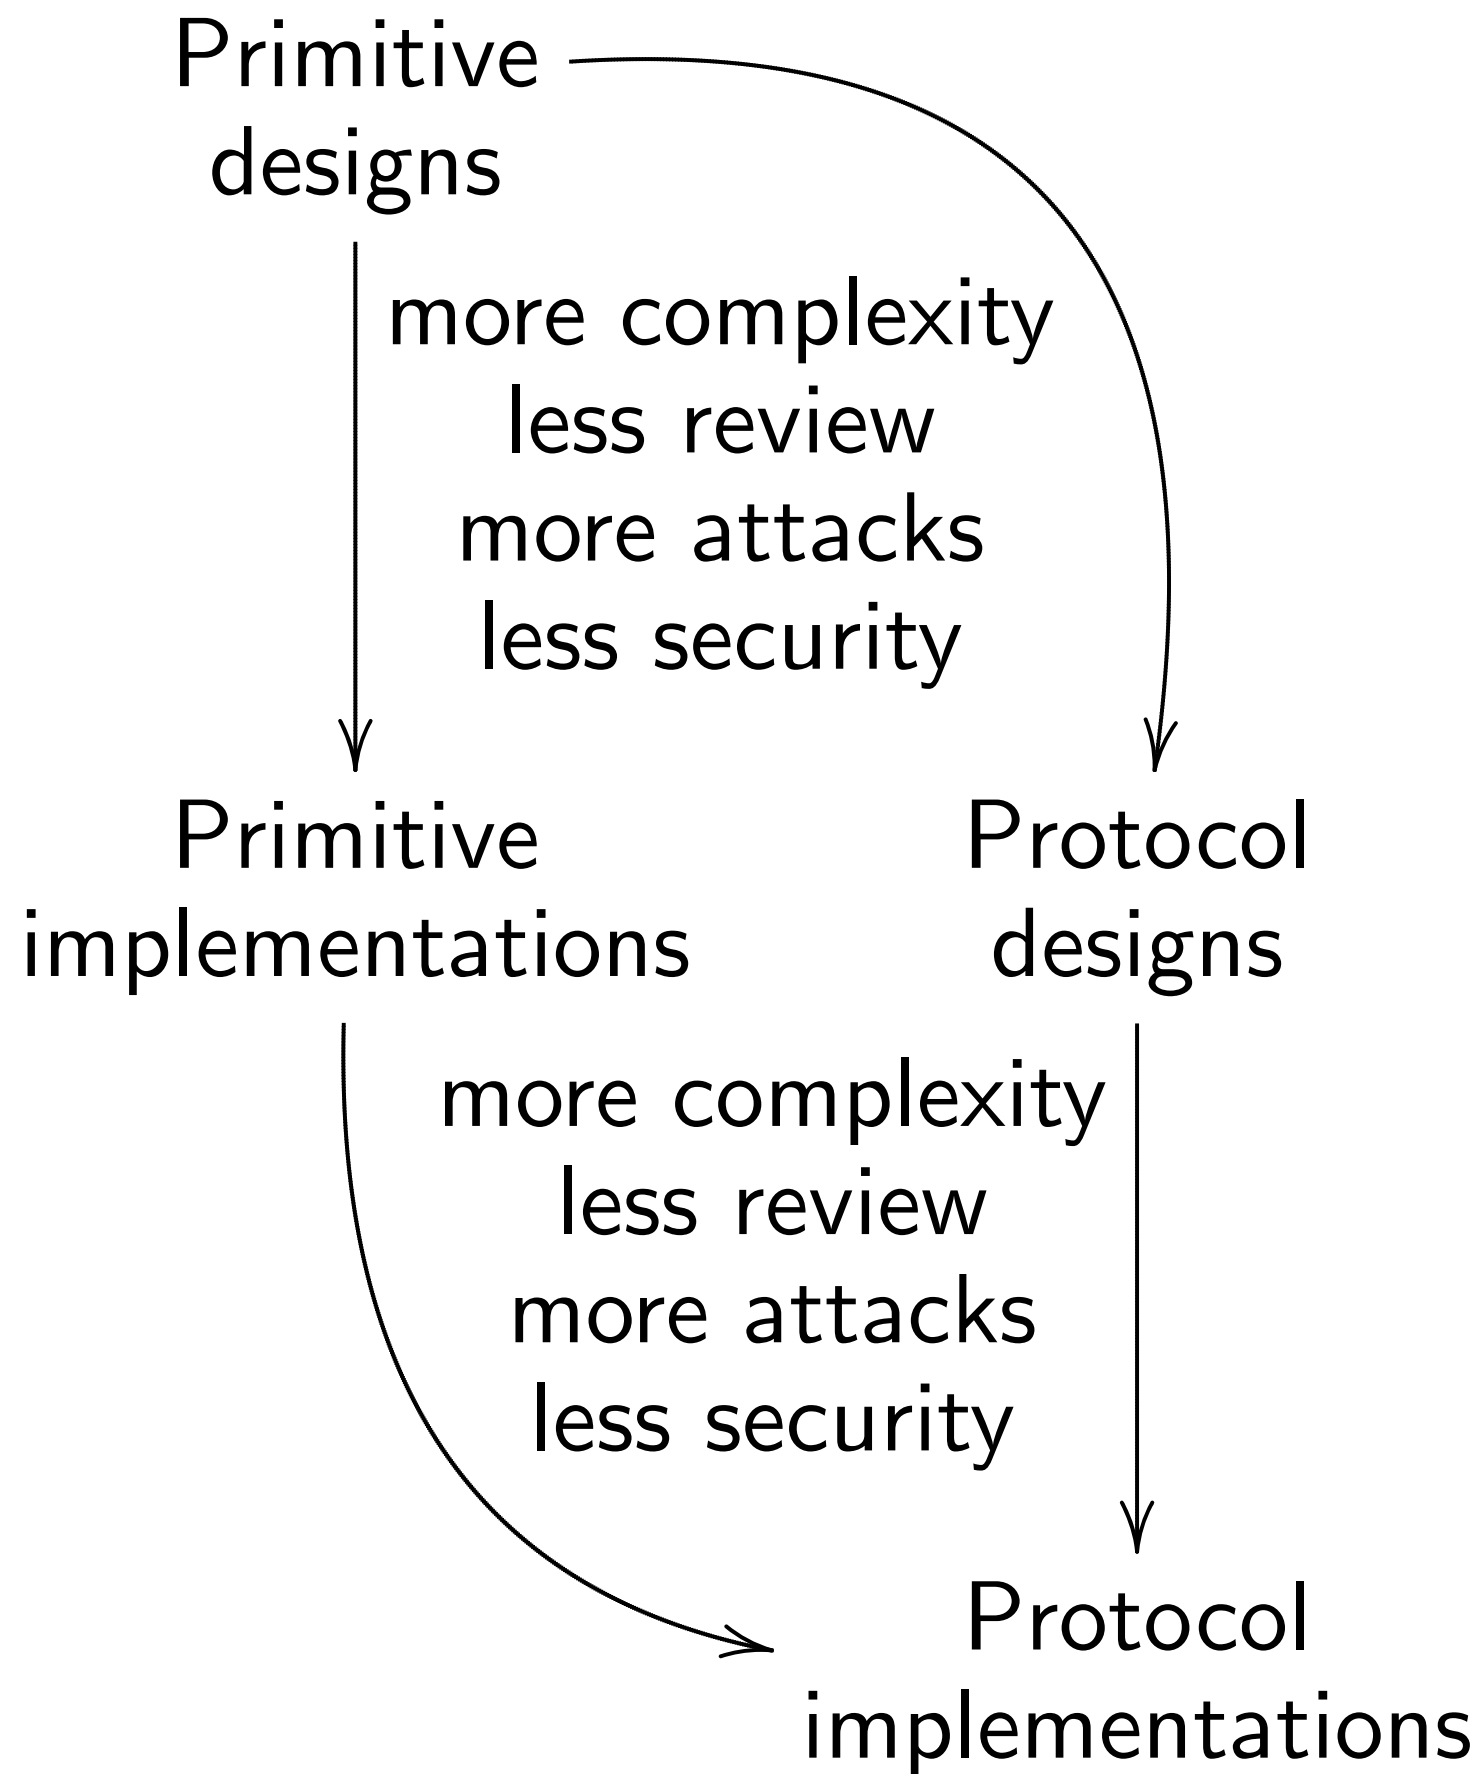
Public review is na
towards the simple
ignoring complicat
protocols and imp
There's much mor
of, e.g., discrete lo
than of ECDSA sig

The big picture



Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations. There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

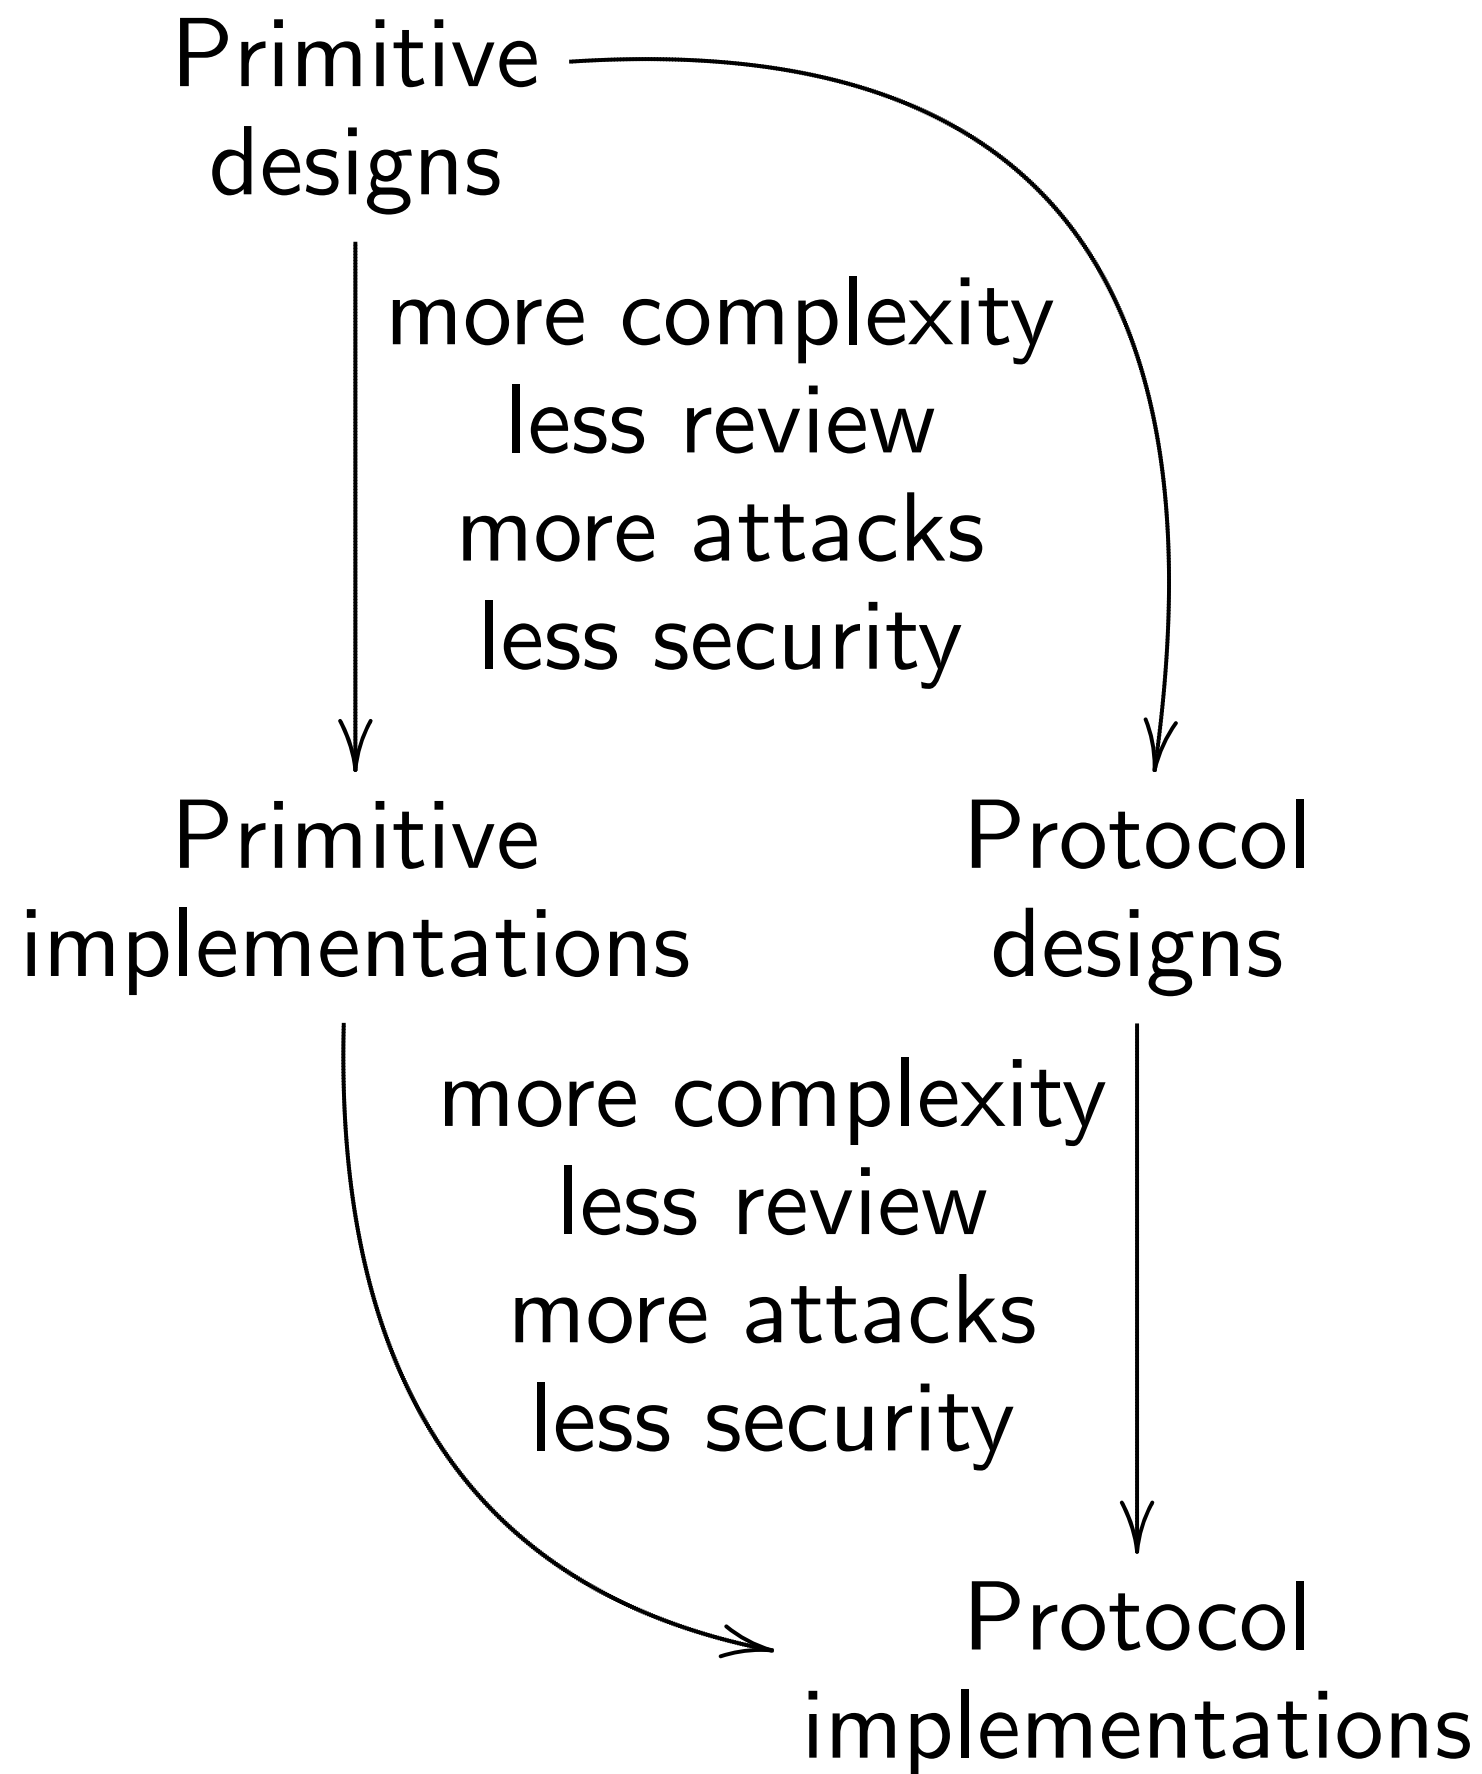
The big picture



Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

The big picture

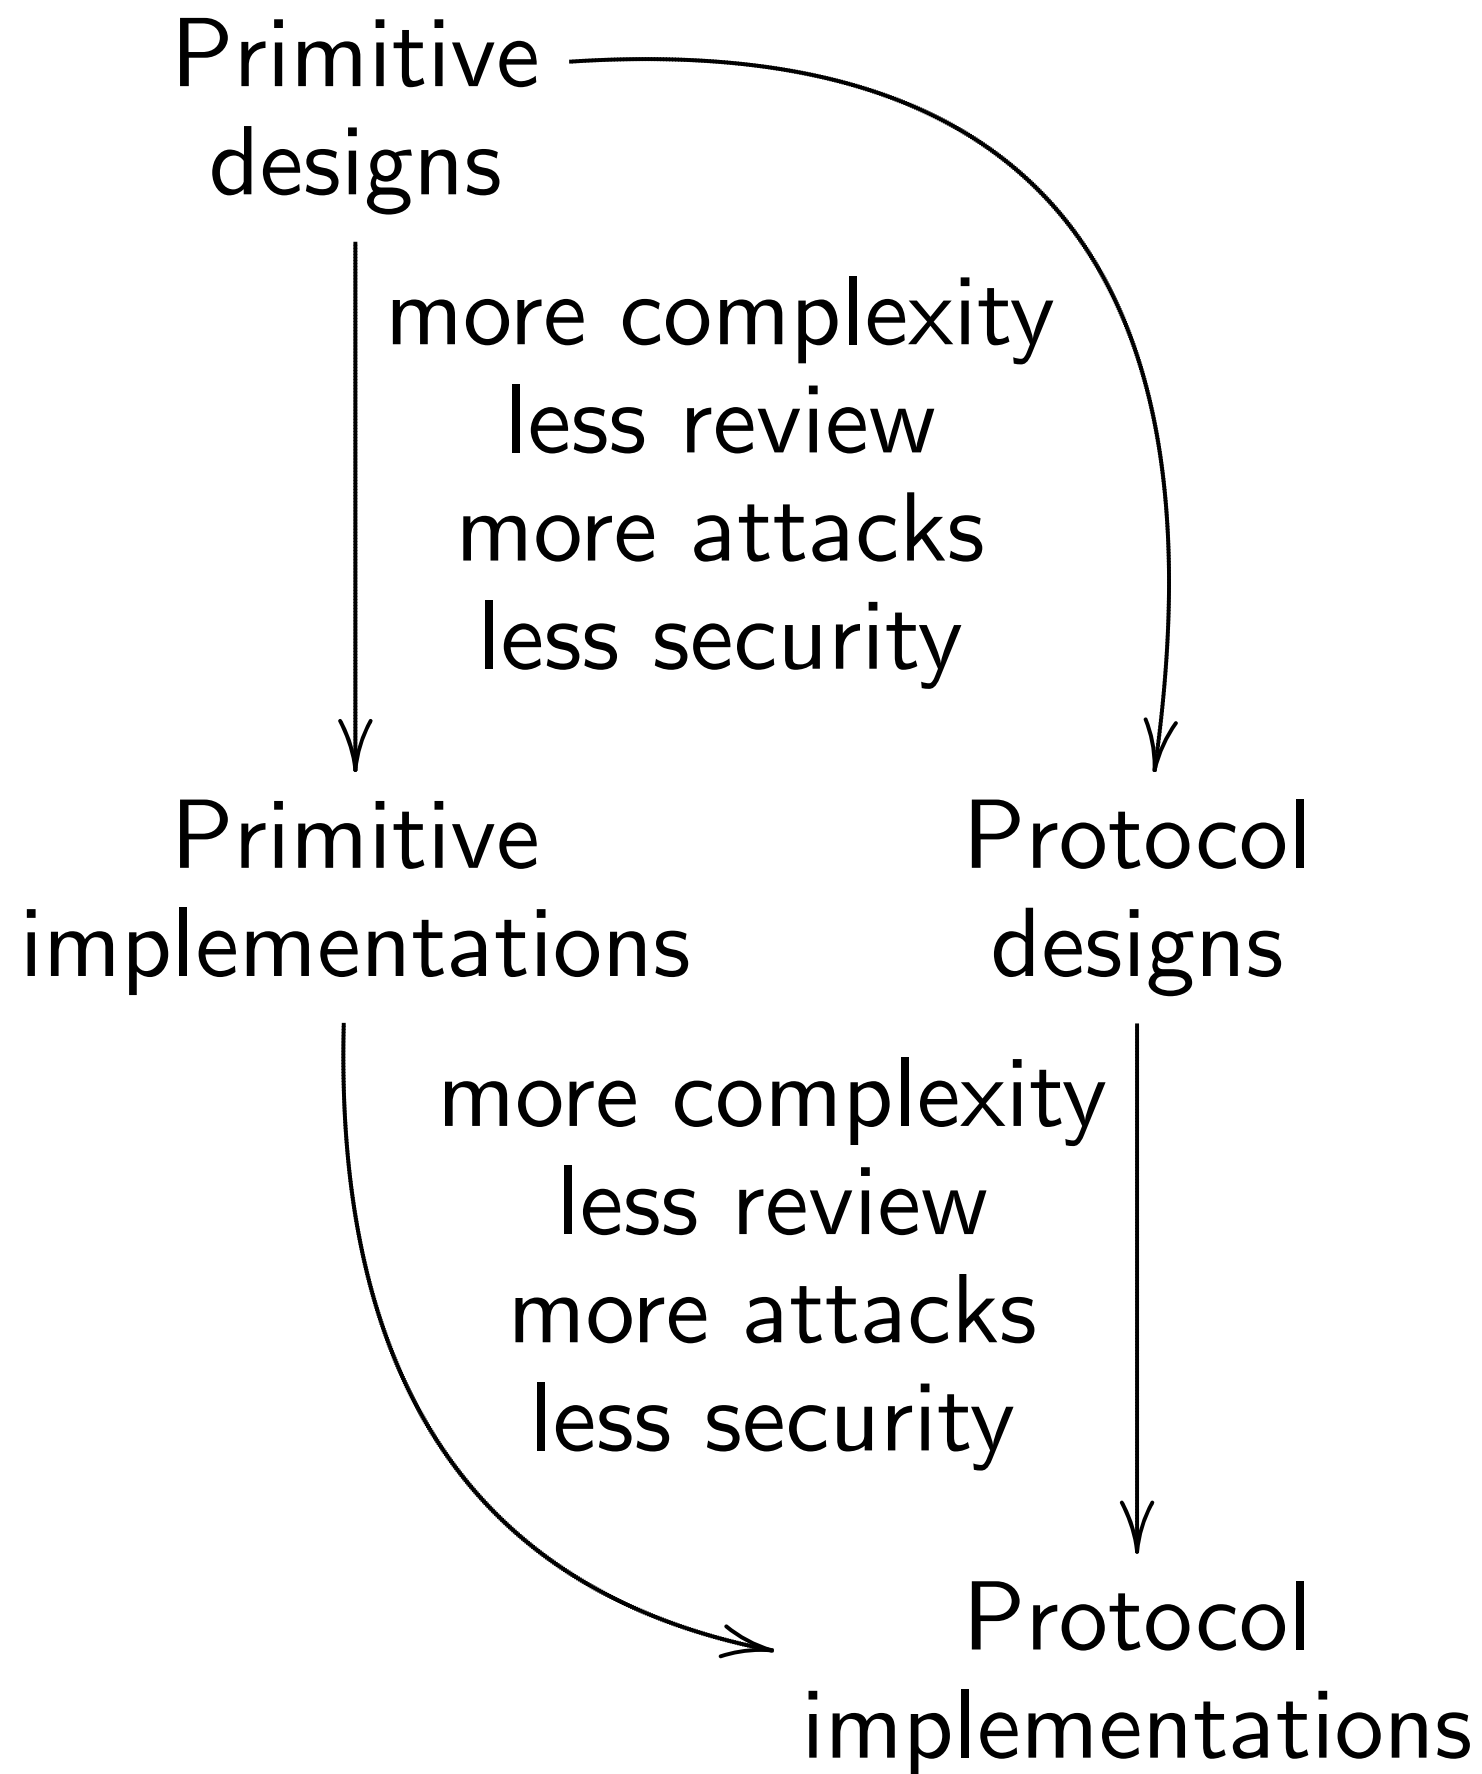


Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

There's much more public review of the ECDSA design than of ECDSA implementations.

The big picture



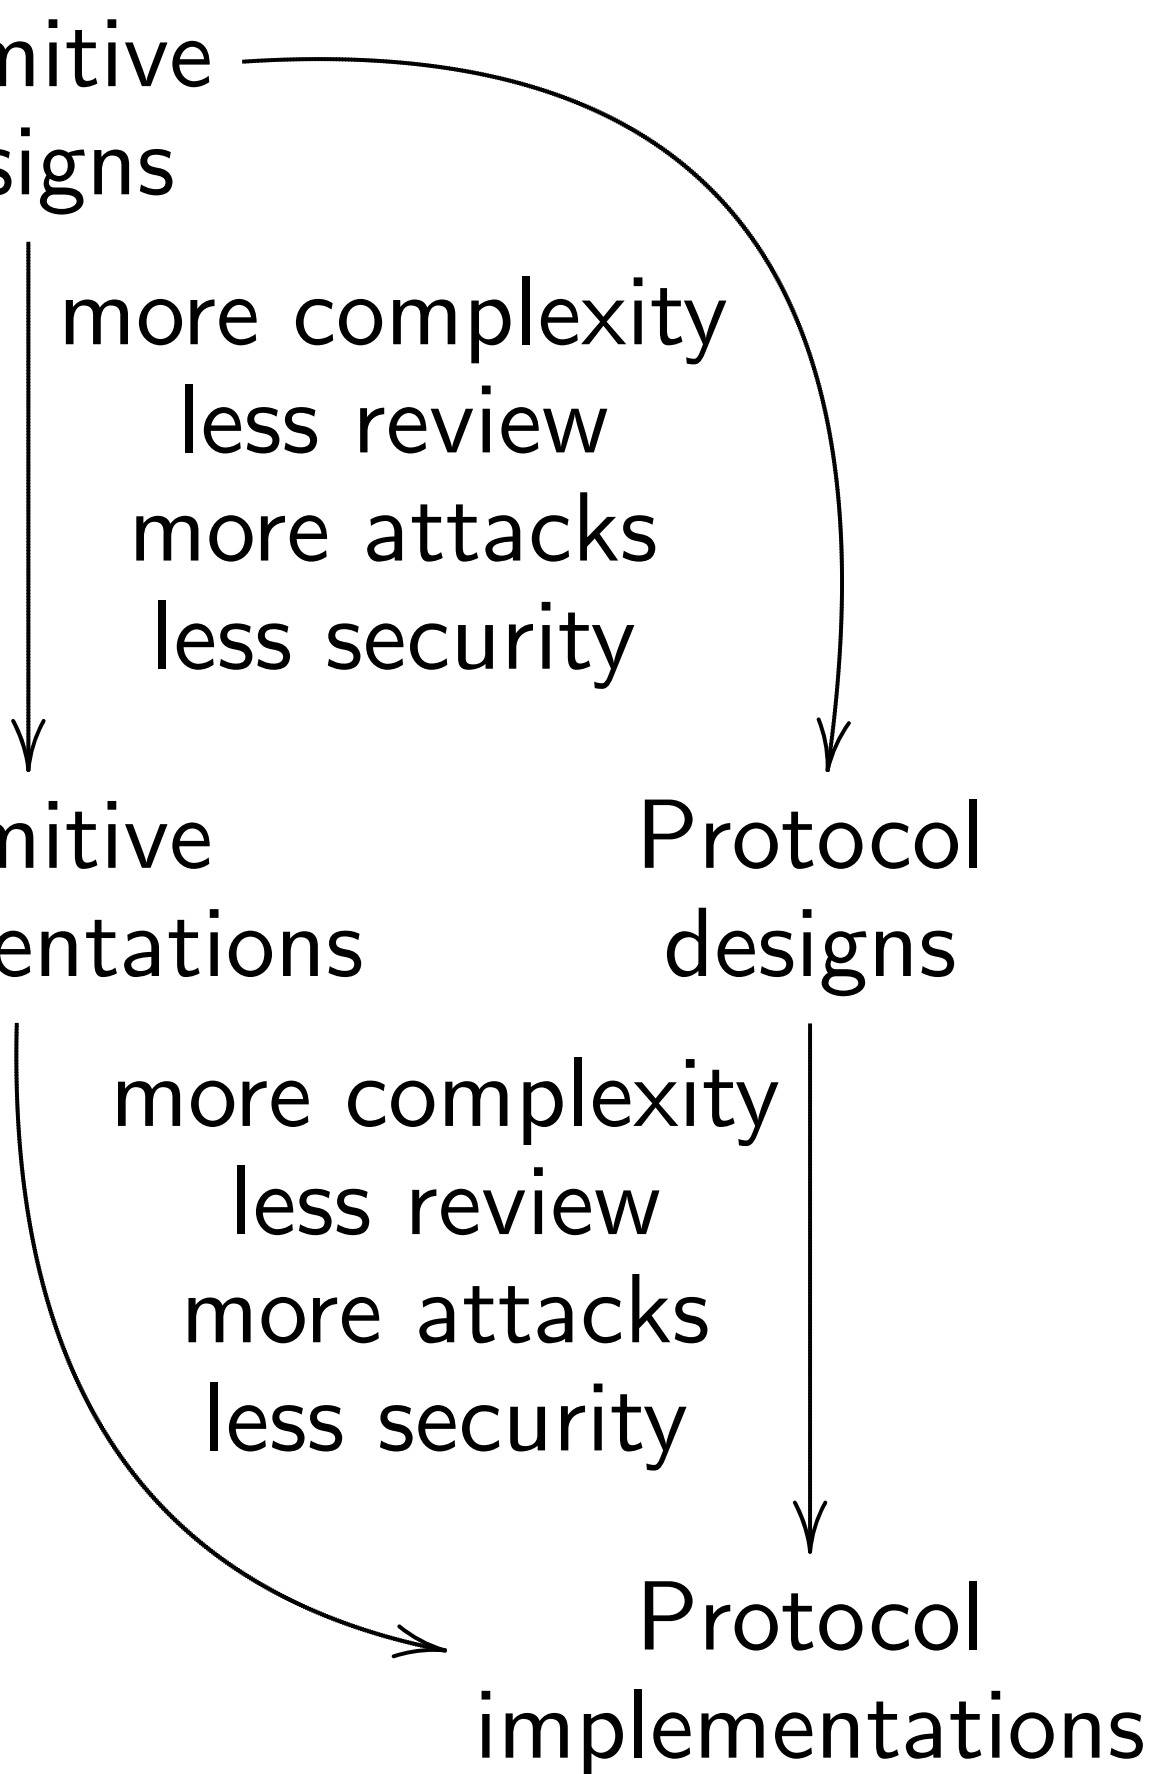
Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

There's much more public review of the ECDSA design than of ECDSA implementations.

There's much more public review of ECDSA implementations than of ECDSA applications.

picture



Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

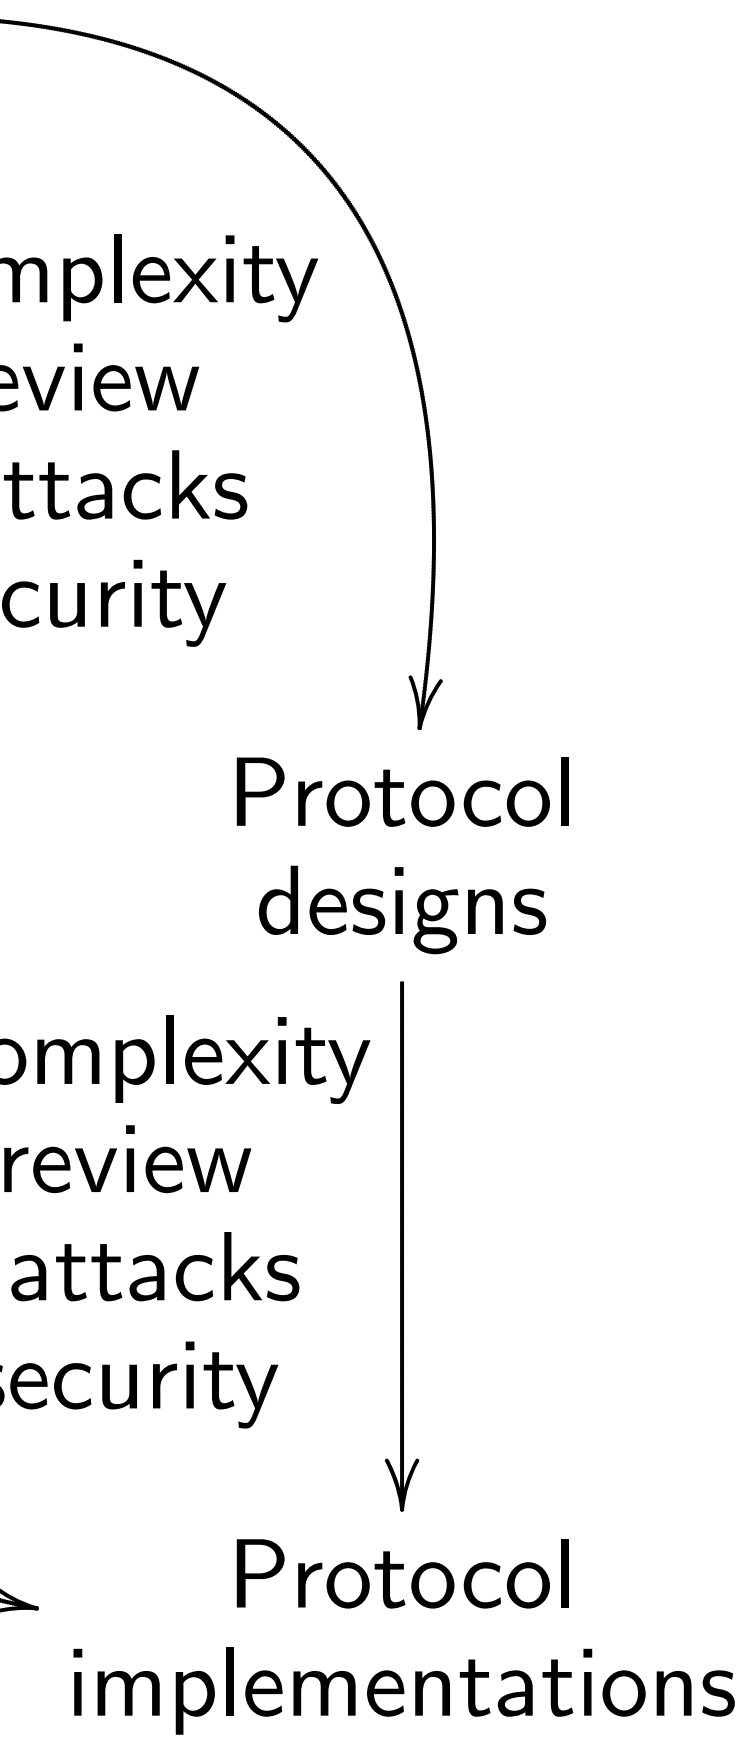
There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

There's much more public review of the ECDSA design than of ECDSA implementations.

There's much more public review of ECDSA implementations than of ECDSA applications.

What ab

The func
of "prov
Prove th
is as sec
i.e.: Pro
is as sec
Prove th
is as sec
Then it's
to focus



Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

There's much more public review of the ECDSA design than of ECDSA implementations.

There's much more public review of ECDSA implementations than of ECDSA applications.

What about security

The fundamental goal of "provable security" is to prove that the implementation is as secure as the design. i.e.: Prove that the implementation is as secure as the design. Prove that the implementation is as secure as the design. Then it's safe for us to focus on the practical

Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

There's much more public review of the ECDSA design than of ECDSA implementations.

There's much more public review of ECDSA implementations than of ECDSA applications.

What about security proofs?

The fundamental goal of “provable security”:
Prove that the whole system is as secure as the primitive.

i.e.: Prove that the protocol is as secure as the primitive.

Prove that the implementation is as secure as the design.

Then it's safe for reviewers to focus on the primitive design.

ocol
gns

ocol
ntations

Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

There's much more public review of the ECDSA design than of ECDSA implementations.

There's much more public review of ECDSA implementations than of ECDSA applications.

What about security proofs?

The fundamental goal of “provable security”:
Prove that the whole system is as secure as the primitive.

i.e.: Prove that the protocol is as secure as the primitive.

Prove that the implementation is as secure as the design.

Then it's safe for reviewers to focus on the primitive design.

Public review is naturally biased towards the simplest targets, ignoring complications of protocols and implementations.

There's much more public review of, e.g., discrete logarithms than of ECDSA signatures.

There's much more public review of the ECDSA design than of ECDSA implementations.

There's much more public review of ECDSA implementations than of ECDSA applications.

What about security proofs?

The fundamental goal of “provable security” :
Prove that the whole system is as secure as the primitive.

i.e.: Prove that the protocol is as secure as the primitive.

Prove that the implementation is as secure as the design.

Then it's safe for reviewers to focus on the primitive design.

Maybe will succeed someday, but needs to overcome huge problems.

review is naturally biased
the simplest targets,
complications of
s and implementations.

much more public review
discrete logarithms
ECDSA signatures.

much more public review
CDSA design
ECDSA implementations.

much more public review
SA implementations
ECDSA applications.

What about security proofs?

The fundamental goal
of “provable security” :
Prove that the whole system
is as secure as the primitive.

i.e.: Prove that the protocol
is as secure as the primitive.

Prove that the implementation
is as secure as the design.

Then it’s safe for reviewers
to focus on the primitive design.

Maybe will succeed someday, but
needs to overcome huge problems.

Problem
Proofs a
rarely re

naturally biased
est targets,
tions of
plementations.

re public review
gorithms
gnatures.

re public review
sign
plementations.

re public review
ementations
pplications.

What about security proofs?

The fundamental goal
of “provable security”:
Prove that the whole system
is as secure as the primitive.

i.e.: Prove that the protocol
is as secure as the primitive.

Prove that the implementation
is as secure as the design.

Then it’s safe for reviewers
to focus on the primitive design.

Maybe will succeed someday, but
needs to overcome huge problems.

Problem 1: “Proo
Proofs are increasi
rarely reviewed, ra

What about security proofs?

The fundamental goal
of “provable security”:

Prove that the whole system
is as secure as the primitive.

i.e.: Prove that the protocol
is as secure as the primitive.

Prove that the implementation
is as secure as the design.

Then it’s safe for reviewers
to focus on the primitive design.

Maybe will succeed someday, but
needs to overcome huge problems.

Problem 1: “Proofs” have e

Proofs are increasingly comp

rarely reviewed, rarely autom

What about security proofs?

The fundamental goal of “provable security”:

Prove that the whole system is as secure as the primitive.

i.e.: Prove that the protocol is as secure as the primitive.

Prove that the implementation is as secure as the design.

Then it’s safe for reviewers to focus on the primitive design.

Maybe will succeed someday, but needs to overcome huge problems.

Problem 1: “Proofs” have errors. Proofs are increasingly complex, rarely reviewed, rarely automated.

What about security proofs?

The fundamental goal of “provable security”:

Prove that the whole system is as secure as the primitive.

i.e.: Prove that the protocol is as secure as the primitive.

Prove that the implementation is as secure as the design.

Then it's safe for reviewers to focus on the primitive design.

Maybe will succeed someday, but needs to overcome huge problems.

Problem 1: “Proofs” have errors. Proofs are increasingly complex, rarely reviewed, rarely automated.

Problem 2: Most proofs are of security bounds that aren't tight: e.g. forking-lemma “security” is pure deception for typical sizes.

What about security proofs?

The fundamental goal of “provable security”:

Prove that the whole system is as secure as the primitive.

i.e.: Prove that the protocol is as secure as the primitive.

Prove that the implementation is as secure as the design.

Then it’s safe for reviewers to focus on the primitive design.

Maybe will succeed someday, but needs to overcome huge problems.

Problem 1: “Proofs” have errors. Proofs are increasingly complex, rarely reviewed, rarely automated.

Problem 2: Most proofs are of security bounds that aren’t tight: e.g. forking-lemma “security” is pure deception for typical sizes.

Problem 3: “Security” definitions prioritize simplicity over accuracy. e.g. is MAC-pad-encrypt secure?

What about security proofs?

The fundamental goal of “provable security”:

Prove that the whole system is as secure as the primitive.

i.e.: Prove that the protocol is as secure as the primitive.

Prove that the implementation is as secure as the design.

Then it’s safe for reviewers to focus on the primitive design.

Maybe will succeed someday, but needs to overcome huge problems.

Problem 1: “Proofs” have errors. Proofs are increasingly complex, rarely reviewed, rarely automated.

Problem 2: Most proofs are of security bounds that aren’t tight: e.g. forking-lemma “security” is pure deception for typical sizes.

Problem 3: “Security” definitions prioritize simplicity over accuracy. e.g. is MAC-pad-encrypt secure?

Problem 4: Maybe the only way to achieve the fundamental goal is to switch to weak primitives.

About security proofs?

fundamental goal

"achievable security":

that the whole system

is secure as the primitive.

to prove that the protocol

is secure as the primitive.

that the implementation

is secure as the design.

is safe for reviewers

to rely on the primitive design.

will succeed someday, but

to overcome huge problems.

Problem 1: "Proofs" have errors.

Proofs are increasingly complex,
rarely reviewed, rarely automated.

Problem 2: Most proofs are of
security bounds that aren't tight:
e.g. forking-lemma "security"
is pure deception for typical sizes.

Problem 3: "Security" definitions
prioritize simplicity over accuracy.
e.g. is MAC-pad-encrypt secure?

Problem 4: Maybe the only way
to achieve the fundamental goal
is to switch to weak primitives.

Some advice

Creating

Think about

implementing

What will

What errors

to appear

Can you

Security proofs?

goal

“Security”:

role system

primitive.

the protocol

primitive.

implementation

design.

reviewers

primitive design.

and someday, but

the huge problems.

Problem 1: “Proofs” have errors.

Proofs are increasingly complex,
rarely reviewed, rarely automated.

Problem 2: Most proofs are of
security bounds that aren't tight:
e.g. forking-lemma “security”
is pure deception for typical sizes.

Problem 3: “Security” definitions
prioritize simplicity over accuracy.
e.g. is MAC-pad-encrypt secure?

Problem 4: Maybe the only way
to achieve the fundamental goal
is to switch to weak primitives.

Some advice to create

Creating or evaluating

**Think about the
implementations**

What will the implementation

What errors are likely

to appear in implementations

Can you compensate

Problem 1: “Proofs” have errors.
Proofs are increasingly complex,
rarely reviewed, rarely automated.

Problem 2: Most proofs are of
security bounds that aren’t tight:
e.g. forking-lemma “security”
is pure deception for typical sizes.

Problem 3: “Security” definitions
prioritize simplicity over accuracy.
e.g. is MAC-pad-encrypt secure?

Problem 4: Maybe the only way
to achieve the fundamental goal
is to switch to weak primitives.

Some advice to crypto design

Creating or evaluating a design

**Think about the
implementations.**

What will the implementors

What errors are likely

to appear in implementation

Can you compensate for this

Problem 1: “Proofs” have errors.
Proofs are increasingly complex,
rarely reviewed, rarely automated.

Problem 2: Most proofs are of
security bounds that aren’t tight:
e.g. forking-lemma “security”
is pure deception for typical sizes.

Problem 3: “Security” definitions
prioritize simplicity over accuracy.
e.g. is MAC-pad-encrypt secure?

Problem 4: Maybe the only way
to achieve the fundamental goal
is to switch to weak primitives.

Some advice to crypto designers

Creating or evaluating a design?

**Think about the
implementations.**

What will the implementors do?

What errors are likely
to appear in implementations?

Can you compensate for this?

Problem 1: “Proofs” have errors.
Proofs are increasingly complex,
rarely reviewed, rarely automated.

Problem 2: Most proofs are of
security bounds that aren’t tight:
e.g. forking-lemma “security”
is pure deception for typical sizes.

Problem 3: “Security” definitions
prioritize simplicity over accuracy.
e.g. is MAC-pad-encrypt secure?

Problem 4: Maybe the only way
to achieve the fundamental goal
is to switch to weak primitives.

Some advice to crypto designers

Creating or evaluating a design?

**Think about the
implementations.**

What will the implementors do?

What errors are likely
to appear in implementations?

Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

**Think about the
higher-level protocols.**

Will the system be secure?

1: “Proofs” have errors.
are increasingly complex,
viewed, rarely automated.

2: Most proofs are of
bounds that aren't tight:
ing-lemma “security”
deception for typical sizes.

3: “Security” definitions
simplicity over accuracy.
MAC-pad-encrypt secure?

4: Maybe the only way
ve the fundamental goal
tch to weak primitives.

Some advice to crypto designers

Creating or evaluating a design?

**Think about the
implementations.**

What will the implementors do?

What errors are likely

to appear in implementations?

Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

**Think about the
higher-level protocols.**

Will the system be secure?

Crypto h

HTTPS.

fs” have errors.
ingly complex,
rely automated.

proofs are of
at aren't tight:
a “security”
for typical sizes.

“security” definitions
y over accuracy.
ncrypt secure?

e the only way
damental goal
ak primitives.

Some advice to crypto designers

Creating or evaluating a design?

**Think about the
implementations.**

What will the implementors do?

What errors are likely
to appear in implementations?

Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

**Think about the
higher-level protocols.**

Will the system be secure?

Crypto horror stories

HTTPS.

errors.
plex,
nated.
of
tight:
"
sizes.
itions
uracy.
ure?
way
goal
es.

Some advice to crypto designers

Creating or evaluating a design?

**Think about the
implementations.**

What will the implementors do?

What errors are likely
to appear in implementations?
Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

**Think about the
higher-level protocols.**

Will the system be secure?

Crypto horror story #3

HTTPS.

Some advice to crypto designers

Creating or evaluating a design?

Think about the implementations.

What will the implementors do?

What errors are likely to appear in implementations?

Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

Think about the higher-level protocols.

Will the system be secure?

Crypto horror story #3

HTTPS.

Some advice to crypto designers

Creating or evaluating a design?

Think about the implementations.

What will the implementors do?

What errors are likely to appear in implementations?

Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

Think about the higher-level protocols.

Will the system be secure?

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge improvement in HTTPS usability; will obviously be widely used.

But is HTTPS actually secure?

Some advice to crypto designers

Creating or evaluating a design?

Think about the implementations.

What will the implementors do?

What errors are likely to appear in implementations?

Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

Think about the higher-level protocols.

Will the system be secure?

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge improvement in HTTPS usability; will obviously be widely used.

But is HTTPS actually secure?

“It’s not so bad against passive eavesdroppers!”
—The eavesdroppers will start forging (more) packets.

Some advice to crypto designers

Creating or evaluating a design?

Think about the implementations.

What will the implementors do?

What errors are likely to appear in implementations?

Can you compensate for this?

Is the design a primitive?

Think about the protocols.

Is the design a protocol?

Think about the higher-level protocols.

Will the system be secure?

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge improvement in HTTPS usability; will obviously be widely used.

But is HTTPS actually secure?

“It’s not so bad against passive eavesdroppers!”
—The eavesdroppers will start forging (more) packets.

“Then we’ll know they’re there!”
—Yes, we knew that already.

What we want is *security*.

Advice to crypto designers

g or evaluating a design?

about the

implementations.

Will the implementors do?

Errors are likely

to occur in implementations?

Can they compensate for this?

Can you design a primitive?

about the protocols.

Can you design a protocol?

about the

low-level protocols.

Can the system be secure?

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge
improvement in HTTPS usability;
it will obviously be widely used.

But is HTTPS actually secure?

“It’s not so bad
against passive eavesdroppers!”

—The eavesdroppers will start
forging (more) packets.

“Then we’ll know they’re there!”

—Yes, we knew that already.

What we want is *security*.

2013.01

State-of-

are obvious

e.g. Defe

requires

assumpt

Crypto designers

ating a design?

lementors do?

kely

mentations?

ate for this?

imitive?

protocols.

otocol?

ocols.

e secure?

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge
improvement in HTTPS usability;
will obviously be widely used.

But is HTTPS actually secure?

“It’s not so bad
against passive eavesdroppers!”
—The eavesdroppers will start
forging (more) packets.

“Then we’ll know they’re there!”
—Yes, we knew that already.

What we want is *security*.

2013.01 Green:

State-of-the-art T
are obviously unsa

e.g. Defense vs. B
requires “goofy m
assumption” for p

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge improvement in HTTPS usability; will obviously be widely used.

But is HTTPS actually secure?

“It’s not so bad against passive eavesdroppers!”
—The eavesdroppers will start forging (more) packets.

“Then we’ll know they’re there!”

—Yes, we knew that already.

What we want is *security*.

2013.01 Green:

State-of-the-art TLS proofs are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher requires “goofy made-up assumption” for proof.

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge improvement in HTTPS usability; will obviously be widely used. But is HTTPS actually secure?

“It’s not so bad against passive eavesdroppers!”
—The eavesdroppers will start forging (more) packets.

“Then we’ll know they’re there!”
—Yes, we knew that already.
What we want is *security*.

2013.01 Green:

State-of-the-art TLS proofs are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher requires “goofy made-up assumption” for proof.

Crypto horror story #3

HTTPS.

letsencrypt.org is a huge improvement in HTTPS usability; will obviously be widely used. But is HTTPS actually secure?

“It’s not so bad against passive eavesdroppers!”
—The eavesdroppers will start forging (more) packets.

“Then we’ll know they’re there!”
—Yes, we knew that already. What we want is *security*.

2013.01 Green:

State-of-the-art TLS proofs are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher requires “goofy made-up assumption” for proof.

But that was “the *good news* . . . The problem with TLS is that we are cursed with *implementations*.”

e.g. Defense vs. Bleichenbacher is in wrong order in OpenSSL. Does this allow timing attacks?

horror story #3

crypt.org is a huge
ment in HTTPS usability;
ously be widely used.
HTTPS actually secure?

so bad
passive eavesdroppers!”
eavesdroppers will start
(more) packets.

ve’ll know they’re there!”
ve knew that already.
e want is *security*.

2013.01 Green:

State-of-the-art TLS proofs
are obviously unsatisfactory.
e.g. Defense vs. Bleichenbacher
requires “goofy made-up
assumption” for proof.

But that was “the *good
news* . . . The problem with
TLS is that we are cursed with
implementations.”

e.g. Defense vs. Bleichenbacher
is in wrong order in OpenSSL.
Does this allow timing attacks?

2014.08
Weiss–S
Successf
exploitin
variation
NITROX

y #3

g is a huge
HTTPS usability;
widely used.
actually secure?

vesdroppers!”
ers will start
ckets.

they’re there!”
hat already.
security.

2013.01 Green:

State-of-the-art TLS proofs
are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher
requires “goofy made-up
assumption” for proof.

But that was “the *good
news* . . . The problem with
TLS is that we are cursed with
implementations.”

e.g. Defense vs. Bleichenbacher
is in wrong order in OpenSSL.
Does this allow timing attacks?

2014.08 Meyer–So
Weiss–Schwenk–S
Successful Bleiche
exploiting analogo
variations in Java
NITROX SSL acce

2013.01 Green:

State-of-the-art TLS proofs
are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher
requires “goofy made-up
assumption” for proof.

But that was “the *good
news* . . . The problem with
TLS is that we are cursed with
implementations.”

e.g. Defense vs. Bleichenbacher
is in wrong order in OpenSSL.
Does this allow timing attacks?

2014.08 Meyer–Somorovsky–
Weiss–Schwenk–Schinzel–Te
Successful Bleichenbacher at
exploiting analogous timing
variations in Java SSE, Cavi
NITROX SSL accelerator ch

2013.01 Green:

State-of-the-art TLS proofs are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher requires “goofy made-up assumption” for proof.

But that was “the *good news* . . . The problem with TLS is that we are cursed with *implementations*.”

e.g. Defense vs. Bleichenbacher is in wrong order in OpenSSL.
Does this allow timing attacks?

2014.08 Meyer–Somorovsky–Weiss–Schwenk–Schinzel–Tews: Successful Bleichenbacher attacks, exploiting analogous timing variations in Java SSE, Cavium NITROX SSL accelerator chip.

2013.01 Green:

State-of-the-art TLS proofs are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher requires “goofy made-up assumption” for proof.

But that was “the *good news* . . . The problem with TLS is that we are cursed with *implementations*.”

e.g. Defense vs. Bleichenbacher is in wrong order in OpenSSL.
Does this allow timing attacks?

2014.08 Meyer–Somorovsky–Weiss–Schwenk–Schinzel–Tews: Successful Bleichenbacher attacks, exploiting analogous timing variations in Java SSE, Cavium NITROX SSL accelerator chip.

The whole concept of a “public-key cryptosystem” is a historical accident, dangerously unauthenticated.

2013.01 Green:

State-of-the-art TLS proofs are obviously unsatisfactory.

e.g. Defense vs. Bleichenbacher requires “goofy made-up assumption” for proof.

But that was “the *good news* . . . The problem with TLS is that we are cursed with *implementations*.”

e.g. Defense vs. Bleichenbacher is in wrong order in OpenSSL. Does this allow timing attacks?

2014.08 Meyer–Somorovsky–Weiss–Schwenk–Schinzel–Tews: Successful Bleichenbacher attacks, exploiting analogous timing variations in Java SSE, Cavium NITROX SSL accelerator chip.

The whole concept of a “public-key cryptosystem” is a historical accident, dangerously unauthenticated.

Do we seriously believe that we’ll make HTTPS secure by fixing the implementations?
Fix the bad crypto design.

Green:
-the-art TLS proofs
ously unsatisfactory.
ense vs. Bleichenbacher
“goofy made-up
ion” for proof.
t was “the *good*
The problem with
hat we are cursed with
entations.”
ense vs. Bleichenbacher
ong order in OpenSSL.
is allow timing attacks?

2014.08 Meyer–Somorovsky–
Weiss–Schwenk–Schinzel–Tews:
Successful Bleichenbacher attacks,
exploiting analogous timing
variations in Java SSE, Cavium
NITROX SSL accelerator chip.
The whole concept of a
“public-key cryptosystem”
is a historical accident,
dangerously unauthenticated.
Do we seriously believe that
we’ll make HTTPS secure
by fixing the implementations?
Fix the bad crypto design.

Exercise
failures o
Renegot
Diginota
BEAST
Trustway
CRIME
Lucky 13
RC4 key
TLS tru
gotofail
Triple H
Heartble
POODL
Winshoc

LS proofs
atisfactory.

leichenbacher
ade-up
roof.

good
blem with
e cursed with

leichenbacher
n OpenSSL.
ming attacks?

2014.08 Meyer–Somorovsky–
Weiss–Schwenk–Schinzel–Tews:
Successful Bleichenbacher attacks,
exploiting analogous timing
variations in Java SSE, Cavium
NITROX SSL accelerator chip.

The whole concept of a
“public-key cryptosystem”
is a historical accident,
dangerously unauthenticated.

Do we seriously believe that
we’ll make HTTPS secure
by fixing the implementations?
Fix the bad crypto design.

Exercise: How ma
failures can a *desig*
Renegotiation attac
Diginotar CA com
BEAST CBC attac
Trustwave HTTPS
CRIME compressio
Lucky 13 padding,
RC4 keystream bia
TLS truncation.
gotofail signature-
Triple Handshake.
Heartbleed buffer
POODLE padding
Winshock buffer o

2014.08 Meyer–Somorovsky–
Weiss–Schwenk–Schinzel–Tews:
Successful Bleichenbacher attacks,
exploiting analogous timing
variations in Java SSE, Cavium
NITROX SSL accelerator chip.

The whole concept of a
“public-key cryptosystem”
is a historical accident,
dangerously unauthenticated.

Do we seriously believe that
we’ll make HTTPS secure
by fixing the implementations?
Fix the bad crypto design.

Exercise: How many of these
failures can a *designer* address?

- Renegotiation attack.
- Diginotar CA compromise.
- BEAST CBC attack.
- Trustwave HTTPS intercept.
- CRIME compression attack.
- Lucky 13 padding/timing attack.
- RC4 keystream bias.
- TLS truncation.
- gotofail signature-verification.
- Triple Handshake.
- Heartbleed buffer overread.
- POODLE padding-oracle attack.
- Winshock buffer overflow.

2014.08 Meyer–Somorovsky–
Weiss–Schwenk–Schinzel–Tews:
Successful Bleichenbacher attacks,
exploiting analogous timing
variations in Java SSE, Cavium
NITROX SSL accelerator chip.

The whole concept of a
“public-key cryptosystem”
is a historical accident,
dangerously unauthenticated.

Do we seriously believe that
we’ll make HTTPS secure
by fixing the implementations?
Fix the bad crypto design.

Exercise: How many of these TLS
failures can a *designer* address?

Renegotiation attack.

Diginotar CA compromise.

BEAST CBC attack.

Trustwave HTTPS interception.

CRIME compression attack.

Lucky 13 padding/timing attack.

RC4 keystream bias.

TLS truncation.

gotofail signature-verification bug.

Triple Handshake.

Heartbleed buffer overread.

POODLE padding-oracle attack.

Winshock buffer overflow.