

How cryptographic benchmarking goes wrong

Daniel J. Bernstein

Thanks to NIST 60NANB12D261 for funding this work, and for not reviewing these slides in advance.

PRESERVE, ending 2015.06.30, was a European project “Preparing Secure Vehicle-to-X Communication Systems” .

Project cost: 5383431 EUR, including 3850000 EUR from the European Commission.

“**About PRESERVE**”: “The mission of PRESERVE is, to *design, implement, and test a secure and scalable V2X Security Subsystem for realistic deployment scenarios.* . . . [Expected Results:] 1. Harmonized V2X Security Architecture. 2. Implementation of V2X Security Subsystem. 3. Cheap and scalable security ASIC for V2X. 4. Testing results VSS under realistic conditions. 5. Research results for deployment challenges.”

ptographic benchmarking
ong

. Bernstein

to NIST 60NANB12D261
ing this work, and for not
g these slides in advance.

PREERVE, ending 2015.06.30,
European project
ng Secure Vehicle-to-X
nication Systems” .

cost: 5383431 EUR,
g 3850000 EUR from
European Commission.

1

“**About PRESERVE**”: “The mission of PRESERVE is, to *design, implement, and test a secure and scalable V2X Security Subsystem for realistic deployment scenarios.*

... [Expected Results:] 1. Harmonized V2X Security Architecture. 2. Implementation of V2X Security Subsystem. 3. Cheap and scalable security ASIC for V2X. 4. Testing results VSS under realistic conditions. 5. Research results for deployment challenges.”

2

Cars already
Why build
PRESERVE
“Security
Security
“Process
second a
ms can be
hardware
a Pentium
needs ab
a verifica
cryptogr
likely to

1

“About PRESERVE”: “The mission of PRESERVE is, to *design, implement, and test a secure and scalable V2X Security Subsystem for realistic deployment scenarios.* ... [Expected Results:] 1. Harmonized V2X Security Architecture. 2. Implementation of V2X Security Subsystem. 3. Cheap and scalable security ASIC for V2X. 4. Testing results VSS under realistic conditions. 5. Research results for deployment challenges.”

2

Cars already include
Why build an ASIC
PRESERVE delivered
“Security Requirements
Security Architecture
“Processing 1,000
second and proces
ms **can hardly be
hardware.** As disc
a Pentium D 3.4 C
**needs about 5 tim
a verification ... a
cryptographic co-p
likely to be necess**

1

“About PRESERVE”: “The mission of PRESERVE is, to *design, implement, and test a secure and scalable V2X Security Subsystem for realistic deployment scenarios.* . . . [Expected Results:] 1. Harmonized V2X Security Architecture. 2. Implementation of V2X Security Subsystem. 3. Cheap and scalable security ASIC for V2X. 4. Testing results VSS under realistic conditions. 5. Research results for deployment challenges.”

2

Cars already include many C
Why build an ASIC?
PRESERVE deliverable 1.1,
“Security Requirements of V
Security Architecture”, 2011
“Processing 1,000 packets p
second and processing each
ms **can hardly be met by cu
hardware.** As discussed in [3
a Pentium D 3.4 GHz proces
**needs about 5 times as long
a verification . . . a dedicated
cryptographic co-processor is
likely to be necessary.”**

“About PRESERVE”: “The mission of PRESERVE is, to *design, implement, and test a secure and scalable V2X Security Subsystem for realistic deployment scenarios.* . . . [Expected Results:] 1. Harmonized V2X Security Architecture. 2. Implementation of V2X Security Subsystem. 3. Cheap and scalable security ASIC for V2X. 4. Testing results VSS under realistic conditions. 5. Research results for deployment challenges.”

Cars already include many CPUs. Why build an ASIC?

PRESERVE deliverable 1.1, “Security Requirements of Vehicle Security Architecture”, 2011: “Processing 1,000 packets per second and processing each in 1 ms **can hardly be met by current hardware.** As discussed in [32], a Pentium D 3.4 GHz processor **needs about 5 times as long for a verification . . . a dedicated cryptographic co-processor is likely to be necessary.**”

PRESERVE: “The
of PRESERVE is,
n, implement, and
ecure and scalable
curity Subsystem for
deployment scenarios.
ected Results:] 1.
ized V2X Security
ture. 2. Implementation
Security Subsystem. 3.
nd scalable security ASIC
4. Testing results VSS
alistic conditions. 5.
n results for deployment
es.”

2

Cars already include many CPUs.
Why build an ASIC?

PRESERVE deliverable 1.1,
“Security Requirements of Vehicle
Security Architecture”, 2011:

“Processing 1,000 packets per
second and processing each in 1
ms **can hardly be met by current
hardware.** As discussed in [32],
a Pentium D 3.4 GHz processor
**needs about 5 times as long for
a verification . . . a dedicated
cryptographic co-processor is
likely to be necessary.”**

3

PRESERVE
“Deploy
V4”, 201
ECC sign
second i
factor fo
environm
4mm×4
technolo
space fo
90nm wi
cores an
more.”
max 100

2

Cars already include many CPUs.
Why build an ASIC?

PRESERVE deliverable 1.1,
“Security Requirements of Vehicle
Security Architecture”, 2011:

“Processing 1,000 packets per
second and processing each in 1
ms **can hardly be met by current
hardware**. As discussed in [32],
a Pentium D 3.4 GHz processor
**needs about 5 times as long for
a verification . . . a dedicated
cryptographic co-processor is
likely to be necessary.**”

3

PRESERVE deliverable 1.1,
“Deployment Issues for V4”, 2016: “the number of
ECC signature verification operations per
second is the key performance
factor for ASICs in a high-speed
environment . . . [On a 4mm×4mm chip]
current technology **may only allow
space for one ECC signature
verification**. 90nm will allow for
100 cores and 55nm will allow for
more.” For 180nm technology,
max 100MHz, 1000 operations per second

2

Cars already include many CPUs.
Why build an ASIC?

PRESERVE deliverable 1.1,
“Security Requirements of Vehicle
Security Architecture”, 2011:
“Processing 1,000 packets per
second and processing each in 1
ms **can hardly be met by current
hardware**. As discussed in [32],
a Pentium D 3.4 GHz processor
**needs about 5 times as long for
a verification . . . a dedicated
cryptographic co-processor is
likely to be necessary.**”

3

PRESERVE deliverable 5.4,
“Deployment Issues Report
V4”, 2016: “the number of
ECC signature verifications per
second is the key performance
factor for ASICs in a C2C
environment . . . [On a
4mm×4mm chip] the 180nm
technology **may only yield enough
space for one ECC core**, where
90nm will allow for up to ten
cores and 55nm will allow for
more.” For 180nm core says
max 100MHz, 100 verif/second

Cars already include many CPUs.
Why build an ASIC?

PRESERVE deliverable 1.1,
“Security Requirements of Vehicle
Security Architecture”, 2011:
“Processing 1,000 packets per
second and processing each in 1
ms **can hardly be met by current
hardware**. As discussed in [32],
a Pentium D 3.4 GHz processor
**needs about 5 times as long for
a verification . . . a dedicated
cryptographic co-processor is
likely to be necessary.**”

PRESERVE deliverable 5.4,
“Deployment Issues Report
V4”, 2016: “the number of
ECC signature verifications per
second is the key performance
factor for ASICs in a C2C
environment . . . [On a
4mm×4mm chip] the 180nm
technology **may only yield enough
space for one ECC core**, whereas
90nm will allow for up to ten ECC
cores and 55nm will allow for even
more.” For 180nm core says
max 100MHz, 100 verif/second.

already include many CPUs.
Could an ASIC?

PRESERVE deliverable 1.1,
“Functional Requirements of Vehicle
Architecture”, 2011:
“... processing 1,000 packets per
second and processing each in 1
microsecond... **hardly be met by current
technology. As discussed in [32],
a 3.4 GHz processor
may take about 5 times as long for
signature verification... a dedicated
ASIC or graphic co-processor is
likely to be necessary.**”

3

PRESERVE deliverable 5.4,
“Deployment Issues Report
V4”, 2016: “the number of
ECC signature verifications per
second is the key performance
factor for ASICs in a C2C
environment... [On a
4mm×4mm chip] the 180nm
technology **may only yield enough
space for one ECC core**, whereas
90nm will allow for up to ten ECC
cores and 55nm will allow for even
more.” For 180nm core says
max 100MHz, 100 verif/second.

4

Compare
IAIK NIS
858 scal
in 11162
at 180nm
technolo
standard
9.3744 μ
condition
core volt
Signature
somewha
Still clos
than the

de many CPUs.
C?

table 1.1,
ments of Vehicle
ure”, 2011:

packets per
sing each in 1
met by current
ussed in [32],
GHz processor
es as long for
a dedicated
processor is
ary.”

3

PRESERVE deliverable 5.4,
“Deployment Issues Report
V4”, 2016: “the number of
ECC signature verifications per
second is the key performance
factor for ASICs in a C2C
environment ... [On a
4mm×4mm chip] the 180nm
technology **may only yield enough
space for one ECC core**, whereas
90nm will allow for up to ten ECC
cores and 55nm will allow for even
more.” For 180nm core says
max 100MHz, 100 verif/second.

4

Compare to, e.g.,
IAIK NIST P-256
858 scalarmult/sec
in 111620 GE at 1
at 180nm (“UMC
technology using F
standard cell library
9.3744 $\mu\text{m}^2/\text{GE}$; v
conditions (tempe
core voltage 1.62V
Signature verificat
somewhat slower t
Still close to 100×
than the PRESER

3

PRESERVE deliverable 5.4, “Deployment Issues Report V4”, 2016: “the number of ECC signature verifications per second is the key performance factor for ASICs in a C2C environment . . . [On a 4mm×4mm chip] the 180nm technology **may only yield enough space for one ECC core**, whereas 90nm will allow for up to ten ECC cores and 55nm will allow for even more.” For 180nm core says max 100MHz, 100 verif/second.

4

Compare to, e.g., [IAIK NIST P-256 ECC Mod](#) 858 scalarmult/second in 111620 GE at 192 MHz at 180nm (“UMC L180GII technology using Faraday f1 standard cell library (FSA0A 9.3744 $\mu\text{m}^2/\text{GE}$; worst case conditions (temperature 125 core voltage 1.62V)”).

Signature verification will be somewhat slower than scalar Still close to 100× more efficient than the PRESERVE estimate

PRESERVE deliverable 5.4, “Deployment Issues Report V4”, 2016: “the number of ECC signature verifications per second is the key performance factor for ASICs in a C2C environment . . . [On a 4mm×4mm chip] the 180nm technology **may only yield enough space for one ECC core**, whereas 90nm will allow for up to ten ECC cores and 55nm will allow for even more.” For 180nm core says max 100MHz, 100 verif/second.

Compare to, e.g., **IAIK NIST P-256 ECC Module**: 858 scalarmult/second in 111620 GE at 192 MHz at 180nm (“UMC L180GII technology using Faraday f180 standard cell library (FSA0A_C), 9.3744 $\mu\text{m}^2/\text{GE}$; worst case conditions (temperature 125°C, core voltage 1.62V)”).

Signature verification will be somewhat slower than scalarmult. Still close to 100× more efficient than the PRESERVE estimates.

PREERVE deliverable 5.4,
Implementation Issues Report
16: “the number of
signature verifications per
is the key performance
for ASICs in a C2C
ment ... [On a
mm chip] the 180nm
ogy **may only yield enough
for one ECC core**, whereas
will allow for up to ten ECC
and 55nm will allow for even
For 180nm core says
0MHz, 100 verif/second.

4

Compare to, e.g.,
IAIK NIST P-256 ECC Module:
858 scalarmult/second
in 111620 GE at 192 MHz
at 180nm (“UMC L180GII
technology using Faraday f180
standard cell library (FSA0A_C),
9.3744 $\mu\text{m}^2/\text{GE}$; worst case
conditions (temperature 125°C,
core voltage 1.62V”).

Signature verification will be
somewhat slower than scalarmult.
Still close to 100× more efficient
than the PRESERVE estimates.

5

Let’s go
core arg
Central
in [32], a
processo
(i.e., 17
for signa
[32] is “
Z., ‘Ana
overhead
Third Jo
Mobile M
(WMNC

4

Figure 5.4,
 Report
 number of
 operations per
 performance
 on a C2C
 On a
 the 180nm
 only yield enough
 core, whereas
 up to ten ECC
 will allow for even
 core says
) verif/second.

Compare to, e.g.,
IAIK NIST P-256 ECC Module:
 858 scalarmult/second
 in 111620 GE at 192 MHz
 at 180nm (“UMC L180GII
 technology using Faraday f180
 standard cell library (FSA0A_C),
 9.3744 $\mu\text{m}^2/\text{GE}$; worst case
 conditions (temperature 125°C,
 core voltage 1.62V”).

Signature verification will be
 somewhat slower than scalarmult.
 Still close to 100× more efficient
 than the PRESERVE estimates.

5

Let’s go back to P
 core argument for
 Central claim: “As
 in [32], a Pentium
 processor **needs ab**
 (i.e., 17 million CF
 for signature verifi
 [32] is “Petit, J., M
 Z., ‘Analysis of au
 overhead in vehicu
 Third Joint IFIP V
 Mobile Networking
 (WMNC), 2010.”

4

Compare to, e.g.,

IAIK NIST P-256 ECC Module:

858 scalarmult/second

in 111620 GE at 192 MHz

at 180nm (“UMC L180GII

technology using Faraday f180

standard cell library (FSA0A_C),

$9.3744 \mu\text{m}^2/\text{GE}$; worst case

conditions (temperature 125°C ,
core voltage 1.62V”).

Signature verification will be
somewhat slower than scalarmult.
Still close to $100\times$ more efficient
than the PRESERVE estimates.

5

Let’s go back to PRESERVE
core argument for an ASIC.

Central claim: “As discussed
in [32], a Pentium D 3.4 GHz
processor **needs about**” 5ms
(i.e., 17 million CPU cycles)
for signature verification.

[32] is “Petit, J., Mammeri,
Z., ‘Analysis of authentication
overhead in vehicular network
Third Joint IFIP Wireless and
Mobile Networking Conference
(WMNC), 2010.”

Compare to, e.g.,

IAIK NIST P-256 ECC Module:

858 scalarmult/second
in 111620 GE at 192 MHz
at 180nm (“UMC L180GII
technology using Faraday f180
standard cell library (FSA0A_C),
 $9.3744 \mu\text{m}^2/\text{GE}$; worst case
conditions (temperature 125°C ,
core voltage 1.62V)”).

Signature verification will be
somewhat slower than scalarmult.
Still close to $100\times$ more efficient
than the PRESERVE estimates.

Let’s go back to PRESERVE’s
core argument for an ASIC.

Central claim: “As discussed
in [32], a Pentium D 3.4 GHz
processor **needs about**” 5ms
(i.e., 17 million CPU cycles)
for signature verification.

[32] is “Petit, J., Mammeri,
Z., ‘Analysis of authentication
overhead in vehicular networks’,
Third Joint IFIP Wireless and
Mobile Networking Conference
(WMNC), 2010.”

e to, e.g.,

ST P-256 ECC Module:

armult/second

20 GE at 192 MHz

m (“UMC L180GII

gy using Faraday f180

l cell library (FSA0A_C),

$\mu\text{m}^2/\text{GE}$; worst case

ns (temperature 125°C ,

age 1.62V)”).

re verification will be

at slower than scalarmult.

se to $100\times$ more efficient

e PRESERVE estimates.

5

Let’s go back to PRESERVE’s core argument for an ASIC.

Central claim: “As discussed in [32], a Pentium D 3.4 GHz processor **needs about**” 5ms (i.e., 17 million CPU cycles) for signature verification.

[32] is “Petit, J., Mammeri, Z., ‘Analysis of authentication overhead in vehicular networks’, Third Joint IFIP Wireless and Mobile Networking Conference (WMNC), 2010.”

6

[32] says to the h
economy
from veh
governm
compani
have ma
vehicular
[1]. On
collisions
and 790
United S
economy
[2]. . . .
costing +

5

ECC Module:

cond

92 MHz

L180GII

Faraday f180

y (FSA0A_C),

worst case

perature 125°C,

/)").

ion will be

than scalarmult.

k more efficient

VE estimates.

Let's go back to PRESERVE's core argument for an ASIC.

Central claim: "As discussed in [32], a Pentium D 3.4 GHz processor **needs about**" 5ms (i.e., 17 million CPU cycles) for signature verification.

[32] is "Petit, J., Mammeri, Z., 'Analysis of authentication overhead in vehicular networks', Third Joint IFIP Wireless and Mobile Networking Conference (WMNC), 2010."

6

[32] says "1. Intro to the huge life loss economic impacts from vehicular collisions governments, auto companies, and ins have made the rec vehicular fatalities [1]. On average, v collisions cause 10 and 7900 injuries o United States, leav economic impact o [2]. . . . [Similar st costing €160 billic

5

Let's go back to PRESERVE's core argument for an ASIC.

Central claim: "As discussed in [32], a Pentium D 3.4 GHz processor **needs about**" 5ms (i.e., 17 million CPU cycles) for signature verification.

[32] is "Petit, J., Mammeri, Z., 'Analysis of authentication overhead in vehicular networks', Third Joint IFIP Wireless and Mobile Networking Conference (WMNC), 2010."

6

[32] says "1. Introduction. ... to the huge life losses and the economic impacts resulting from vehicular collisions, many governments, automotive companies, and industry corporations have made the reduction of vehicular fatalities a top priority [1]. On average, vehicular collisions cause 102 deaths and 7900 injuries daily in the United States, leaving an economic impact of \$230 billion [2]. ... [Similar story for EU costing €160 billion annually]

Let's go back to PRESERVE's core argument for an ASIC.

Central claim: "As discussed in [32], a Pentium D 3.4 GHz processor **needs about**" 5ms (i.e., 17 million CPU cycles) for signature verification.

[32] is "Petit, J., Mammeri, Z., 'Analysis of authentication overhead in vehicular networks', Third Joint IFIP Wireless and Mobile Networking Conference (WMNC), 2010."

[32] says "1. Introduction. Due to the huge life losses and the economic impacts resulting from vehicular collisions, many governments, automotive companies, and industry consortia have made the reduction of vehicular fatalities a top priority [1]. On average, vehicular collisions cause 102 deaths and 7900 injuries daily in the United States, leaving an economic impact of \$230 billion [2]. . . . [Similar story for EU:] costing €160 billion annually [3]."

back to PRESERVE's
document for an ASIC.

claim: "As discussed
a Pentium D 3.4 GHz
or **needs about**" 5ms
million CPU cycles)
ature verification.

Petit, J., Mammeri,
lysis of authentication
d in vehicular networks',
oint IFIP Wireless and
Networking Conference
) , 2010."

6

[32] says "1. Introduction. Due to the huge life losses and the economic impacts resulting from vehicular collisions, many governments, automotive companies, and industry consortia have made the reduction of vehicular fatalities a top priority [1]. On average, vehicular collisions cause 102 deaths and 7900 injuries daily in the United States, leaving an economic impact of \$230 billion [2]. . . . [Similar story for EU:] costing €160 billion annually [3]."

7

Vehicles
informat
of IEEE1
support
Signatur
[8] over
P-224 an
paper, w
and com
the auth
provided
II. Signa
verificati
D 3.4Gh

PRESERVE's
an ASIC.

s discussed
D 3.4 GHz
out" 5ms
PU cycles)
cation.

Mammeri,
thentication
lar networks',
Wireless and
g Conference

6

[32] says "1. Introduction. Due to the huge life losses and the economic impacts resulting from vehicular collisions, many governments, automotive companies, and industry consortia have made the reduction of vehicular fatalities a top priority [1]. On average, vehicular collisions cause 102 deaths and 7900 injuries daily in the United States, leaving an economic impact of \$230 billion [2]. . . . [Similar story for EU:] costing €160 billion annually [3]."

7

Vehicles will comm
information. "All
of IEEE1609.2 sta
support the Elliptic
Signature Algorithm
[8] over the two N
P-224 and P-256.
paper, we assess t
and communication
the authentication
provided by ECDS
II. Signature gener
verification times o
D 3.4Ghz workstat

[32] says “1. Introduction. Due to the huge life losses and the economic impacts resulting from vehicular collisions, many governments, automotive companies, and industry consortia have made the reduction of vehicular fatalities a top priority [1]. On average, vehicular collisions cause 102 deaths and 7900 injuries daily in the United States, leaving an economic impact of \$230 billion [2]. . . . [Similar story for EU:] costing €160 billion annually [3].”

Vehicles will communicate sensitive information. “All implementations of IEEE1609.2 standard [7] support the Elliptic Curve Digital Signature Algorithm (ECDSA) [8] over the two NIST curves P-224 and P-256. . . . In this paper, we assess the processing and communication overhead of the authentication mechanism provided by ECDSA. . . . Table II. Signature generation and verification times on a Pentium D 3.4Ghz workstation [10]”

[32] says “1. Introduction. Due to the huge life losses and the economic impacts resulting from vehicular collisions, many governments, automotive companies, and industry consortia have made the reduction of vehicular fatalities a top priority [1]. On average, vehicular collisions cause 102 deaths and 7900 injuries daily in the United States, leaving an economic impact of \$230 billion [2]. . . . [Similar story for EU:] costing €160 billion annually [3].”

Vehicles will communicate safety information. “All implementations of IEEE1609.2 standard [7] shall support the Elliptic Curve Digital Signature Algorithm (ECDSA) [8] over the two NIST curves P-224 and P-256. . . . In this paper, we assess the processing and communication overhead of the authentication mechanism provided by ECDSA. . . . Table II. Signature generation and verification times on a Pentium D 3.4Ghz workstation [10]”

s “1. Introduction. Due
uge life losses and the
c impacts resulting
nicular collisions, many
ents, automotive
es, and industry consortia
de the reduction of
r fatalities a top priority
average, vehicular
s cause 102 deaths
0 injuries daily in the
States, leaving an
c impact of \$230 billion
[Similar story for EU:]
€160 billion annually [3].”

7

Vehicles will communicate safety information. “All implementations of IEEE1609.2 standard [7] shall support the Elliptic Curve Digital Signature Algorithm (ECDSA) [8] over the two NIST curves P-224 and P-256. . . . In this paper, we assess the processing and communication overhead of the authentication mechanism provided by ECDSA. . . . Table II. Signature generation and verification times on a Pentium D 3.4Ghz workstation [10]”

8

[10] (in
J., ‘Anal
Authent
VANETS
Conferen
Mobility
Cairo, D
[10] says
impleme
and follo
For NIS
“Pentium
2.50ms/
4.97ms/

duction. Due
 sses and the
 resulting
 isions, many
 omotive
 dustry consortia
 duction of
 a top priority
 ehicular
 2 deaths
 daily in the
 ving an
 of \$230 billion
 ory for EU:]
 on annually [3].”

Vehicles will communicate safety information. “All implementations of IEEE1609.2 standard [7] shall support the Elliptic Curve Digital Signature Algorithm (ECDSA) [8] over the two NIST curves P-224 and P-256. . . . In this paper, we assess the processing and communication overhead of the authentication mechanism provided by ECDSA. . . . Table II. Signature generation and verification times on a Pentium D 3.4Ghz workstation [10]”

[10] (in [32]) is “P. J., ‘Analysis of ECDSA Authentication Protocol for VANETs’, 3rd IFIP Conference on New Mobility and Security, Cairo, December 2006.” [10] says “ECDSA implemented using . . . and following the . . . For NIST P-224/P-256 “Pentium D 3.4GHz 2.50ms/3.33ms to 4.97ms/6.63ms to

Vehicles will communicate safety information. “All implementations of IEEE1609.2 standard [7] shall support the Elliptic Curve Digital Signature Algorithm (ECDSA) [8] over the two NIST curves P-224 and P-256. . . . In this paper, we assess the processing and communication overhead of the authentication mechanism provided by ECDSA. . . . Table II. Signature generation and verification times on a Pentium D 3.4Ghz workstation [10]”

[10] (in [32]) is “Petit J., ‘Analysis of ECDSA Authentication Processing in VANETs’, 3rd IFIP International Conference on New Technologies for Mobility and Security (NTMS), Cairo, December 2009.”

[10] says “ECDSA was implemented using MIRACL and following the Fig.1.” For NIST P-224/P-256 on “Pentium D 3.4GHz workstation 2.50ms/3.33ms to sign, 4.97ms/6.63ms to verify.

Vehicles will communicate safety information. “All implementations of IEEE1609.2 standard [7] shall support the Elliptic Curve Digital Signature Algorithm (ECDSA) [8] over the two NIST curves P-224 and P-256. . . . In this paper, we assess the processing and communication overhead of the authentication mechanism provided by ECDSA. . . . Table II. Signature generation and verification times on a Pentium D 3.4Ghz workstation [10]”

[10] (in [32]) is “Petit J., ‘Analysis of ECDSA Authentication Processing in VANETs’, 3rd IFIP International Conference on New Technologies, Mobility and Security (NTMS), Cairo, December 2009.”

[10] says “ECDSA was implemented using MIRACL and following the Fig.1.”

For NIST P-224/P-256 on “Pentium D 3.4GHz workstation” :
2.50ms/3.33ms to sign,
4.97ms/6.63ms to verify.

will communicate safety
tion. “All implementations
1609.2 standard [7] shall
the Elliptic Curve Digital
Signature Algorithm (ECDSA)
the two NIST curves
and P-256. . . . In this
we assess the processing
communication overhead of
authentication mechanism
by ECDSA. . . . Table
signature generation and
verification times on a Pentium
3GHz workstation [10]”

8

[10] (in [32]) is “Petit
J., ‘Analysis of ECDSA
Authentication Processing in
VANETs’, 3rd IFIP International
Conference on New Technologies,
Mobility and Security (NTMS),
Cairo, December 2009.”

[10] says “ECDSA was
implemented using MIRACL
and following the Fig.1.”
For NIST P-224/P-256 on
“Pentium D 3.4GHz workstation”:
2.50ms/3.33ms to sign,
4.97ms/6.63ms to verify.

9

Compare
speeds r
of 14nm
(“2015 I
<https://>
0.015ms
0.049ms

communicate safety
 implementations
 standard [7] shall
 Curve Digital
 m (ECDSA)
 IST curves
 ... In this
 he processing
 on overhead of
 mechanism
 A. ... Table
 ration and
 on a Pentium
 tion [10]"

[10] (in [32]) is "Petit
 J., 'Analysis of ECDSA
 Authentication Processing in
 VANETs', 3rd IFIP International
 Conference on New Technologies,
 Mobility and Security (NTMS),
 Cairo, December 2009."

[10] says "ECDSA was
 implemented using MIRACL
 and following the Fig.1."
 For NIST P-224/P-256 on
 "Pentium D 3.4GHz workstation":
 2.50ms/3.33ms to sign,
 4.97ms/6.63ms to verify.

Compare to, e.g.,
 speeds reported for
 of 14nm 3.31GHz
 ("2015 Intel Core
<https://bench.c>
 0.015ms to sign (4
 0.049ms to verify

8

[10] (in [32]) is “Petit J., ‘Analysis of ECDSA Authentication Processing in VANETs’, 3rd IFIP International Conference on New Technologies, Mobility and Security (NTMS), Cairo, December 2009.”

[10] says “ECDSA was implemented using MIRACL and following the Fig.1.”

For NIST P-224/P-256 on “Pentium D 3.4GHz workstation” :
2.50ms/3.33ms to sign,
4.97ms/6.63ms to verify.

9

Compare to, e.g., Ed25519 speeds reported for single core of 14nm 3.31GHz Skylake (“2015 Intel Core i5-6600”) <https://bench.cr.yp.to>:

0.015ms to sign (49840 cycles)
0.049ms to verify (163206 cycles)

[10] (in [32]) is “Petit J., ‘Analysis of ECDSA Authentication Processing in VANETs’, 3rd IFIP International Conference on New Technologies, Mobility and Security (NTMS), Cairo, December 2009.”

[10] says “ECDSA was implemented using MIRACL and following the Fig.1.”
For NIST P-224/P-256 on “Pentium D 3.4GHz workstation”:
2.50ms/3.33ms to sign,
4.97ms/6.63ms to verify.

Compare to, e.g., Ed25519 speeds reported for single core of 14nm 3.31GHz Skylake (“2015 Intel Core i5-6600”) on <https://bench.cr.yp.to>:
0.015ms to sign (49840 cycles),
0.049ms to verify (163206 cycles).

[10] (in [32]) is “Petit J., ‘Analysis of ECDSA Authentication Processing in VANETs’, 3rd IFIP International Conference on New Technologies, Mobility and Security (NTMS), Cairo, December 2009.”

[10] says “ECDSA was implemented using MIRACL and following the Fig.1.”

For NIST P-224/P-256 on “Pentium D 3.4GHz workstation”:
2.50ms/3.33ms to sign,
4.97ms/6.63ms to verify.

Compare to, e.g., Ed25519 speeds reported for single core of 14nm 3.31GHz Skylake (“2015 Intel Core i5-6600”) on <https://bench.cr.yp.to>:
0.015ms to sign (49840 cycles),
0.049ms to verify (163206 cycles).

This chip didn’t exist in 2009.
Compare instead to single core of 65nm 2.4GHz Core 2 (“2007 Intel Core 2 Quad Q6600”).
0.065ms to sign (156843 cycles),
0.232ms to verify (557082 cycles).

[32]) is “Petit
 Analysis of ECDSA
 Application Processing in
 s’, 3rd IFIP International
 ence on New Technologies,
 and Security (NTMS),
 December 2009.”

s “ECDSA was
 nted using MIRACL
 owing the Fig.1.”
 T P-224/P-256 on
 m D 3.4GHz workstation”:
 3.33ms to sign,
 6.63ms to verify.

Compare to, e.g., Ed25519
 speeds reported for single core
 of 14nm 3.31GHz Skylake
 (“2015 Intel Core i5-6600”) on
<https://bench.cr.yp.to>:
 0.015ms to sign (49840 cycles),
 0.049ms to verify (163206 cycles).

This chip didn’t exist in 2009.
 Compare instead to single core
 of 65nm 2.4GHz Core 2 (“2007
 Intel Core 2 Quad Q6600”).
 0.065ms to sign (156843 cycles),
 0.232ms to verify (557082 cycles).

2012 Be
 on 720M
 0.9ms to
 ARM Co
 1000MH
 in iPad
 1000MH
 in Sams
 1000MH
 Motorola
 800MHz
 Amazon
 Today:
 Cortex-A

Petit
 CDSA
 processing in
 P International
 w Technologies,
 rity (NTMS),
 2009.”

was
 g MIRACL
 Fig.1.”
 P-256 on
 Hz workstation”:
 sign,
 verify.

Compare to, e.g., Ed25519
 speeds reported for single core
 of 14nm 3.31GHz Skylake
 (“2015 Intel Core i5-6600”) on
<https://bench.cr.yp.to>:
 0.015ms to sign (49840 cycles),
 0.049ms to verify (163206 cycles).

This chip didn’t exist in 2009.
 Compare instead to single core
 of 65nm 2.4GHz Core 2 (“2007
 Intel Core 2 Quad Q6600”).
 0.065ms to sign (156843 cycles),
 0.232ms to verify (557082 cycles).

2012 Bernstein–Sc
 on 720MHz ARM
 0.9ms to verify (65
 ARM Cortex-A8 c
 1000MHz Apple A
 in iPad 1, iPhone
 1000MHz Samsun
 in Samsung Galaxy
 1000MHz TI OMA
 Motorola Droid X
 800MHz Freescale
 Amazon Kindle 4
 Today: in CPUs c
 Cortex-A7 is even

Compare to, e.g., Ed25519 speeds reported for single core of 14nm 3.31GHz Skylake (“2015 Intel Core i5-6600”) on <https://bench.cr.yp.to>:
 0.015ms to sign (49840 cycles),
 0.049ms to verify (163206 cycles).

This chip didn't exist in 2009.

Compare instead to single core of 65nm 2.4GHz Core 2 (“2007 Intel Core 2 Quad Q6600”).

0.065ms to sign (156843 cycles),
 0.232ms to verify (557082 cycles).

2012 Bernstein–Schwabe on 720MHz ARM Cortex-A8
 0.9ms to verify (650102 cycles)
 ARM Cortex-A8 cores were
 1000MHz Apple A4 in iPad 1, iPhone 4 (2010);
 1000MHz Samsung Exynos in Samsung Galaxy S (2010);
 1000MHz TI OMAP3630 in Motorola Droid X (2010);
 800MHz Freescale i.MX50 in Amazon Kindle 4 (2011); ...
 Today: in CPUs costing ≈ 2
 Cortex-A7 is even more pop

Compare to, e.g., Ed25519 speeds reported for single core of 14nm 3.31GHz Skylake (“2015 Intel Core i5-6600”) on <https://bench.cr.yp.to>:
 0.015ms to sign (49840 cycles),
 0.049ms to verify (163206 cycles).

This chip didn’t exist in 2009.
 Compare instead to single core of 65nm 2.4GHz Core 2 (“2007 Intel Core 2 Quad Q6600”).
 0.065ms to sign (156843 cycles),
 0.232ms to verify (557082 cycles).

2012 Bernstein–Schwabe on 720MHz ARM Cortex-A8: 0.9ms to verify (650102 cycles).

ARM Cortex-A8 cores were in 1000MHz Apple A4 in iPad 1, iPhone 4 (2010); 1000MHz Samsung Exynos 3110 in Samsung Galaxy S (2010); 1000MHz TI OMAP3630 in Motorola Droid X (2010); 800MHz Freescale i.MX50 in Amazon Kindle 4 (2011); ...
 Today: in CPUs costing ≈ 2 EUR. Cortex-A7 is even more popular.

to, e.g., Ed25519
 reported for single core
 3.31GHz Skylake
 Intel Core i5-6600") on
[/bench.cr.yp.to](http://bench.cr.yp.to):
 to sign (49840 cycles),
 to verify (163206 cycles).
 p didn't exist in 2009.
 e instead to single core
 2.4GHz Core 2 ("2007
 re 2 Quad Q6600").
 to sign (156843 cycles),
 to verify (557082 cycles).

2012 Bernstein–Schwabe
 on 720MHz ARM Cortex-A8:
 0.9ms to verify (650102 cycles).
 ARM Cortex-A8 cores were in
 1000MHz Apple A4
 in iPad 1, iPhone 4 (2010);
 1000MHz Samsung Exynos 3110
 in Samsung Galaxy S (2010);
 1000MHz TI OMAP3630 in
 Motorola Droid X (2010);
 800MHz Freescale i.MX50 in
 Amazon Kindle 4 (2011); ...
 Today: in CPUs costing ≈ 2 EUR.
 Cortex-A7 is even more popular.

180nm 3
 ("2001 I
 0.46ms
 for Curv
 using flo
 Integer r
 Nobody
 adapting
 Would b
 3.4GHz
 same ba
 more ins
 Ed25519
 on one c

Ed25519
 on single core
 Skylake
 i5-6600") on
cr.yp.to:
 49840 cycles),
 (163206 cycles).
 exist in 2009.
 on single core
 Core 2 ("2007
 Q6600").
 156843 cycles),
 (557082 cycles).

2012 Bernstein–Schwabe
 on 720MHz ARM Cortex-A8:
 0.9ms to verify (650102 cycles).
 ARM Cortex-A8 cores were in
 1000MHz Apple A4
 in iPad 1, iPhone 4 (2010);
 1000MHz Samsung Exynos 3110
 in Samsung Galaxy S (2010);
 1000MHz TI OMAP3630 in
 Motorola Droid X (2010);
 800MHz Freescale i.MX50 in
 Amazon Kindle 4 (2011); ...
 Today: in CPUs costing ≈ 2 EUR.
 Cortex-A7 is even more popular.

180nm 32-bit 2GHz
 ("2001 Intel Pentium
 0.46ms (0.9 million
 for Curve25519 sc
 using floating-point
 Integer multiplier i
 Nobody has ever b
 adapting this to si
 Would be ≈ 0.6 ms
 3.4GHz Pentium D
 same basic microa
 more instructions,
 Ed25519 would be
 on one core than l

2012 Bernstein–Schwabe
 on 720MHz ARM Cortex-A8:
 0.9ms to verify (650102 cycles).

ARM Cortex-A8 cores were in
 1000MHz Apple A4
 in iPad 1, iPhone 4 (2010);
 1000MHz Samsung Exynos 3110
 in Samsung Galaxy S (2010);
 1000MHz TI OMAP3630 in
 Motorola Droid X (2010);
 800MHz Freescale i.MX50 in
 Amazon Kindle 4 (2011); ...

Today: in CPUs costing ≈ 2 EUR.
 Cortex-A7 is even more popular.

180nm 32-bit 2GHz Willamette
 (“2001 Intel Pentium 4”):
 0.46ms (0.9 million cycles)
 for Curve25519 scalarmult
 using floating-point multiplication.
 Integer multiplier is much slower.

Nobody has ever bothered
 adapting this to signatures.
 Would be ≈ 0.6 ms for verify.

3.4GHz Pentium D (dual core)
 same basic microarchitecture
 more instructions, faster clock.
 Ed25519 would be $> 10\times$ faster
 on one core than Petit’s software.

2012 Bernstein–Schwabe
 on 720MHz ARM Cortex-A8:
 0.9ms to verify (650102 cycles).

ARM Cortex-A8 cores were in
 1000MHz Apple A4
 in iPad 1, iPhone 4 (2010);
 1000MHz Samsung Exynos 3110
 in Samsung Galaxy S (2010);
 1000MHz TI OMAP3630 in
 Motorola Droid X (2010);
 800MHz Freescale i.MX50 in
 Amazon Kindle 4 (2011); ...

Today: in CPUs costing ≈ 2 EUR.
 Cortex-A7 is even more popular.

180nm 32-bit 2GHz Willamette
 (“2001 Intel Pentium 4”):
 0.46ms (0.9 million cycles)
 for Curve25519 scalarmult
 using floating-point multiplier.
 Integer multiplier is much slower!

Nobody has ever bothered
 adapting this to signatures.
 Would be ≈ 0.6 ms for verify.

3.4GHz Pentium D (dual core):
 same basic microarchitecture,
 more instructions, faster clock.
 Ed25519 would be $> 10\times$ faster
 on one core than Petit’s software.

rnstein–Schwabe
 MHz ARM Cortex-A8:
 to verify (650102 cycles).
 Cortex-A8 cores were in
 GHz Apple A4
 1, iPhone 4 (2010);
 GHz Samsung Exynos 3110
 Samsung Galaxy S (2010);
 GHz TI OMAP3630 in
 a Droid X (2010);
 GHz Freescale i.MX50 in
 Kindle 4 (2011); ...
 in CPUs costing ≈ 2 EUR.
 A7 is even more popular.

180nm 32-bit 2GHz Willamette
 (“2001 Intel Pentium 4”):
 0.46ms (0.9 million cycles)
 for Curve25519 scalarmult
 using floating-point multiplier.
 Integer multiplier is much slower!
 Nobody has ever bothered
 adapting this to signatures.
 Would be ≈ 0.6 ms for verify.
 3.4GHz Pentium D (dual core):
 same basic microarchitecture,
 more instructions, faster clock.
 Ed25519 would be $>10\times$ faster
 on one core than Petit’s software.

Bad ECC
 certainly
 • can’t u
 • can’t u
 • need a
 etc. Typ
 2000 Bro
 Menezes
 4.0ms/6
 cycles) f
 inside N
 2001 Be
 0.7 millio
 for NIST

chwabe
 Cortex-A8:
 50102 cycles).
 cores were in
 A4
 4 (2010);
 g Exynos 3110
 y S (2010);
 AP3630 in
 (2010);
 i.MX50 in
 (2011); ...
 osting ≈ 2 EUR.
 more popular.

180nm 32-bit 2GHz Willamette
 (“2001 Intel Pentium 4”):
 0.46ms (0.9 million cycles)
 for Curve25519 scalarmult
 using floating-point multiplier.
 Integer multiplier is much slower!
 Nobody has ever bothered
 adapting this to signatures.
 Would be ≈ 0.6 ms for verify.
 3.4GHz Pentium D (dual core):
 same basic microarchitecture,
 more instructions, faster clock.
 Ed25519 would be $> 10\times$ faster
 on one core than Petit’s software.

Bad ECDSA-NIST
 certainly has some
 • can’t use fastest
 • can’t use fastest
 • need an annoying
 etc. Typical estim
 2000 Brown–Hank
 Menezes on 400M
 4.0ms/6.4ms (1.6,
 cycles) for double
 inside NIST P-224
 2001 Bernstein, \approx
 0.7 million cycles
 for NIST P-224 sc

180nm 32-bit 2GHz Willamette
 (“2001 Intel Pentium 4”):

0.46ms (0.9 million cycles)
 for Curve25519 scalarmult
 using floating-point multiplier.
 Integer multiplier is much slower!

Nobody has ever bothered
 adapting this to signatures.
 Would be ≈ 0.6 ms for verify.

3.4GHz Pentium D (dual core):
 same basic microarchitecture,
 more instructions, faster clock.
 Ed25519 would be $> 10\times$ faster
 on one core than Petit’s software.

Bad ECDSA-NIST-P-256 de
 certainly has some impact:

- can’t use fastest mulmods
 - can’t use fastest curve form
 - need an annoying inversion
- etc. Typical estimate: $2\times$ sl

2000 Brown–Hankerson–López
 Menezes on 400MHz Pentium
 4.0ms/6.4ms (1.6/2.6 million
 cycles) for double scalarmult
 inside NIST P-224/P-256 ve

2001 Bernstein, $\approx 1.6\times$ faster
 0.7 million cycles on Pentium
 for NIST P-224 scalarmult.

180nm 32-bit 2GHz Willamette
 (“2001 Intel Pentium 4”):

0.46ms (0.9 million cycles)
 for Curve25519 scalarmult
 using floating-point multiplier.
 Integer multiplier is much slower!

Nobody has ever bothered
 adapting this to signatures.
 Would be ≈ 0.6 ms for verify.

3.4GHz Pentium D (dual core):
 same basic microarchitecture,
 more instructions, faster clock.
 Ed25519 would be $>10\times$ faster
 on one core than Petit’s software.

Bad ECDSA-NIST-P-256 design
 certainly has some impact:

- can’t use fastest mulmods;
 - can’t use fastest curve formulas;
 - need an annoying inversion;
- etc. Typical estimate: $2\times$ slower.

2000 Brown–Hankerson–López–
 Menezes on 400MHz Pentium II:
 4.0ms/6.4ms (1.6/2.6 million
 cycles) for double scalarmult
 inside NIST P-224/P-256 verif.

2001 Bernstein, $\approx 1.6\times$ faster:
 0.7 million cycles on Pentium II
 for NIST P-224 scalarmult.

32-bit 2GHz Willamette
Intel Pentium 4”):

(0.9 million cycles)

e25519 scalarmult

adding-point multiplier.

multiplier is much slower!

has ever bothered

g this to signatures.

time $\approx 0.6\text{ms}$ for verify.

Pentium D (dual core):

basic microarchitecture,

instructions, faster clock.

It would be $>10\times$ faster

per core than Petit’s software.

Bad ECDSA-NIST-P-256 design
certainly has some impact:

- can’t use fastest mulmods;
 - can’t use fastest curve formulas;
 - need an annoying inversion;
- etc. Typical estimate: $2\times$ slower.

2000 Brown–Hankerson–López–

Menezes on 400MHz Pentium II:

4.0ms/6.4ms (1.6/2.6 million

cycles) for double scalarmult

inside NIST P-224/P-256 verif.

2001 Bernstein, $\approx 1.6\times$ faster:

0.7 million cycles on Pentium II

for NIST P-224 scalarmult.

2000 Bro

Menezes

cycles on

e.g., P-2

1.2 millio

2.7 millio

2001 Be

0.7 millio

0.8 millio

0.9 millio

using co

OpenSS

2.0 millio

400MHz Willamette
 (Pentium 4"):
 (in cycles)
 scalarmult
 fast multiplier.
 is much slower!

bothered
 signatures.
 for verify.

D (dual core):
 architecture,
 faster clock.
 is $>10\times$ faster
 Petit's software.

Bad ECDSA-NIST-P-256 design
 certainly has some impact:

- can't use fastest mulmods;
- can't use fastest curve formulas;
- need an annoying inversion;

etc. Typical estimate: $2\times$ slower.

2000 Brown–Hankerson–López–
 Menezes on 400MHz Pentium II:
 4.0ms/6.4ms (1.6/2.6 million
 cycles) for double scalarmult
 inside NIST P-224/P-256 verif.

2001 Bernstein, $\approx 1.6\times$ faster:
 0.7 million cycles on Pentium II
 for NIST P-224 scalarmult.

2000 Brown–Hankerson–
 Menezes software
 cycles on P4 than
 e.g., P-224 scalarmult
 1.2 million cycles
 2.7 million cycles

2001 Bernstein P-
 0.7 million cycles
 0.8 million cycles
 0.9 million cycles
 using compressed

OpenSSL 1.0.1, P-
 2.0 million cycles

Bad ECDSA-NIST-P-256 design certainly has some impact:

- can't use fastest mulmods;
 - can't use fastest curve formulas;
 - need an annoying inversion;
- etc. Typical estimate: $2\times$ slower.

2000 Brown–Hankerson–López–Menezes on 400MHz Pentium II: 4.0ms/6.4ms (1.6/2.6 million cycles) for double scalarmult inside NIST P-224/P-256 verif.

2001 Bernstein, $\approx 1.6\times$ faster: 0.7 million cycles on Pentium II for NIST P-224 scalarmult.

2000 Brown–Hankerson–López–Menezes software uses many cycles on P4 than on PII.

e.g., P-224 scalarmult: 1.2 million cycles on Pentium 2.7 million cycles on Pentium

2001 Bernstein P-224 scalar 0.7 million cycles on Pentium 0.8 million cycles on Pentium 0.9 million cycles on Pentium using compressed keys.

OpenSSL 1.0.1, P-224 verif: 2.0 million cycles on Pentium

Bad ECDSA-NIST-P-256 design certainly has some impact:

- can't use fastest mulmods;
 - can't use fastest curve formulas;
 - need an annoying inversion;
- etc. Typical estimate: $2\times$ slower.

2000 Brown–Hankerson–López–Menezes on 400MHz Pentium II: 4.0ms/6.4ms (1.6/2.6 million cycles) for double scalarmult inside NIST P-224/P-256 verif.

2001 Bernstein, $\approx 1.6\times$ faster: 0.7 million cycles on Pentium II for NIST P-224 scalarmult.

2000 Brown–Hankerson–López–Menezes software uses many more cycles on P4 than on PII.

e.g., P-224 scalarmult:

1.2 million cycles on Pentium II.

2.7 million cycles on Pentium 4.

2001 Bernstein P-224 scalarmult:

0.7 million cycles on Pentium II.

0.8 million cycles on Pentium 4.

0.9 million cycles on Pentium 4 using compressed keys.

OpenSSL 1.0.1, P-224 verif:

2.0 million cycles on Pentium D.

ECDSA-NIST-P-256 design
 has some impact:
 use fastest mulmods;
 use fastest curve formulas;
 an annoying inversion;
 typical estimate: $2\times$ slower.
 Brown–Hankerson–López–
 on 400MHz Pentium II:
 .4ms (1.6/2.6 million
 for double scalarmult
 NIST P-224/P-256 verif.
 Bernstein, $\approx 1.6\times$ faster:
 on cycles on Pentium II
 P-224 scalarmult.

2000 Brown–Hankerson–López–
 Menezes software uses many more
 cycles on P4 than on PII.

e.g., P-224 scalarmult:

1.2 million cycles on Pentium II.

2.7 million cycles on Pentium 4.

2001 Bernstein P-224 scalarmult:

0.7 million cycles on Pentium II.

0.8 million cycles on Pentium 4.

0.9 million cycles on Pentium 4

using compressed keys.

OpenSSL 1.0.1, P-224 verif:

2.0 million cycles on Pentium D.

How did
 17 millic
 22 millic
 Presuma
 bad mul
 Why did
 ECDSA,
 underlyi
 Why did
 previous
 Why did
 Why did

-P-256 design
 e impact:
 mulmods;
 curve formulas;
 g inversion;
 ate: $2\times$ slower.
 kerson-López-
 Hz Pentium II:
 /2.6 million
 scalarmult
 /P-256 verif.
 $1.6\times$ faster:
 on Pentium II
 scalarmult.

2000 Brown-Hankerson-López-
 Menezes software uses many more
 cycles on P4 than on PII.

e.g., P-224 scalarmult:

1.2 million cycles on Pentium II.

2.7 million cycles on Pentium 4.

2001 Bernstein P-224 scalarmult:

0.7 million cycles on Pentium II.

0.8 million cycles on Pentium 4.

0.9 million cycles on Pentium 4

using compressed keys.

OpenSSL 1.0.1, P-224 verif:

2.0 million cycles on Pentium D.

How did Petit man
 17 million cycles fo
 22 million cycles fo

Presumably some
 bad mulmod and b

Why did Petit rein
 ECDSA, using MI
 underlying arithme

Why did Petit not
 previous speed lite

Why did Petit cho

Why did BHLM cl

2000 Brown–Hankerson–López–Menezes software uses many more cycles on P4 than on PII.

e.g., P-224 scalarmult:

1.2 million cycles on Pentium II.

2.7 million cycles on Pentium 4.

2001 Bernstein P-224 scalarmult:

0.7 million cycles on Pentium II.

0.8 million cycles on Pentium 4.

0.9 million cycles on Pentium 4

using compressed keys.

OpenSSL 1.0.1, P-224 verific:

2.0 million cycles on Pentium D.

How did Petit manage to use 17 million cycles for P-224 and 22 million cycles for P-256?

Presumably some combination of bad mulmod and bad curve?

Why did Petit reimplement ECDSA, using MIRACL for the underlying arithmetic?

Why did Petit not simply cite previous speed literature?

Why did Petit choose Pentium?

Why did BHLM choose PII?

2000 Brown–Hankerson–López–Menezes software uses many more cycles on P4 than on PII.

e.g., P-224 scalarmult:

1.2 million cycles on Pentium II.

2.7 million cycles on Pentium 4.

2001 Bernstein P-224 scalarmult:

0.7 million cycles on Pentium II.

0.8 million cycles on Pentium 4.

0.9 million cycles on Pentium 4

using compressed keys.

OpenSSL 1.0.1, P-224 verific:

2.0 million cycles on Pentium D.

How did Petit manage to use 17 million cycles for P-224 verific, 22 million cycles for P-256 verific?

Presumably some combination of bad mulmod and bad curve ops.

Why did Petit reimplement ECDSA, using MIRACL for the underlying arithmetic?

Why did Petit not simply cite previous speed literature?

Why did Petit choose Pentium D?

Why did BHLM choose PII?

own–Hankerson–López–
s software uses many more
n P4 than on PII.

224 scalarmult:

on cycles on Pentium II.

on cycles on Pentium 4.

rnstein P-224 scalarmult:

on cycles on Pentium II.

on cycles on Pentium 4.

on cycles on Pentium 4

mpressed keys.

L 1.0.1, P-224 verif:

on cycles on Pentium D.

How did Petit manage to use
17 million cycles for P-224 verif,
22 million cycles for P-256 verif?

Presumably some combination of
bad mulmod and bad curve ops.

Why did Petit reimplement
ECDSA, using MIRACL for the
underlying arithmetic?

Why did Petit not simply cite
previous speed literature?

Why did Petit choose Pentium D?

Why did BHLM choose PII?

Petit: “
cryptogr
OpenSS
Authors
comparis
that MIR
performa
elliptic c

erson–López–
uses many more
on PII.

mult:

on Pentium II.

on Pentium 4.

224 scalarmult:

on Pentium II.

on Pentium 4.

on Pentium 4

keys.

-224 verif:

on Pentium D.

How did Petit manage to use
17 million cycles for P-224 verif,
22 million cycles for P-256 verif?

Presumably some combination of
bad mulmod and bad curve ops.

Why did Petit reimplement
ECDSA, using MIRACL for the
underlying arithmetic?

Why did Petit not simply cite
previous speed literature?

Why did Petit choose Pentium D?

Why did BHLM choose PII?

Petit: “There are
cryptographic libra
OpenSSL and Cry
Authors in [21] pro
comparison and co
that MIRACL has
performance for op
elliptic curves over

How did Petit manage to use 17 million cycles for P-224 verif, 22 million cycles for P-256 verif?

Presumably some combination of bad mulmod and bad curve ops.

Why did Petit reimplement ECDSA, using MIRACL for the underlying arithmetic?

Why did Petit not simply cite previous speed literature?

Why did Petit choose Pentium D?

Why did BHLM choose PII?

Petit: “There are three main cryptographic libraries: MIRACL, OpenSSL and Crypto++.
Authors in [21] proposed a comparison and concluded that MIRACL has the best performance for operations on elliptic curves over binary fields.”

How did Petit manage to use 17 million cycles for P-224 verif, 22 million cycles for P-256 verif?

Presumably some combination of bad mulmod and bad curve ops.

Why did Petit reimplement ECDSA, using MIRACL for the underlying arithmetic?

Why did Petit not simply cite previous speed literature?

Why did Petit choose Pentium D?

Why did BHLM choose PII?

Petit: “There are three main cryptographic libraries: MIRACL, OpenSSL and Crypto++.
Authors in [21] proposed a comparison and concluded that MIRACL has the best performance for operations on elliptic curves over binary fields.”

How did Petit manage to use 17 million cycles for P-224 verif, 22 million cycles for P-256 verif?

Presumably some combination of bad mulmod and bad curve ops.

Why did Petit reimplement ECDSA, using MIRACL for the underlying arithmetic?

Why did Petit not simply cite previous speed literature?

Why did Petit choose Pentium D?

Why did BHLM choose PII?

Petit: “There are three main cryptographic libraries: MIRACL, OpenSSL and Crypto++.

Authors in [21] proposed a comparison and concluded that MIRACL has the best performance for operations on elliptic curves over binary fields.”

But NIST P-224 and NIST P-256 are defined over prime fields!

[21] says “For elliptic curves over prime fields, OpenSSL has the best performance under all platforms.”

Petit manage to use
 on cycles for P-224 verif,
 on cycles for P-256 verif?

ably some combination of
 mod and bad curve ops.

l Petit reimplement
 using MIRACL for the
 ng arithmetic?

l Petit not simply cite
 speed literature?

l Petit choose Pentium D?

l BHLM choose PII?

Petit: “There are three main
 cryptographic libraries: MIRACL,
 OpenSSL and Crypto++.

Authors in [21] proposed a
 comparison and concluded
 that MIRACL has the best
 performance for operations on
 elliptic curves over binary fields.”

But NIST P-224 and NIST P-256
 are defined over prime fields!

[21] says “For elliptic curves
 over prime fields, OpenSSL has
 the best performance under all
 platforms.”

More ge
 Paper an
 crypto u

If the cr
 Why is t
 Why sho

If the cr
 Paper is
 Look, he
 More lik
 More lik
 funding

image to use
 or P-224 verif,
 or P-256 verif?

combination of
 bad curve ops.

implement
 MIRACL for the
 etic?

simply cite
 literature?

use Pentium D?

choose PII?

Petit: “There are three main
 cryptographic libraries: MIRACL,
 OpenSSL and Crypto++.

Authors in [21] proposed a
 comparison and concluded
 that MIRACL has the best
 performance for operations on
 elliptic curves over binary fields.”

But NIST P-224 and NIST P-256
 are defined over prime fields!

[21] says “For elliptic curves
 over prime fields, OpenSSL has
 the best performance under all
 platforms.”

More general situa
 Paper analyzes im
 crypto upon an ap

If the crypto soun
 Why is the paper
 Why should it be

If the crypto soun
 Paper is more inte
 Look, here’s a spe

More likely to be p
 More likely to mot
 funding to fix the

Petit: “There are three main cryptographic libraries: MIRACL, OpenSSL and Crypto++.

Authors in [21] proposed a comparison and concluded that MIRACL has the best performance for operations on elliptic curves over binary fields.”

But NIST P-224 and NIST P-256 are defined over prime fields!

[21] says “For elliptic curves over prime fields, OpenSSL has the best performance under all platforms.”

More general situation:
Paper analyzes impact of crypto upon an application.

If the crypto sounds fast:
Why is the paper interesting?
Why should it be published?

If the crypto sounds slower:
Paper is more interesting.
Look, here’s a speed problem!
More likely to be published.
More likely to motivate funding to fix the problem.

Petit: “There are three main cryptographic libraries: MIRACL, OpenSSL and Crypto++.

Authors in [21] proposed a comparison and concluded that MIRACL has the best performance for operations on elliptic curves over binary fields.”

But NIST P-224 and NIST P-256 are defined over prime fields!

[21] says “For elliptic curves over prime fields, OpenSSL has the best performance under all platforms.”

More general situation:
Paper analyzes impact of crypto upon an application.

If the crypto sounds fast:
Why is the paper interesting?
Why should it be published?

If the crypto sounds slower:
Paper is more interesting.
Look, here’s a speed problem!
More likely to be published.
More likely to motivate funding to fix the problem.

There are three main
 graphic libraries: MIRACL,
 L and Crypto++.

in [21] proposed a
 son and concluded
 RACL has the best
 ance for operations on
 curves over binary fields.”

T P-224 and NIST P-256
 ed over prime fields!

s “For elliptic curves
 me fields, OpenSSL has
 performance under all
 s.”

More general situation:
 Paper analyzes impact of
 crypto upon an application.

If the crypto sounds fast:
 Why is the paper interesting?
 Why should it be published?

If the crypto sounds slower:
 Paper is more interesting.
 Look, here’s a speed problem!
 More likely to be published.
 More likely to motivate
 funding to fix the problem.

Obvious
 applicati
 deploym

Many ra
 answerin
 CPU to
 literature
 mulmod

Slowest,
 are most

Situation
 random
 There’s
deliberat

three main
aries: MIRACL,
pto++.

posed a
oncluded
the best
operations on
binary fields.”

and NIST P-256
prime fields!

otic curves
OpenSSL has
nce under all

More general situation:
Paper analyzes impact of
crypto upon an application.

If the crypto sounds fast:
Why is the paper interesting?
Why should it be published?

If the crypto sounds slower:
Paper is more interesting.
Look, here's a speed problem!
More likely to be published.
More likely to motivate
funding to fix the problem.

Obvious question v
application consid
deployment: “Is it

Many random met
answering this que
CPU to test? Wha
literature and libra
mulmod, or curve

Slowest, least com
are most likely to

Situation is fully e
randomness + nat
There's no evidenc
deliberately slowe

More general situation:
 Paper analyzes impact of
 crypto upon an application.

If the crypto sounds fast:
 Why is the paper interesting?
 Why should it be published?

If the crypto sounds slower:
 Paper is more interesting.
 Look, here's a speed problem!
 More likely to be published.
 More likely to motivate
 funding to fix the problem.

Obvious question whenever
 application considers crypto
 deployment: "Is it fast enough?"

Many random methodologies
 answering this question. What
 CPU to test? What to take
 literature and libraries? Reu-
 mulmod, or curve ops, or mo-

Slowest, least competent and
 are most likely to be published.

Situation is fully explainable
 randomness + natural selection.
 There's no evidence that Pe-
deliberately slowed down cry-

More general situation:

Paper analyzes impact of crypto upon an application.

If the crypto sounds fast:

Why is the paper interesting?

Why should it be published?

If the crypto sounds slower:

Paper is more interesting.

Look, here's a speed problem!

More likely to be published.

More likely to motivate funding to fix the problem.

Obvious question whenever an application considers crypto deployment: "Is it fast enough?"

Many random methodologies for answering this question. Which CPU to test? What to take from literature and libraries? Reuse mulmod, or curve ops, or more?

Slowest, least competent answers are most likely to be published.

Situation is fully explainable by randomness + natural selection.

There's no evidence that Petit *deliberately* slowed down crypto.

General situation:
 analyzes impact of
 upon an application.
 crypto sounds fast:
 the paper interesting?
 could it be published?
 crypto sounds slower:
 more interesting.
 there's a speed problem!
 likely to be published.
 likely to motivate
 to fix the problem.

Obvious question whenever an
 application considers crypto
 deployment: "Is it fast enough?"
 Many random methodologies for
 answering this question. Which
 CPU to test? What to take from
 literature and libraries? Reuse
 mulmod, or curve ops, or more?
 Slowest, least competent answers
 are most likely to be published.
 Situation is fully explainable by
 randomness + natural selection.
 There's no evidence that Petit
deliberately slowed down crypto.

Paper in
 software
 incentive
 slow, an
 report it
 Paper w
 function
 lengths,
 timing m
 maximiz
 from old
 This is r
 what ma

ation:
 pact of
 application.
 ds fast:
 interesting?
 published?
 ds slower:
 resting.
 ed problem!
 ublished.
 ivate
 problem.

Obvious question whenever an application considers crypto deployment: “Is it fast enough?”

Many random methodologies for answering this question. Which CPU to test? What to take from literature and libraries? Reuse mulmod, or curve ops, or more?

Slowest, least competent answers are most likely to be published.

Situation is fully explainable by randomness + natural selection. There’s no evidence that Petit *deliberately* slowed down crypto.

Paper introducing software or hardware incentive to report slow, and analogous report its own crypto. Paper will naturally functions, parameters, lengths, platforms, timing mechanism maximize reported from old to new. This is not the same what matters most

Obvious question whenever an application considers crypto deployment: “Is it fast enough?”

Many random methodologies for answering this question. Which CPU to test? What to take from literature and libraries? Reuse mulmod, or curve ops, or more?

Slowest, least competent answers are most likely to be published.

Situation is fully explainable by randomness + natural selection. There’s no evidence that Petit *deliberately* slowed down crypto.

Paper introducing new crypto software or hardware has same incentive to report older crypto as slow, and analogous incentive to report its own crypto as fast.

Paper will naturally select functions, parameters, input lengths, platforms, I/O forms, timing mechanism, etc. that maximize reported improvement from old to new.

This is not the same as selecting what matters most for the user.

Obvious question whenever an application considers crypto deployment: “Is it fast enough?”

Many random methodologies for answering this question. Which CPU to test? What to take from literature and libraries? Reuse mulmod, or curve ops, or more?

Slowest, least competent answers are most likely to be published.

Situation is fully explainable by randomness + natural selection. There’s no evidence that Petit *deliberately* slowed down crypto.

Paper introducing new crypto software or hardware has same incentive to report older crypto as slow, and analogous incentive to report its own crypto as fast.

Paper will naturally select functions, parameters, input lengths, platforms, I/O format, timing mechanism, etc. that maximize reported improvement from old to new.

This is not the same as selecting what matters most for the users.

question whenever an
 on considers crypto
 ent: “Is it fast enough?”
 ndom methodologies for
 g this question. Which
 test? What to take from
 e and libraries? Reuse
 , or curve ops, or more?
 least competent answers
 t likely to be published.
 n is fully explainable by
 ness + natural selection.
 no evidence that Petit
 tely slowed down crypto.

Paper introducing new crypto
 software or hardware has same
 incentive to report older crypto as
 slow, and analogous incentive to
 report its own crypto as fast.

Paper will naturally select
 functions, parameters, input
 lengths, platforms, I/O format,
 timing mechanism, etc. that
 maximize reported improvement
 from old to new.

This is not the same as selecting
 what matters most for the users.

Bit oper
 (assumin
 as listed

key	ops
-----	-----

128	88
-----	----

128	100
-----	-----

128	117
-----	-----

256	144
-----	-----

128	147
-----	-----

256	156
-----	-----

128	162
-----	-----

128	202
-----	-----

256	283
-----	-----

whenever an
 ers crypto
 fast enough?"
 methodologies for
 estion. Which
 at to take from
 ries? Reuse
 ops, or more?
 petent answers
 be published.
 explainable by
 tural selection.
 ce that Petit
 d down crypto.

Paper introducing new crypto
 software or hardware has same
 incentive to report older crypto as
 slow, and analogous incentive to
 report its own crypto as fast.

Paper will naturally select
 functions, parameters, input
 lengths, platforms, I/O format,
 timing mechanism, etc. that
 maximize reported improvement
 from old to new.

This is not the same as selecting
 what matters most for the users.

Bit operations per
 (assuming precom
 as listed in recent

key	ops/bit	ciph
128	88	Simc
128	100	NOE
128	117	Skin
256	144	Simc
128	147.2	PRE
256	156	Skin
128	162.75	Picc
128	202.5	AES
256	283.5	AES

Paper introducing new crypto software or hardware has same incentive to report older crypto as slow, and analogous incentive to report its own crypto as fast.

Paper will naturally select functions, parameters, input lengths, platforms, I/O format, timing mechanism, etc. that maximize reported improvement from old to new.

This is not the same as selecting what matters most for the users.

Bit operations per bit of plaintext (assuming precomputed subkeys) as listed in recent Skinny paper

key	ops/bit	cipher
128	88	Simon: 60 ops
128	100	NOEKEON
128	117	Skinny
256	144	Simon: 106 ops
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Paper introducing new crypto software or hardware has same incentive to report older crypto as slow, and analogous incentive to report its own crypto as fast.

Paper will naturally select functions, parameters, input lengths, platforms, I/O format, timing mechanism, etc. that maximize reported improvement from old to new.

This is not the same as selecting what matters most for the users.

Bit operations per bit of plaintext (assuming precomputed subkeys), as listed in recent Skinny paper:

key	ops/bit	cipher
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Paper introducing new crypto software or hardware has same incentive to report older crypto as slow, and analogous incentive to report its own crypto as fast.

Paper will naturally select functions, parameters, input lengths, platforms, I/O format, timing mechanism, etc. that maximize reported improvement from old to new.

This is not the same as selecting what matters most for the users.

Bit operations per bit of plaintext (assuming precomputed subkeys), not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Introducing new crypto
 or hardware has same
 incentive to report older crypto as
 slow and analogous incentive to
 report own crypto as fast.

Will naturally select
 parameters, input
 platforms, I/O format,
 mechanism, etc. that
 are reported improvement
 relative to new.

Not the same as selecting
 matters most for the users.

Bit operations per bit of plaintext
 (assuming precomputed subkeys),
 not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Many ba
 backed b

e.g. Do
 optimize
 the olde
 Rely on

“We cor
 most arc
 do much
 complet
 heuristic
 get little
 where th
 slightly v

new crypto
are has same
older crypto as
us incentive to
to as fast.

y select
ters, input
, I/O format,
, etc. that
improvement

me as selecting
t for the users.

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Many bad examples
backed by tons of
e.g. Do we bother
optimized impleme
the older crypto?

Rely on “optimizin

“We come so close
most architectures
do much more with
complete algorithm
heuristics. We can
get little niggles h
where the heuristic
slightly wrong ans

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Many bad examples to imitate
backed by tons of misinformation
e.g. Do we bother searching for
optimized implementations of
the older crypto? Take any
Rely on “optimizing” compilers

“We come so close to optimizing
most architectures that we can
do much more without using
complete algorithms instead of
heuristics. We can only try to
get little niggles here and there
where the heuristics get
slightly wrong answers.”

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Many bad examples to imitate,
backed by tons of misinformation.

e.g. Do we bother searching for
optimized implementations of
the older crypto? Take any code!
Rely on “optimizing” compiler!

“We come so close to optimal on
most architectures that we can’t
do much more without using NP
complete algorithms instead of
heuristics. We can only try to
get little niggles here and there
where the heuristics get
slightly wrong answers.”

ations per bit of plaintext
 ng precomputed subkeys),
 rely listed in Skinny paper:

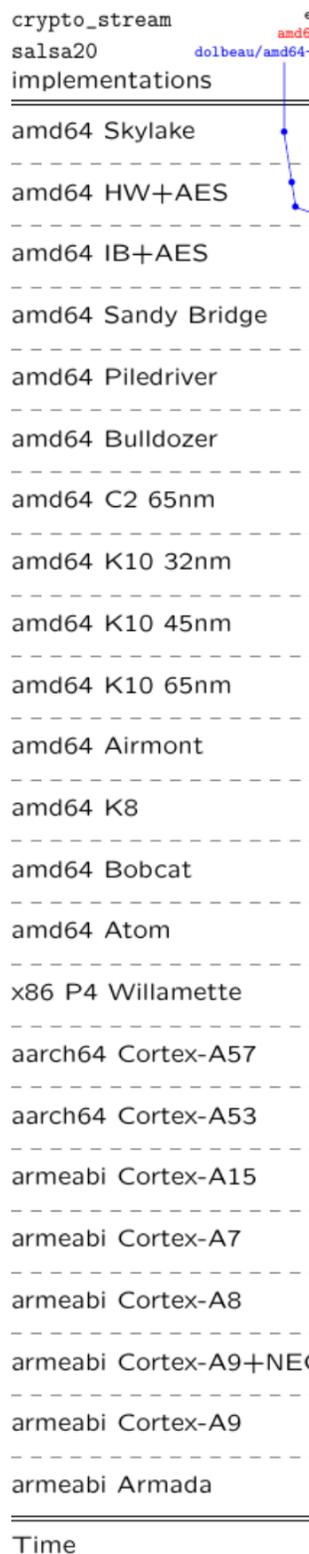
ops/bit	cipher
4	Salsa20/8
3	Salsa20/12
3	Simon: 60 ops broken
0	NOEKEON
7	Skinny
5	Salsa20
4	Simon: 106 ops broken
7.2	PRESENT
5	Skinny
2.75	Piccolo
2.5	AES
3.5	AES

Many bad examples to imitate,
 backed by tons of misinformation.

e.g. Do we bother searching for
 optimized implementations of
 the older crypto? Take any code!
 Rely on “optimizing” compiler!

“We come so close to optimal on
 most architectures that we can’t
 do much more without using NP
 complete algorithms instead of
 heuristics. We can only try to
 get little niggles here and there
 where the heuristics get
 slightly wrong answers.”

Reality i



bit of plaintext
(computed subkeys),
in Skinny paper:

er

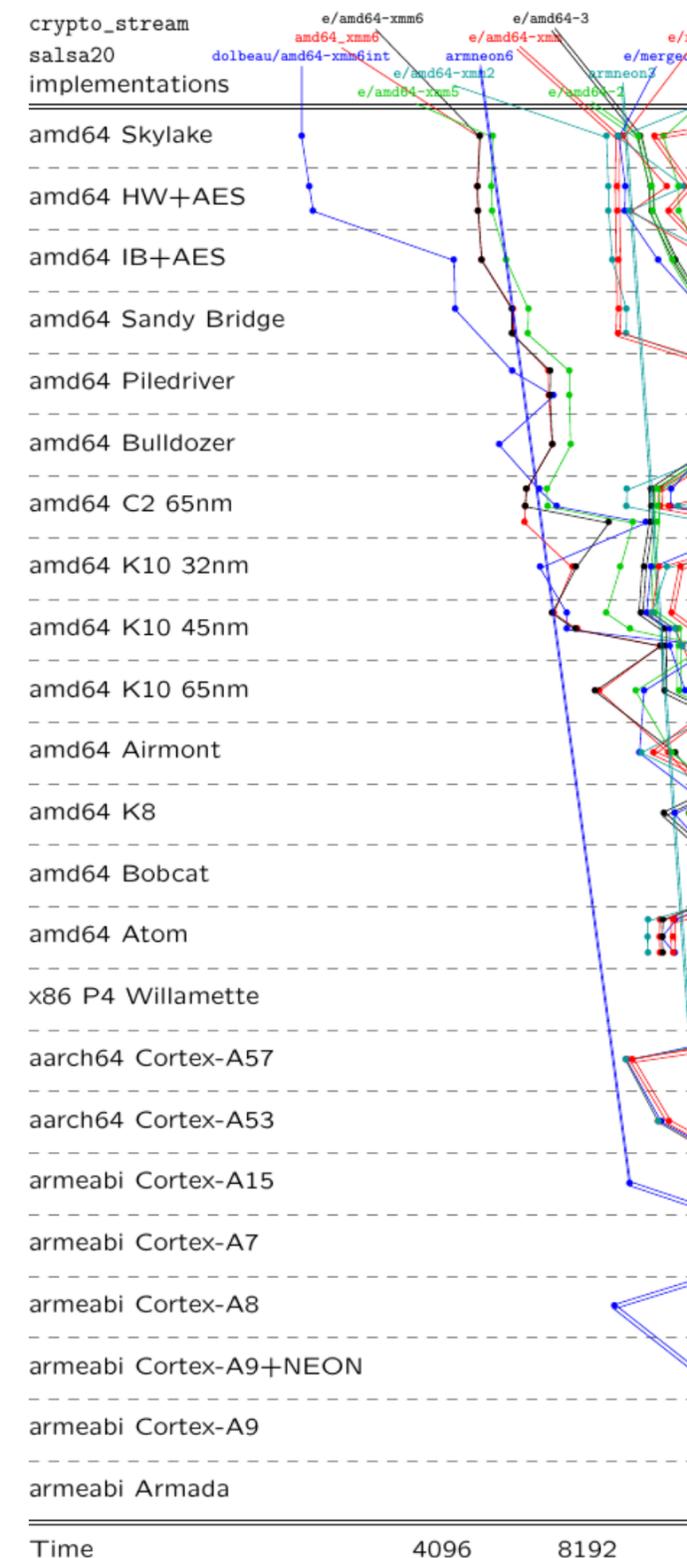
a20/8
a20/12
on: 60 ops broken
EKEON
ny
a20
on: 106 ops broken
SENT
ny
olo

Many bad examples to imitate,
backed by tons of misinformation.

e.g. Do we bother searching for
optimized implementations of
the older crypto? Take any code!
Rely on “optimizing” compiler!

“We come so close to optimal on
most architectures that we can’t
do much more without using NP
complete algorithms instead of
heuristics. We can only try to
get little niggles here and there
where the heuristics get
slightly wrong answers.”

Reality is more complex



intext
(keys),
paper:

Many bad examples to imitate,
backed by tons of misinformation.

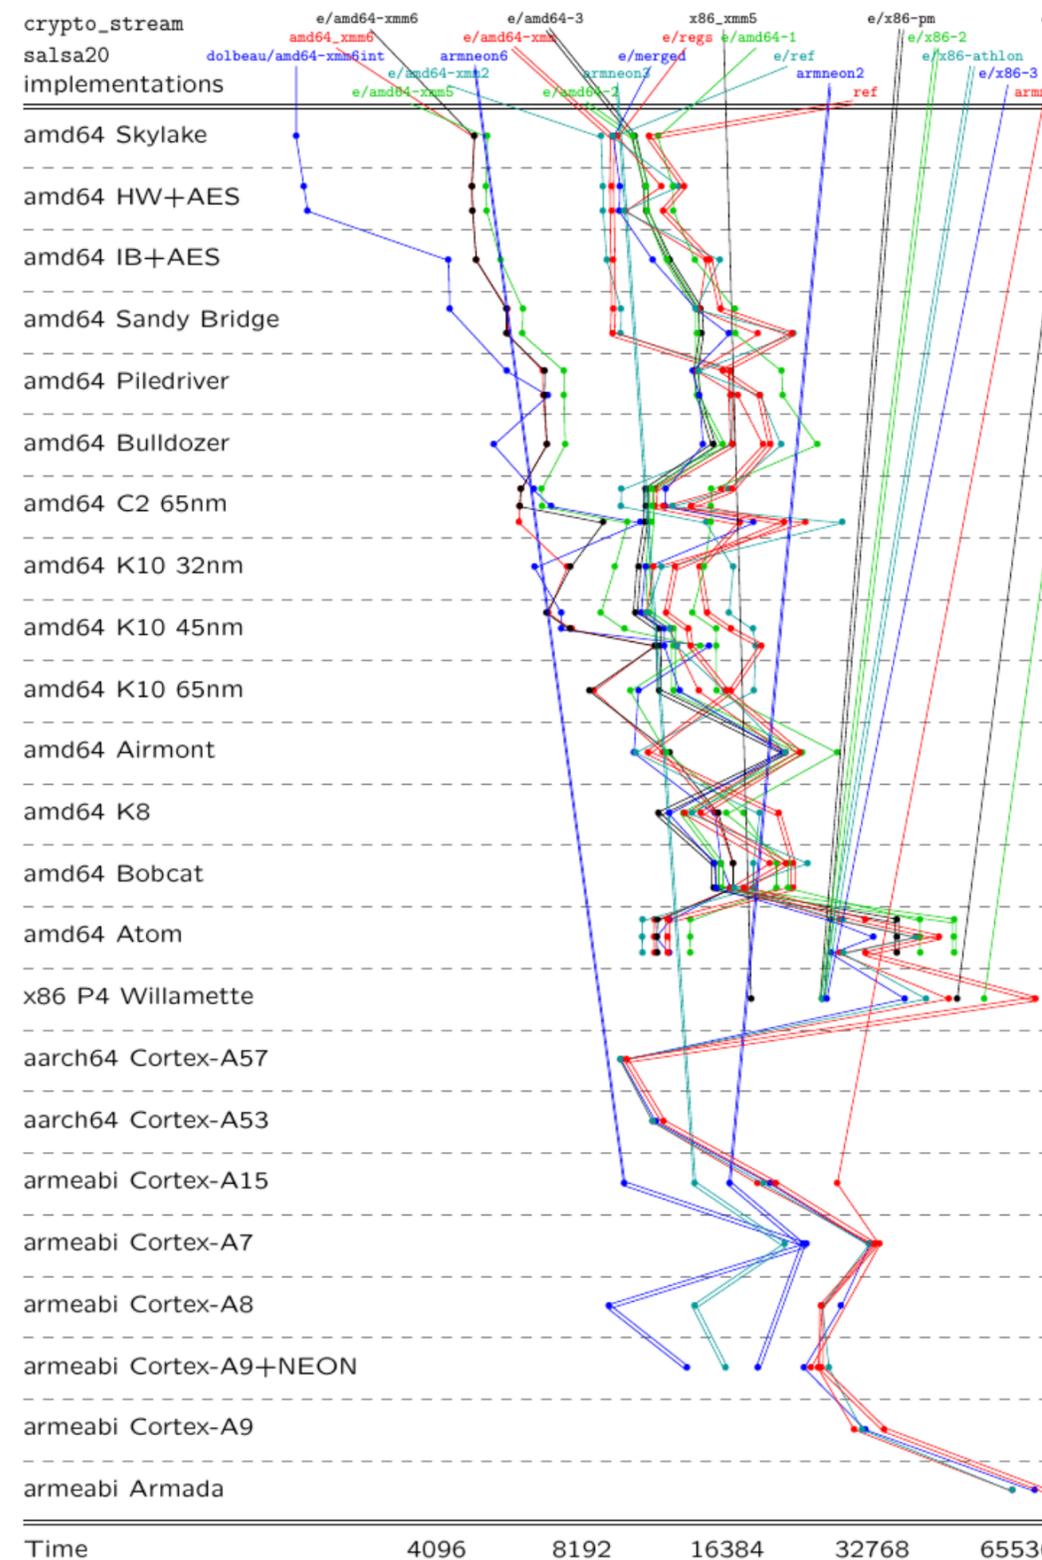
e.g. Do we bother searching for
optimized implementations of
the older crypto? Take any code!
Rely on “optimizing” compiler!

“We come so close to optimal on
most architectures that we can’t
do much more without using NP
complete algorithms instead of
heuristics. We can only try to
get little niggles here and there
where the heuristics get
slightly wrong answers.”

broken

os broken

Reality is more complicated:

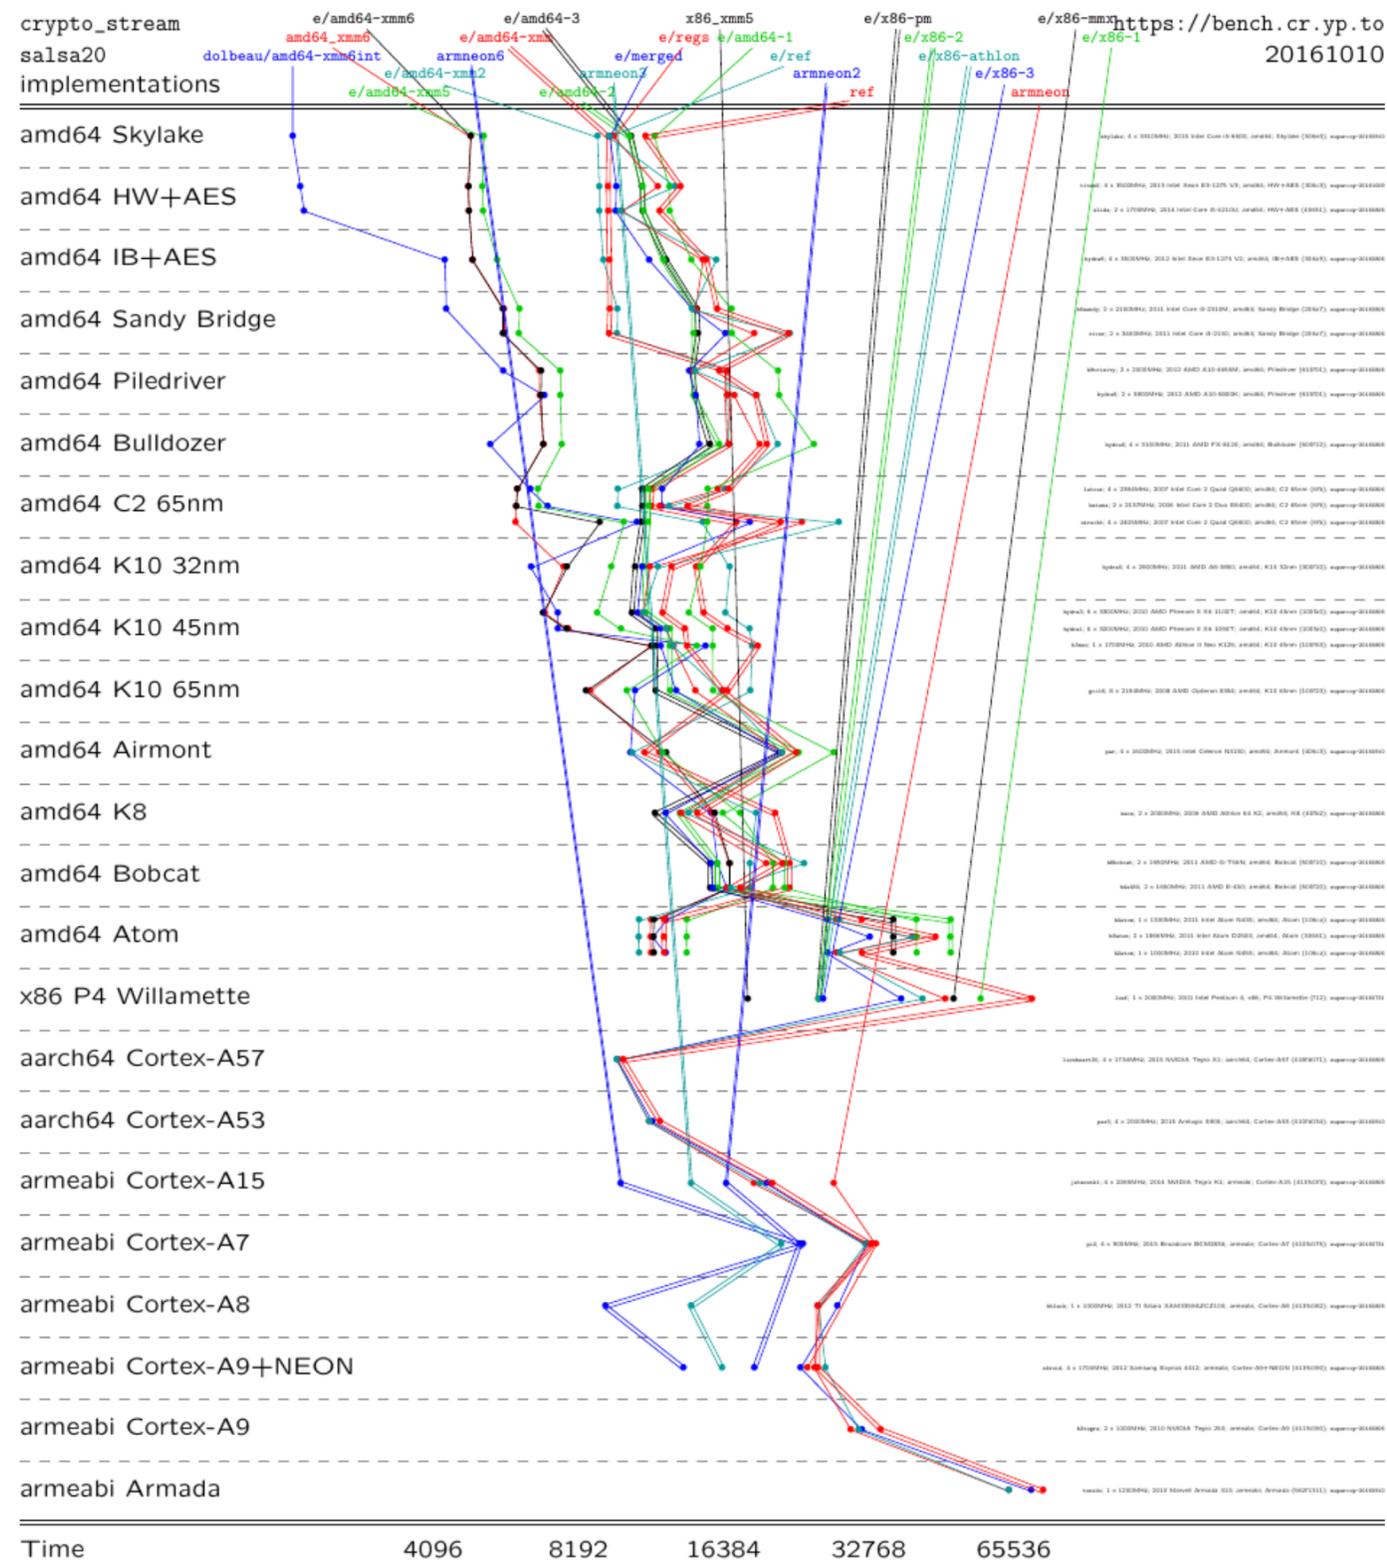


Many bad examples to imitate,
backed by tons of misinformation.

e.g. Do we bother searching for
optimized implementations of
the older crypto? Take any code!
Rely on “optimizing” compiler!

“We come so close to optimal on
most architectures that we can’t
do much more without using NP
complete algorithms instead of
heuristics. We can only try to
get little niggles here and there
where the heuristics get
slightly wrong answers.”

Reality is more complicated:

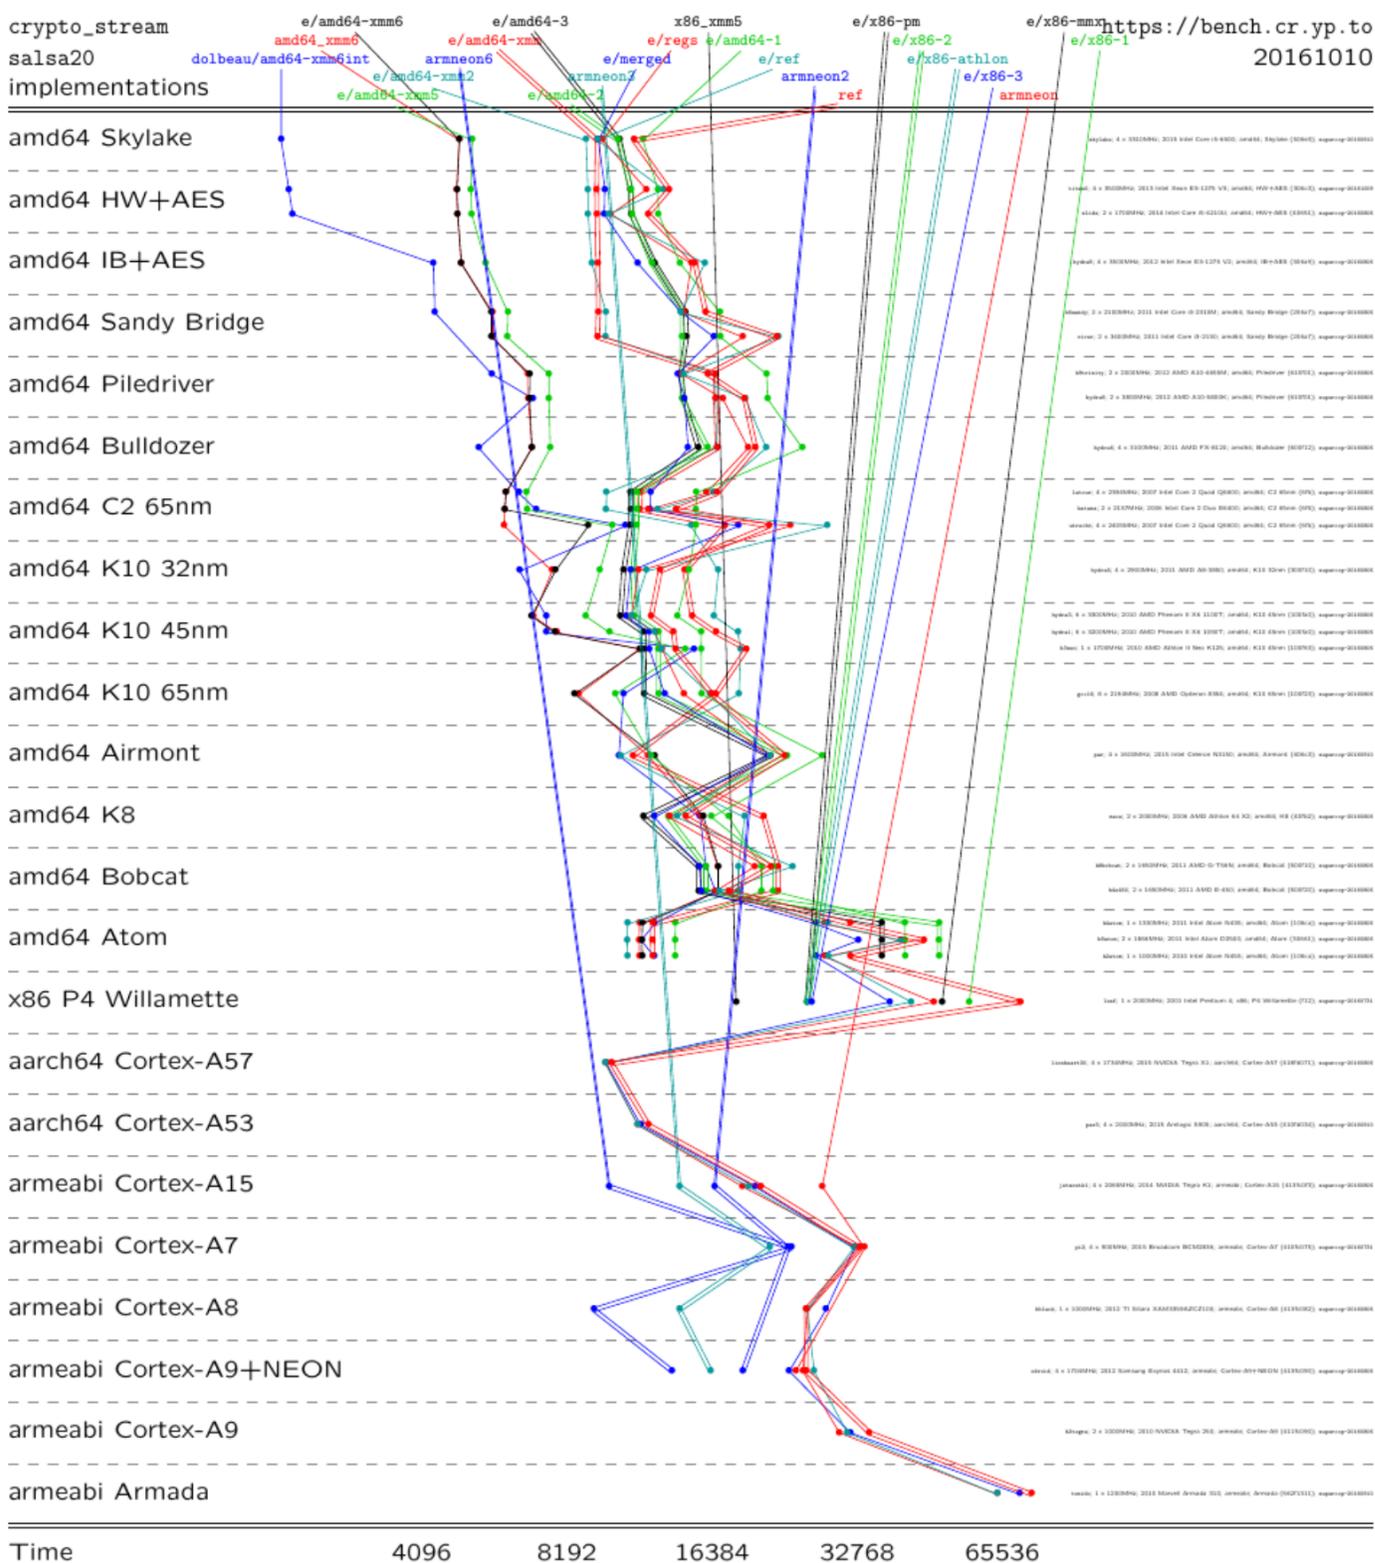


ad examples to imitate,
by tons of misinformation.

we bother searching for
ed implementations of
r crypto? Take any code!
“optimizing” compiler!

me so close to optimal on
chitectures that we can't
n more without using NP
e algorithms instead of
s. We can only try to
e niggles here and there
ne heuristics get
wrong answers.”

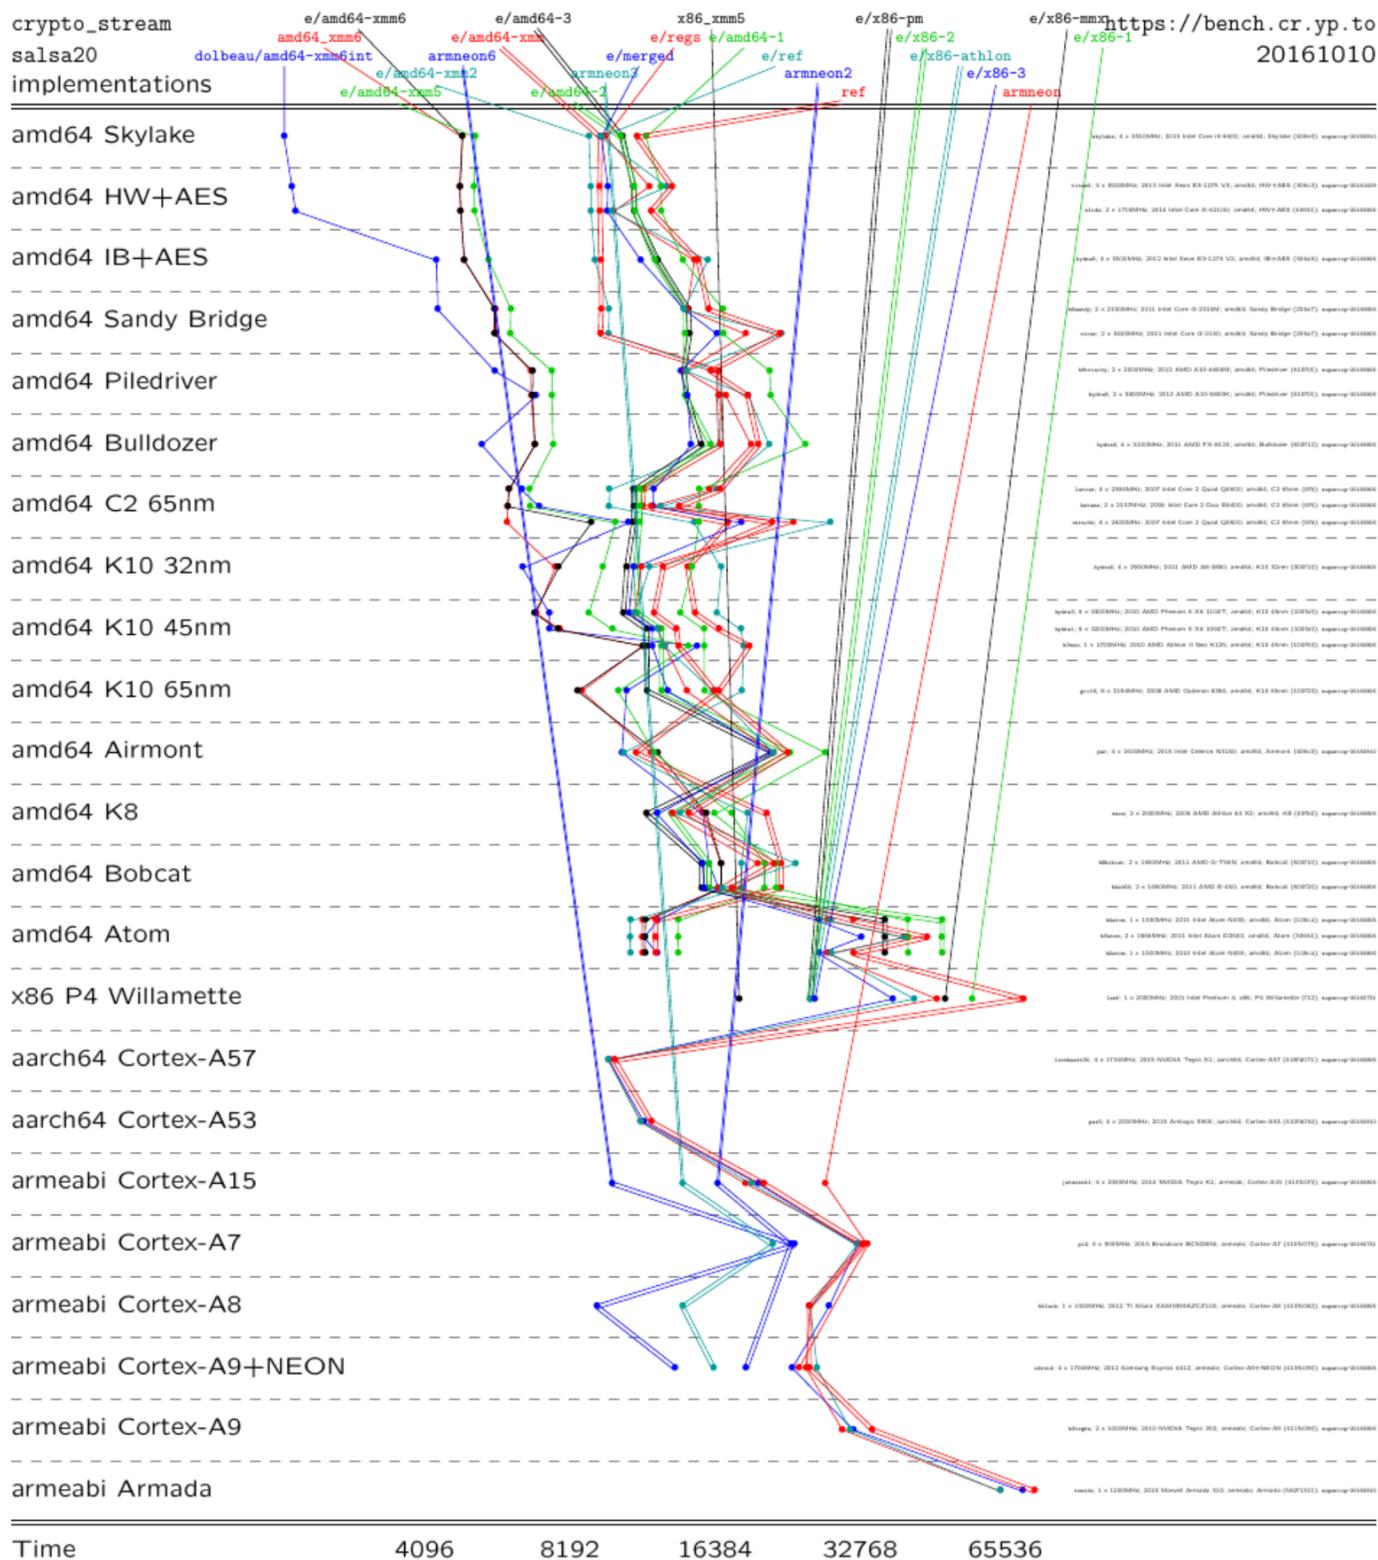
Reality is more complicated:



SUPERCO
includes
of 595 c
>20 imp
Haswell:
impleme
gcc -O3
is 6.15x
Salsa20
merged
with “m
optimiza
compiler

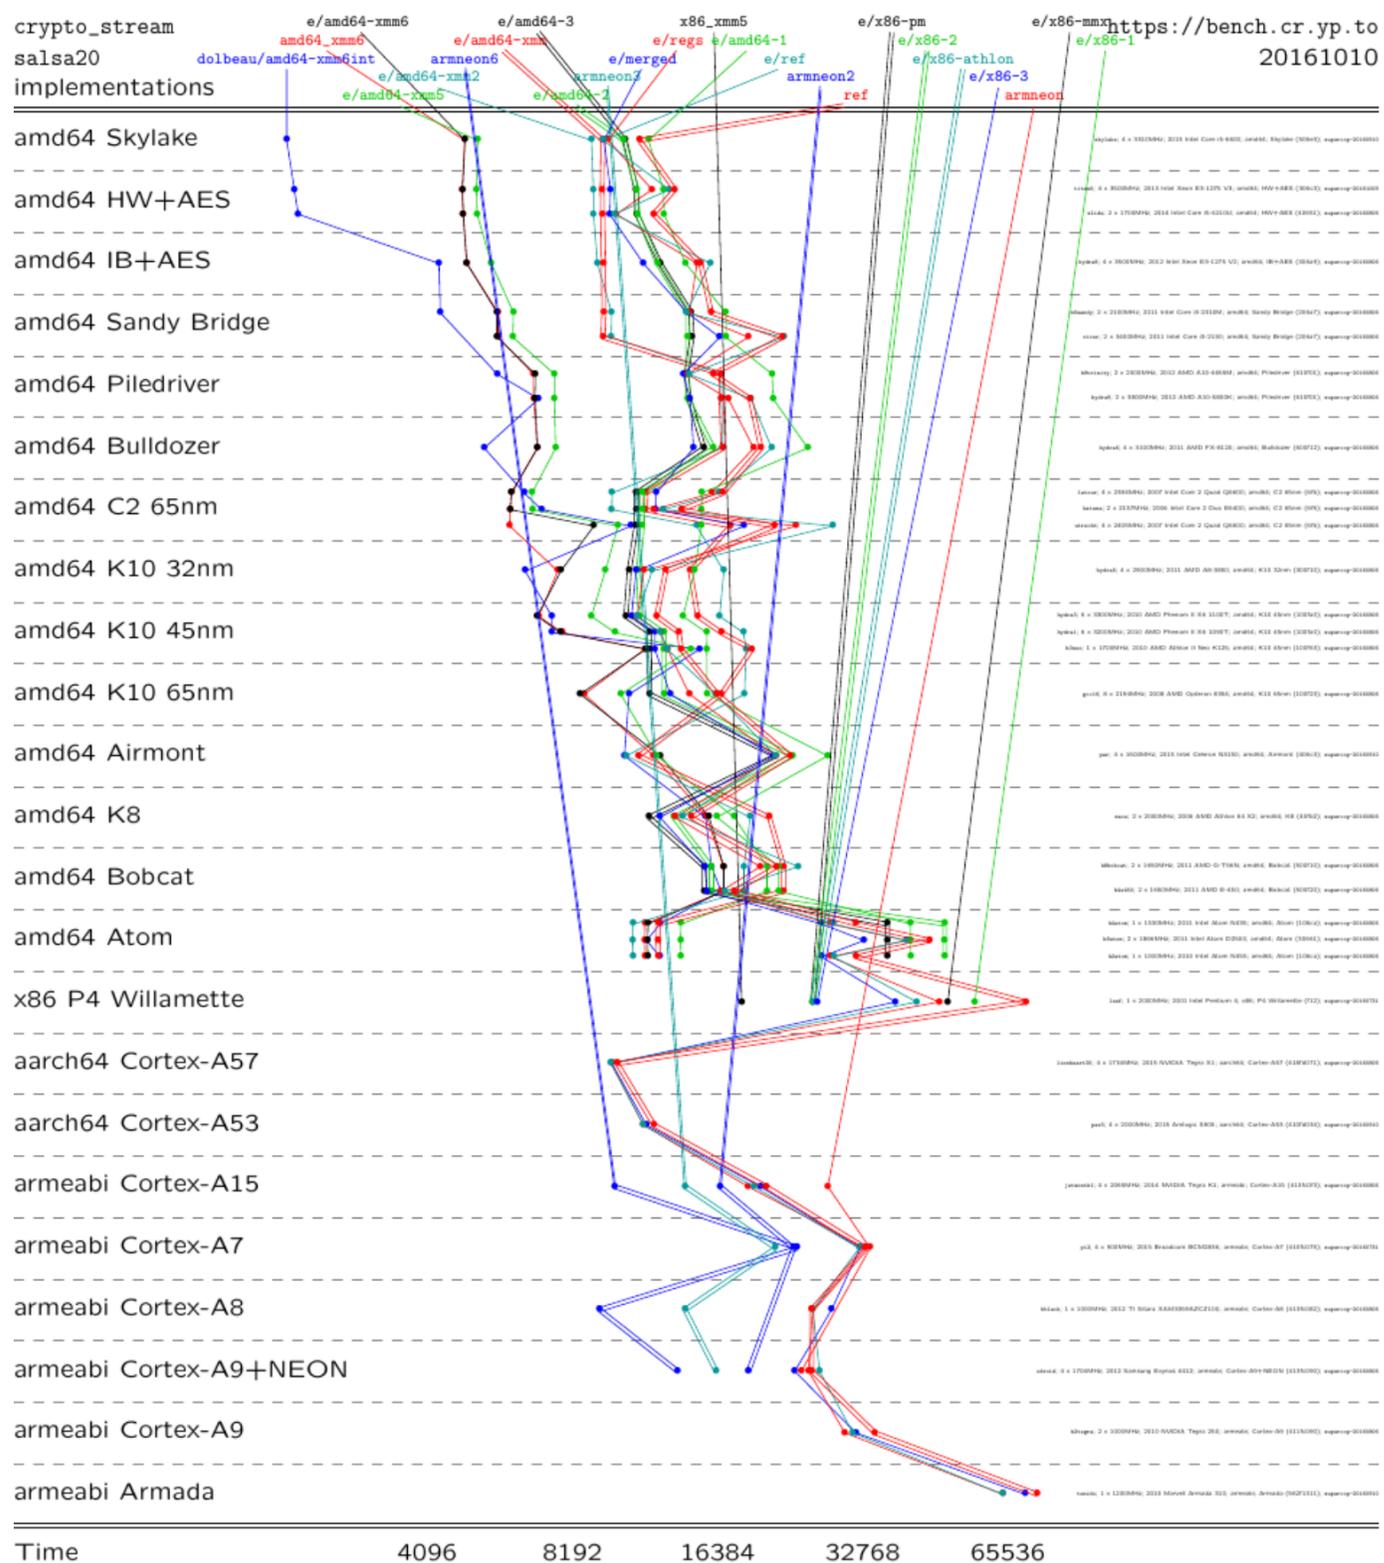
ate,
ation.
for
of
code!
ler!
al on
can't
g NP
of
to
ere

Reality is more complicated:



SUPERCOP benchmarking tool includes 2155 implementations of 595 cryptographic primitives. >20 implementations of Salsa20. Haswell: Reasonably simple implementation compiled with gcc -O3 -fomit-frame-pointer is 6.15x slower than fastest Salsa20 implementation. merged implementation with "machine-independent" optimizations and best of 12 compiler options: 4.52x slower

Reality is more complicated:

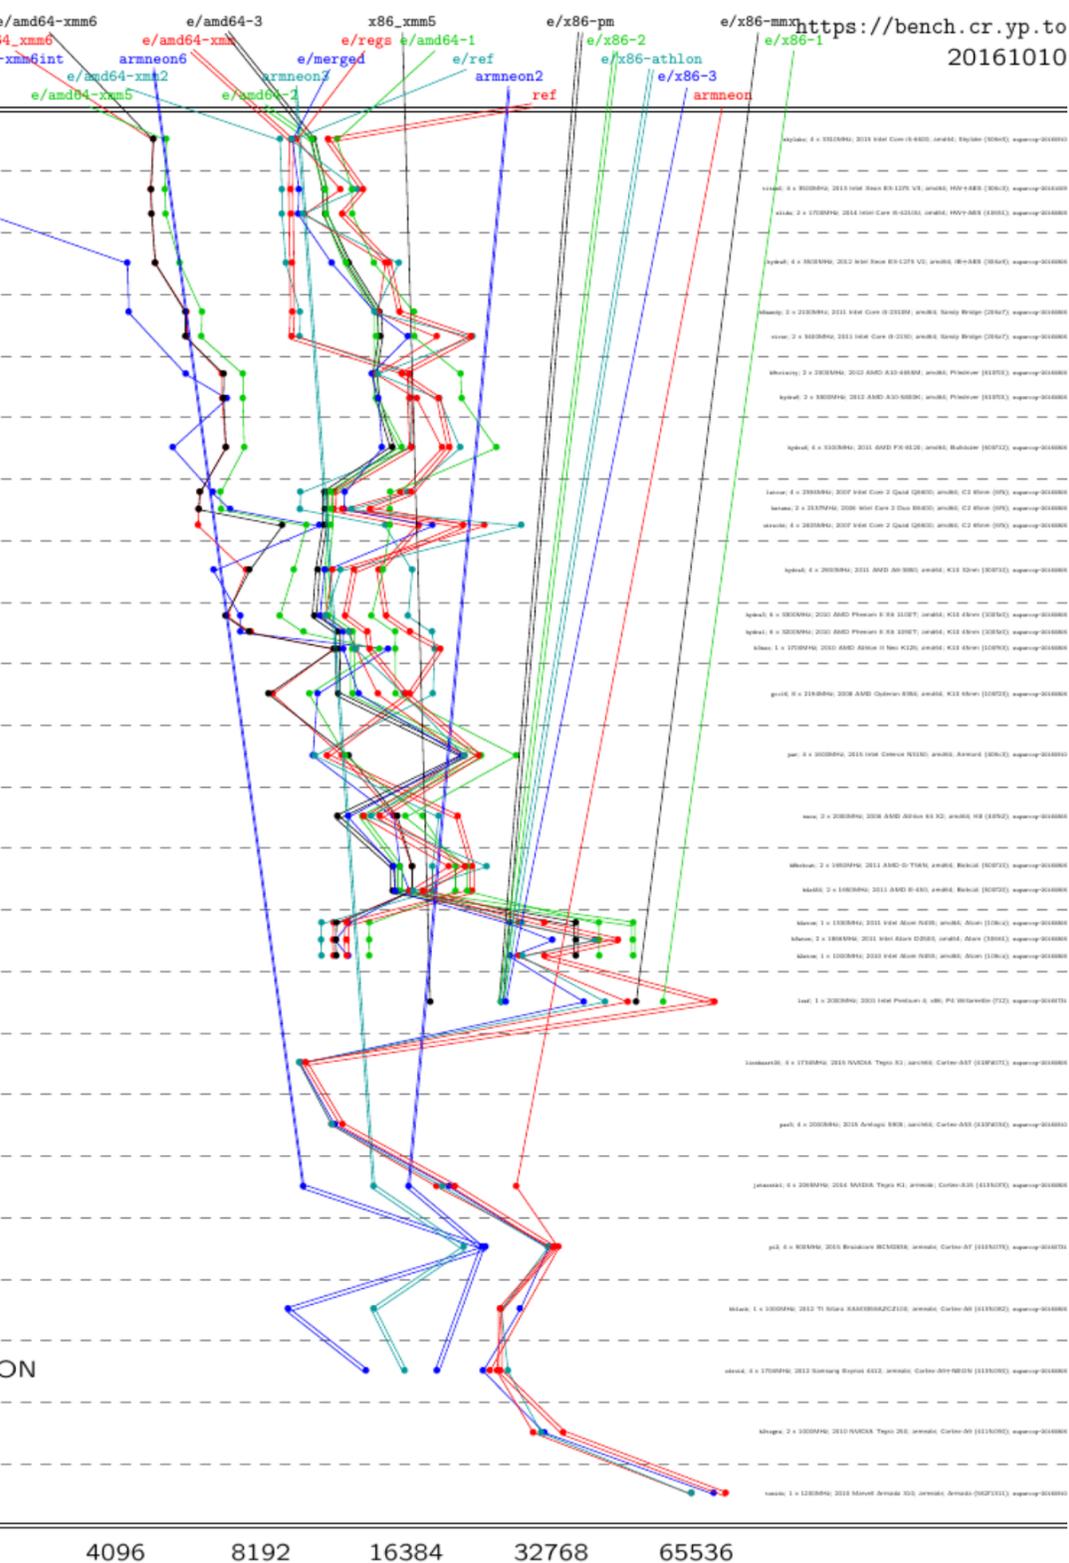


SUPERCOP benchmarking toolkit includes 2155 implementations of 595 cryptographic primitives. >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

s more complicated:



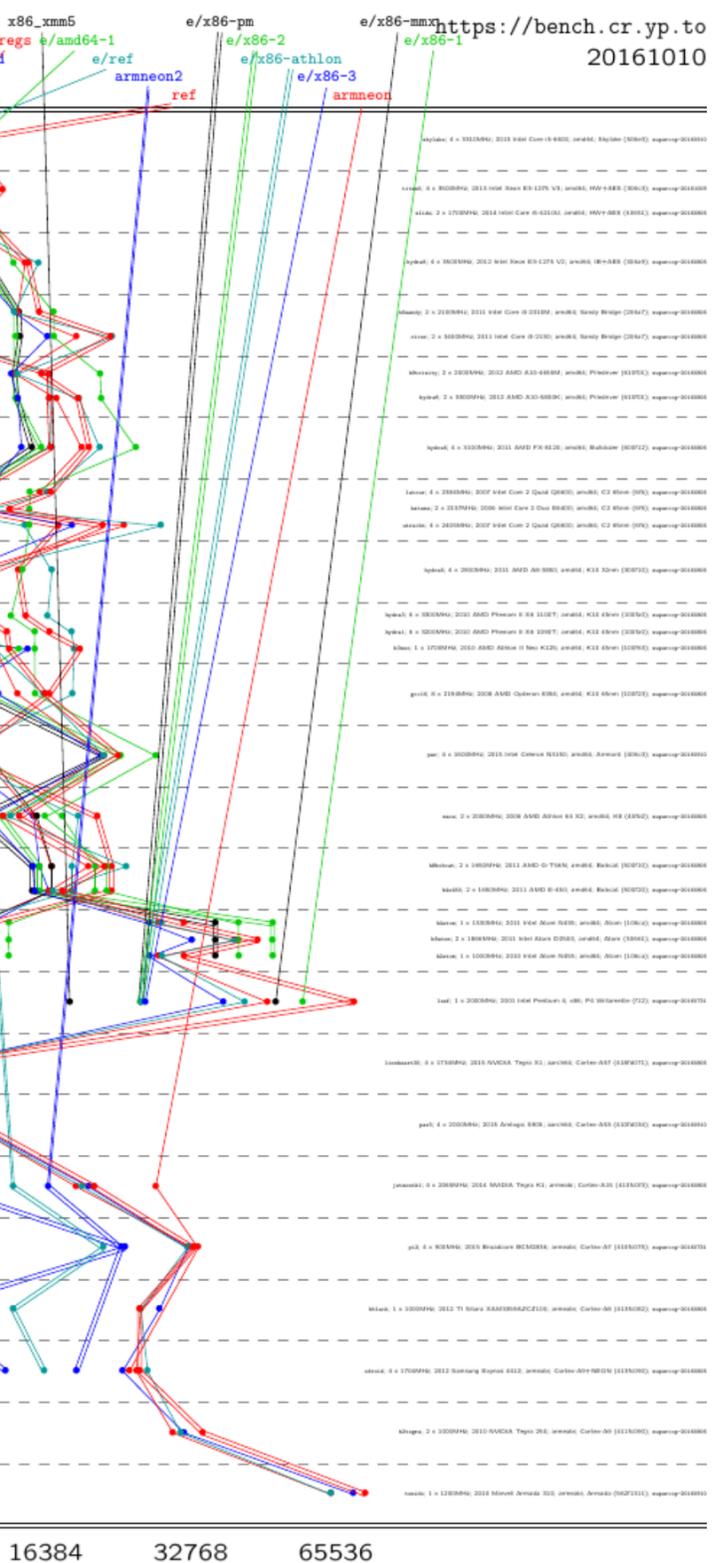
SUPERCOP benchmarking toolkit includes 2155 implementations of 595 cryptographic primitives. >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Another lattice-based means generation of random 2017.03 Valencia Regazzo sources of discrete benchmarking. Qualitative choice of sampling

complicated:



SUPERCOP benchmarking toolkit includes 2155 implementations of 595 cryptographic primitives. >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Another interesting lattice-based signing means generating of random Gaussian 2017.03 Brannigan Valencia–O’Sullivan Regazzoni “An inv sources of random discrete Gaussian benchmarks for RM Qualitatively large choice of RNG \Rightarrow sampling \Rightarrow cost c

SUPERCOP benchmarking toolkit includes 2155 implementations of 595 cryptographic primitives.
 >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Another interesting example: lattice-based signing typically means generating a huge number of random Gaussian samples.

2017.03 Brannigan–Smyth–Oder–Valencia–O’Sullivan–Güneysu–Regazzoni “An investigation of sources of randomness within discrete Gaussian sampling”: benchmarks for RNGs, samplers.

Qualitatively large impacts:
 choice of RNG \Rightarrow cost of sampling \Rightarrow cost of signing.

COP benchmarking toolkit
2155 implementations
cryptographic primitives.

Implementations of Salsa20.

Reasonably simple ref
implementation compiled with
-fomit-frame-pointer
slower than fastest
implementation.

implementation
"machine-independent"

options and best of 121

options: $4.52\times$ slower.

Another interesting example:
lattice-based signing typically
means generating a huge number
of random Gaussian samples.

2017.03 Brannigan–Smyth–Oder–
Valencia–O’Sullivan–Güneysu–
Regazzoni “An investigation of
sources of randomness within
discrete Gaussian sampling” :
benchmarks for RNGs, samplers.

Qualitatively large impacts:
choice of RNG \Rightarrow cost of
sampling \Rightarrow cost of signing.

Two exa
in this 2
Skylake

383.69 M
cycles/b
using AE
(32 cycle

benchmarking toolkit
 implementations
 basic primitives.
 versions of Salsa20.
 only simple ref
 compiled with
 frame-pointer
 an fastest
 tation.
 tation
 dependent”
 best of 121
 4.52× slower.

Another interesting example:
 lattice-based signing typically
 means generating a huge number
 of random Gaussian samples.

2017.03 Brannigan–Smyth–Oder–
 Valencia–O’Sullivan–Güneysu–
 Regazzoni “An investigation of
 sources of randomness within
 discrete Gaussian sampling”:
 benchmarks for RNGs, samplers.

Qualitatively large impacts:
 choice of RNG \Rightarrow cost of
 sampling \Rightarrow cost of signing.

Two examples of s
 in this 2017 paper
 Skylake (Intel Core
 383.69 MByte/sec
 cycles/byte) for A
 using AES-NI; 106
 (32 cycles/byte) fo

toolkit
ons
ves.
sa20.
ref
th
inter

21
wer.

Another interesting example:
lattice-based signing typically
means generating a huge number
of random Gaussian samples.

2017.03 Brannigan–Smyth–Oder–
Valencia–O’Sullivan–Güneysu–
Regazzoni “An investigation of
sources of randomness within
discrete Gaussian sampling” :
benchmarks for RNGs, samplers.

Qualitatively large impacts:
choice of RNG \Rightarrow cost of
sampling \Rightarrow cost of signing.

Two examples of speed reports
in this 2017 paper for a 3.4GHz
Skylake (Intel Core i7-6700):
383.69 MByte/sec (8.86
cycles/byte) for AES CTR-DRBG
using AES-NI; 106.07 MByte/sec
(32 cycles/byte) for ChaCha

Another interesting example:
lattice-based signing typically
means generating a huge number
of random Gaussian samples.

2017.03 Brannigan–Smyth–Oder–
Valencia–O’Sullivan–Güneysu–
Regazzoni “An investigation of
sources of randomness within
discrete Gaussian sampling” :
benchmarks for RNGs, samplers.

Qualitatively large impacts:
choice of RNG \Rightarrow cost of
sampling \Rightarrow cost of signing.

Two examples of speed reported
in this 2017 paper for a 3.4GHz
Skylake (Intel Core i7-6700):

383.69 MByte/sec (8.86
cycles/byte) for AES CTR-DRBG
using AES-NI; 106.07 MByte/sec
(32 cycles/byte) for ChaCha20.

Another interesting example: lattice-based signing typically means generating a huge number of random Gaussian samples.

2017.03 Brannigan–Smyth–Oder–Valencia–O’Sullivan–Güneysu–Regazzoni “An investigation of sources of randomness within discrete Gaussian sampling”: benchmarks for RNGs, samplers.

Qualitatively large impacts: choice of RNG \Rightarrow cost of sampling \Rightarrow cost of signing.

Two examples of speed reported in this 2017 paper for a 3.4GHz Skylake (Intel Core i7-6700):

383.69 MByte/sec (8.86 cycles/byte) for AES CTR-DRBG using AES-NI; 106.07 MByte/sec (32 cycles/byte) for ChaCha20.

But wait. eBACS reports 0.92 cycles/byte for AES-256-CTR, 1.18 cycles/byte for ChaCha20.

Author non-response: “essential for us to examine standard open implementations”. Slow ones?

interesting example:
 based signing typically
 generating a huge number
 of Gaussian samples.

Brannigan–Smyth–Oder–
 O’Sullivan–Güneysu–
 ni “An investigation of
 of randomness within
 Gaussian sampling”:
 marks for RNGs, samplers.

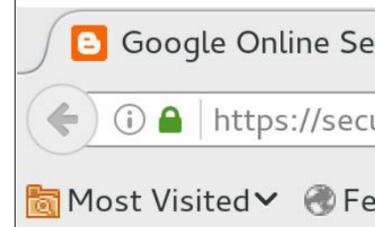
ively large impacts:
 of RNG \Rightarrow cost of
 g \Rightarrow cost of signing.

Two examples of speed reported
 in this 2017 paper for a 3.4GHz
 Skylake (Intel Core i7-6700):

383.69 MByte/sec (8.86
 cycles/byte) for AES CTR-DRBG
 using AES-NI; 106.07 MByte/sec
 (32 cycles/byte) for ChaCha20.

But wait. eBACS reports
 0.92 cycles/byte for AES-256-CTR,
 1.18 cycles/byte for ChaCha20.

Author non-response: “essential
 for us to examine standard open
 implementations”. Slow ones?



Speeding up
 connection
 April 24, 2014

Posted by Elie Burs

Earlier this year,
 operates three t
 AES hardware a
 devices such as
 experience, redu
 amount of time

To make this ha
 began implemen
 encryption and
 March 2013. It v
 abstraction laye

g example:
ng typically
a huge number
an samples.

n–Smyth–Oder–
n–Güneysu–
vestigation of
ness within
sampling” :
NGs, samplers.

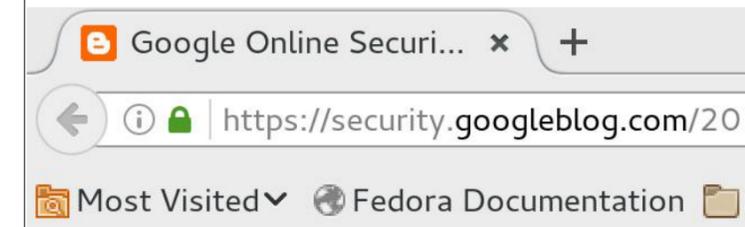
impacts:
cost of
of signing.

Two examples of speed reported
in this 2017 paper for a 3.4GHz
Skylake (Intel Core i7-6700):

383.69 MByte/sec (8.86
cycles/byte) for AES CTR-DRBG
using AES-NI; 106.07 MByte/sec
(32 cycles/byte) for ChaCha20.

But wait. eBACS reports
0.92 cycles/byte for AES-256-CTR,
1.18 cycles/byte for ChaCha20.

Author non-response: “essential
for us to examine standard open
implementations”. Slow ones?



Speeding up and strengthening
connections for Chrome on
April 24, 2014

Posted by Elie Bursztein, Anti-Abuse Researcher

Earlier this year, we deployed a new TLS implementation that
operates three times faster than AES-CTR. This is a significant
AES hardware acceleration, including on mobile devices such as Google Glass and older Android
experience, reducing latency and saving a significant amount of time spent encrypting and

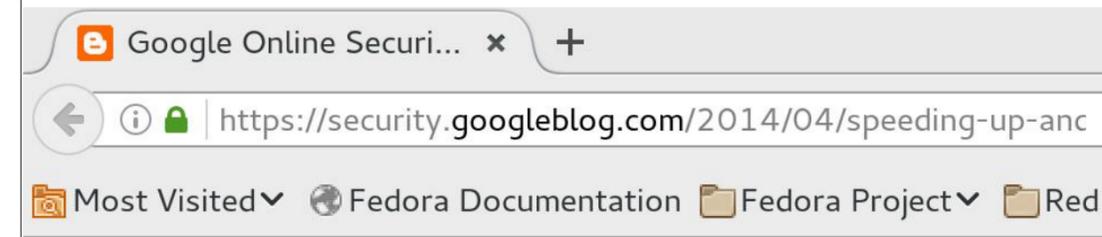
To make this happen, Adam Langley, Google's lead TLS expert, began implementing new algorithms for
encryption and Poly1305 for authentication. This effort started in
March 2013. It was a complex effort that required a new
abstraction layer in OpenSSL in order

Two examples of speed reported in this 2017 paper for a 3.4GHz Skylake (Intel Core i7-6700):

383.69 MByte/sec (8.86 cycles/byte) for AES CTR-DRBG using AES-NI; 106.07 MByte/sec (32 cycles/byte) for ChaCha20.

But wait. eBACS reports 0.92 cycles/byte for AES-256-CTR, 1.18 cycles/byte for ChaCha20.

Author non-response: “essential for us to examine standard open implementations”. Slow ones?



Speeding up and strengthening HTTPS connections for Chrome on Android

April 24, 2014

Posted by Elie Bursztein, Anti-Abuse Research Lead

Earlier this year, we deployed a new TLS cipher suite in Chrome that operates three times faster than AES-GCM on devices that lack AES hardware acceleration, including most Android phones and devices such as Google Glass and older computers. This improves user experience, reducing latency and saving battery life by cutting down the amount of time spent encrypting and decrypting data.

To make this happen, Adam Langley, Wan-Teh Chang, Ben Laurie, and I began implementing new algorithms – ChaCha 20 for symmetric encryption and Poly1305 for authentication – in OpenSSL in March 2013. It was a complex effort that required implementing an abstraction layer in OpenSSL in order to support the Authen

Two examples of speed reported in this 2017 paper for a 3.4GHz Skylake (Intel Core i7-6700):

383.69 MByte/sec (8.86 cycles/byte) for AES CTR-DRBG using AES-NI; 106.07 MByte/sec (32 cycles/byte) for ChaCha20.

But wait. eBACS reports 0.92 cycles/byte for AES-256-CTR, 1.18 cycles/byte for ChaCha20.

Author non-response: “essential for us to examine standard open implementations”. Slow ones?

The screenshot shows a Chrome browser window with the following details:

- Address Bar:** <https://security.googleblog.com/2014/04/speeding-up-anc>
- Page Title:** Speeding up and strengthening HTTPS connections for Chrome on Android
- Date:** April 24, 2014
- Author:** Posted by Elie Bursztein, Anti-Abuse Research Lead
- Content:**

Earlier this year, we deployed a new TLS cipher suite in Chrome that operates three times faster than AES-GCM on devices that don't have AES hardware acceleration, including most Android phones, wearable devices such as Google Glass and older computers. This improves user experience, reducing latency and saving battery life by cutting down the amount of time spent encrypting and decrypting data.

To make this happen, Adam Langley, Wan-Teh Chang, Ben Laurie and I began implementing new algorithms – ChaCha 20 for symmetric encryption and Poly1305 for authentication – in OpenSSL and NSS in March 2013. It was a complex effort that required implementing a new abstraction layer in OpenSSL in order to support the Authenticated

Examples of speed reported
 2017 paper for a 3.4GHz
 (Intel Core i7-6700):

106.07 MByte/sec (8.86
 byte) for AES CTR-DRBG
 ES-NI; 106.07 MByte/sec
 es/byte) for ChaCha20.

eBACS reports
 es/byte for AES-256-CTR,
 es/byte for ChaCha20.

non-response: “essential
 to examine standard open
 ntations”. Slow ones?

The image shows a side-by-side comparison of two browser windows. The left window displays a Google Security Blog article from April 24, 2014, titled "Speeding up and strengthening HTTPS connections for Chrome on Android". The article is by Elie Bursztein, Anti-Abuse Research Lead. The text discusses the deployment of a new TLS cipher suite in Chrome for devices without AES hardware acceleration, mentioning ChaCha 20 for symmetric encryption and Poly1305 for authentication. The right window shows a partial view of another article titled "Do the ChaCha: b" with a URL starting with "https://blog".

Google Online Security

https://security.googleblog.com/2014/04/speeding-up-anc

Most Visited Fedora Documentation Fedora Project Red Hat

Speeding up and strengthening HTTPS connections for Chrome on Android

April 24, 2014

Posted by Elie Bursztein, Anti-Abuse Research Lead

Earlier this year, we deployed a new TLS cipher suite in Chrome that operates three times faster than AES-GCM on devices that don't have AES hardware acceleration, including most Android phones, wearable devices such as Google Glass and older computers. This improves user experience, reducing latency and saving battery life by cutting down the amount of time spent encrypting and decrypting data.

To make this happen, Adam Langley, Wan-Teh Chang, Ben Laurie and I began implementing new algorithms – ChaCha 20 for symmetric encryption and Poly1305 for authentication – in OpenSSL and NSS in March 2013. It was a complex effort that required implementing a new abstraction layer in OpenSSL in order to support the Authenticated

Do the ChaCha: b

https://blog

Most Visited Fe

Today we ar
 form of encr
 performance
 today, Goog
 the Internet
 all sites on C
 mobile brow
 visiting sites

As of the lau
 10% of https
 ciphersuites
 when we tur

CloudFlare ciphers
 February 23, 2015

40

speed reported
for a 3.4GHz
e i7-6700):

(8.86
ES CTR-DRBG
5.07 MByte/sec
or ChaCha20.

reports
or AES-256-CTR,
or ChaCha20.

se: “essential
standard open
Slow ones?



Speeding up and strengthening HTTPS connections for Chrome on Android

April 24, 2014

Posted by Elie Bursztein, Anti-Abuse Research Lead

Earlier this year, we deployed a new TLS cipher suite in Chrome that operates three times faster than AES-GCM on devices that don't have AES hardware acceleration, including most Android phones, wearable devices such as Google Glass and older computers. This improves user experience, reducing latency and saving battery life by cutting down the amount of time spent encrypting and decrypting data.

To make this happen, Adam Langley, Wan-Teh Chang, Ben Laurie and I began implementing new algorithms – ChaCha 20 for symmetric encryption and Poly1305 for authentication – in OpenSSL and NSS in March 2013. It was a complex effort that required implementing a new abstraction layer in OpenSSL in order to support the Authenticated



Today we are adding a new form of encryption — that improves performance: ChaCha20-Poly1305. As of the launch today (February 23, 2015), 10% of https connections to sites on CloudFlare supported ChaCha20-Poly1305. All sites on CloudFlare supported ChaCha20-Poly1305 today, Google services were the first. This means that all sites on CloudFlare supported ChaCha20-Poly1305 today, Google services were the first. This means that all sites on CloudFlare supported ChaCha20-Poly1305 today, Google services were the first.

As of the launch today (February 23, 2015), 10% of https connections to sites on CloudFlare supported ChaCha20-Poly1305. All sites on CloudFlare supported ChaCha20-Poly1305 today, Google services were the first. This means that all sites on CloudFlare supported ChaCha20-Poly1305 today, Google services were the first.

CloudFlare ciphersuite chosen by percentage
February 23, 2015



orted
GHz

:

DRBG

e/sec

20.

6-CTR,

20.

ntial

open

es?

Google Online Securi... x +

https://security.googleblog.com/2014/04/speeding-up-anc

Most Visited Fedora Documentation Fedora Project Red Hat

Speeding up and strengthening HTTPS connections for Chrome on Android

April 24, 2014

Posted by Elie Bursztein, Anti-Abuse Research Lead

Earlier this year, we deployed a new TLS cipher suite in Chrome that operates three times faster than AES-GCM on devices that don't have AES hardware acceleration, including most Android phones, wearable devices such as Google Glass and older computers. This improves user experience, reducing latency and saving battery life by cutting down the amount of time spent encrypting and decrypting data.

To make this happen, Adam Langley, Wan-Teh Chang, Ben Laurie and I began implementing new algorithms – ChaCha 20 for symmetric encryption and Poly1305 for authentication – in OpenSSL and NSS in March 2013. It was a complex effort that required implementing a new abstraction layer in OpenSSL in order to support the Authenticated

Do the ChaCha: bett... x +

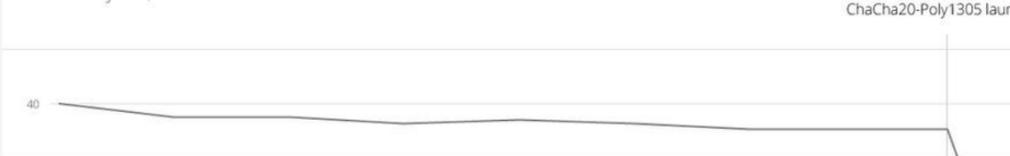
https://blog.cloudflare.com/do-the-chacha-better-mobile-p

Most Visited Fedora Documentation Fedora Project Red

Today we are adding a new feature — actual form of encryption — that improves mobile performance: ChaCha20-Poly1305 cipher suite. Today, Google services were the only major sites on the Internet that supported this new algorithm. Now, all sites on CloudFlare support it, too. This means mobile browsers get a better experience when visiting sites using CloudFlare.

As of the launch today (February 23, 2015), 10% of https connections to CloudFlare use ChaCha20-Poly1305 ciphersuites. The following graph shows the percentage of connections when we turned ChaCha20/Poly1305 on globally.

CloudFlare ciphersuite chosen by percentage
February 23, 2015



Google Online Securi... x +

https://security.googleblog.com/2014/04/speeding-up-anc

Most Visited Fedora Documentation Fedora Project Red Hat

Speeding up and strengthening HTTPS connections for Chrome on Android

April 24, 2014

Posted by Elie Bursztein, Anti-Abuse Research Lead

Earlier this year, we deployed a new TLS cipher suite in Chrome that operates three times faster than AES-GCM on devices that don't have AES hardware acceleration, including most Android phones, wearable devices such as Google Glass and older computers. This improves user experience, reducing latency and saving battery life by cutting down the amount of time spent encrypting and decrypting data.

To make this happen, Adam Langley, Wan-Teh Chang, Ben Laurie and I began implementing new algorithms – ChaCha 20 for symmetric encryption and Poly1305 for authentication – in OpenSSL and NSS in March 2013. It was a complex effort that required implementing a new abstraction layer in OpenSSL in order to support the Authenticated

Do the ChaCha: bett... x +

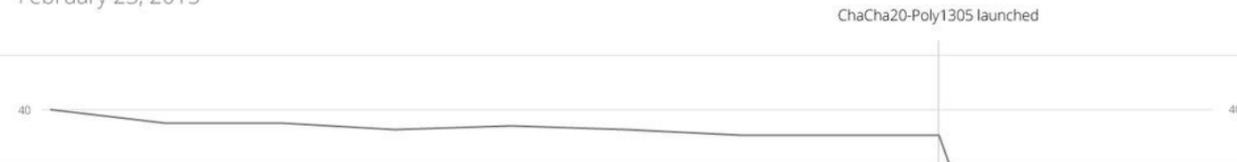
https://blog.cloudflare.com/do-the-chacha-better-mobile-p

Most Visited Fedora Documentation Fedora Project Red Hat

Today we are adding a new feature — actually a new form of encryption — that improves mobile performance: ChaCha20-Poly1305 cipher suites. Until today, Google services were the only major sites on the Internet that supported this new algorithm. Now all sites on CloudFlare support it, too. This means mobile browsers get a better experience when visiting sites using CloudFlare.

As of the launch today (February 23, 2015), nearly 10% of https connections to CloudFlare use the new ciphersuites. The following graph shows the uptick when we turned ChaCha20/Poly1305 on globally:

CloudFlare ciphersuite chosen by percentage
February 23, 2015



curi... x +

curity.googleblog.com/2014/04/speeding-up-anc

Fedora Documentation Fedora Project Red Hat

Speeding up and strengthening HTTPS connections for Chrome on Android

Wan-Teh Chang, Anti-Abuse Research Lead

We recently deployed a new TLS cipher suite in Chrome that is up to 2x faster than AES-GCM on devices that don't have hardware acceleration, including most Android phones, wearable devices like Google Glass and older computers. This improves user experience by reducing latency and saving battery life by cutting down the amount of time spent encrypting and decrypting data.

This effort was a complex effort that required implementing a new cipher suite in OpenSSL in order to support the Authenticated Encryption with Associated Data (AEAD) mode. Adam Langley, Wan-Teh Chang, Ben Laurie and I worked on this effort. We implemented ChaCha 20 for symmetric encryption and Poly1305 for authentication – in OpenSSL and NSS in order to support the Authenticated Encryption with Associated Data (AEAD) mode.

Do the ChaCha: better performance on mobile devices x +

https://blog.cloudflare.com/do-the-chacha-better-mobile-p

Most Visited Fedora Documentation Fedora Project Red Hat

Today we are adding a new feature — actually a new form of encryption — that improves mobile performance: ChaCha20-Poly1305 cipher suites. Until today, Google services were the only major sites on the Internet that supported this new algorithm. Now all sites on CloudFlare support it, too. This means mobile browsers get a better experience when visiting sites using CloudFlare.

As of the launch today (February 23, 2015), nearly 10% of https connections to CloudFlare use the new cipher suites. The following graph shows the uptick when we turned ChaCha20/Poly1305 on globally:



ImperialViolet - Maybe Skip SHA-1

https://www.imperialviolet.org/2014/04/28/maybe-skip-sha1.html

Most Visited Fedora Documentation Fedora Project Red Hat

Maybe Skip SHA-1

In 2005 and 2006, [1][2][3]. These reports were written by cryptographers at all. After all, many people have been using SHA-1 for years.

In the wake of this, in order to hedge the risk of a "ket-chak", I believe we proved that we do not need SHA-1. SHA-1 didn't extend to hash functions, all it existed, it was not a tendency to assume the number is bigger.

As I've mentioned, it contributes to the tested and harder platforms typically code-size, which is even slower than crypto primitives.

14/04/speeding-up-anc

Fedora Project Red Hat

enning HTTPS

n Android

h Lead

LS cipher suite in Chrome that
-GCM on devices that don't have
most Android phones, wearable
der computers. This improves user
ing battery life by cutting down the
decrypting data.

Wan-Teh Chang, Ben Laurie and I
- ChaCha 20 for symmetric
ication - in OpenSSL and NSS in
that required implementing a new
to support the Authenticated

Do the ChaCha: bett...

https://blog.cloudflare.com/do-the-chacha-better-mobile-p

Most Visited Fedora Documentation Fedora Project Red Hat

Today we are adding a new feature — actually a new form of encryption — that improves mobile performance: ChaCha20-Poly1305 cipher suites. Until today, Google services were the only major sites on the Internet that supported this new algorithm. Now all sites on CloudFlare support it, too. This means mobile browsers get a better experience when visiting sites using CloudFlare.

As of the launch today (February 23, 2015), nearly 10% of https connections to CloudFlare use the new ciphersuites. The following graph shows the uptick when we turned ChaCha20/Poly1305 on globally:



ImperialViolet - Maybe S...

https://www.imperialviolet.org/201

Most Visited Fedora Documentation

Maybe Skip SHA-3 (31 May 2017)

In 2005 and 2006, a series of significant [1][2][3]. These repeated break-throughs as cryptographers questioned whether at all. After all, many hash functions from

In the wake of this, NIST announced (PDF) order to hedge the risk of SHA-2 fa “ket-chak”, I believe) won (PDF) and be proved that we *do* know how to build h 2005 didn't extend to SHA-2 and the S hash functions, all of which are secure a it existed, it was no longer clear that SH tendency to assume that SHA-3 must be ber is bigger.

As I've mentioned before, diversity of It contributes to the exponential num tested and hardened; it draws on lim platforms typically need separate, op code-size, which is a worry again in the even slower than SHA-2 which is already crypto primitives

Do the ChaCha: bett... x +

https://blog.cloudflare.com/do-the-chacha-better-mobile-pi

Hat v Most Visited v Fedora Documentation Fedora Project v Red Hat v

Today we are adding a new feature — actually a new form of encryption — that improves mobile performance: ChaCha20-Poly1305 cipher suites. Until today, Google services were the only major sites on the Internet that supported this new algorithm. Now all sites on CloudFlare support it, too. This means mobile browsers get a better experience when visiting sites using CloudFlare.

As of the launch today (February 23, 2015), nearly 10% of https connections to CloudFlare use the new ciphersuites. The following graph shows the uptick when we turned ChaCha20/Poly1305 on globally:



ImperialViolet - Maybe S... x +

https://www.imperialviolet.org/2017/05/31/skipsha3.html

Hat v Most Visited v Fedora Documentation Fedora Project v Red Hat v

Maybe Skip SHA-3 (31 May 2017)

In 2005 and 2006, a series of significant results were published [1][2][3]. These repeated break-throughs caused something of a panic as cryptographers questioned whether we knew how to build hash functions at all. After all, many hash functions from the 1990's had not

In the wake of this, NIST announced (PDF) a competition to order to hedge the risk of SHA-2 falling. In 2012, Keccak (“ket-chak”, I believe) won (PDF) and became SHA-3. But the competition proved that we *do* know how to build hash functions: the series of results in 2005 didn't extend to SHA-2 and the SHA-3 process produced many secure hash functions, all of which are secure as far as we can tell. If SHA-3 existed, it was no longer clear that SHA-3 was needed. Yet the common tendency to assume that SHA-3 must be better than SHA-2 because its key size is bigger.

As I've mentioned before, diversity of cryptographic primitives is good. It contributes to the exponential number of combinations of primitives tested and hardened; it draws on limited developer resources; and it means that platforms typically need separate, optimised code; and it means that code-size, which is a worry again in the mobile age. SHA-3 is even slower than SHA-2 which is already a comparative disadvantage for crypto primitives.



Today we are adding a new feature — actually a new form of encryption — that improves mobile performance: ChaCha20-Poly1305 cipher suites. Until today, Google services were the only major sites on the Internet that supported this new algorithm. Now all sites on CloudFlare support it, too. This means mobile browsers get a better experience when visiting sites using CloudFlare.

As of the launch today (February 23, 2015), nearly 10% of https connections to CloudFlare use the new ciphersuites. The following graph shows the uptick when we turned ChaCha20/Poly1305 on globally:



Maybe Skip SHA-3 (31 May 2017)

In 2005 and 2006, a series of significant results were published against SHA-1 [1][2][3]. These repeated break-throughs caused something of a crisis of faith as cryptographers questioned whether we knew how to build hash functions at all. After all, many hash functions from the 1990's had not aged well [1][2].

In the wake of this, NIST announced (PDF) a competition to develop SHA-3 in order to hedge the risk of SHA-2 falling. In 2012, Keccak (pronounced “ket-chak”, I believe) won (PDF) and became SHA-3. But the competition itself proved that we *do* know how to build hash functions: the series of results in 2005 didn't extend to SHA-2 and the SHA-3 process produced a number of hash functions, all of which are secure as far as we can tell. Thus, by the time it existed, it was no longer clear that SHA-3 was needed. Yet there is a natural tendency to assume that SHA-3 must be better than SHA-2 because the number is bigger.

As I've mentioned before, diversity of cryptographic primitives is expensive. It contributes to the exponential number of combinations that need to be tested and hardened; it draws on limited developer resources as multiple platforms typically need separate, optimised code; and it contributes to code-size, which is a worry again in the mobile age. SHA-3 is also slow, and is even slower than SHA-2 which is already a comparative laggard amongst crypto primitives.