Building a battlefield for authenticated encryption

D. J. Bernstein

University of Illinois at Chicago

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Building a battlefield for authenticated encryption

D. J. Bernstein

University of Illinois at Chicago

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3, which is faster. *Quel surprise!*

a battlefield

enticated encryption

ernstein

ty of Illinois at Chicago

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were th
the state

Not ever

- better
  (e.g., 2

eld
encryption

is at Chicago

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:
3.73 on Core i5-650.
3.88 in 32-bit mode.
10.9 without AES insns.
39.3 on UltraSPARC III.
50.8 on ARM Cortex A8.
53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were these AES-G
the state of the ar

Not even close. Pa
- better AES impl
  (e.g., 2008 Bern

ago

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were these AES-GCM speed
the state of the art?

Not even close. Paper is ign

- better AES implementatio
  (e.g., 2008 Bernstein–Schw

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Krovetz–Rogaway, tomorrow:

Look at how slow AES-GCM is!

Cycles/byte for 4096-byte
authenticated encryption:

3.73 on Core i5-650.

3.88 in 32-bit mode.

10.9 without AES insns.

39.3 on UltraSPARC III.

50.8 on ARM Cortex A8.

53.5 on PowerPC 970.

Paper advertises AES-OCB3,
which is faster. *Quel surprise!*

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

–Rogaway, tomorrow:

how slow AES-GCM is!

byte for 4096-byte
cated encryption:
n Core i5-650.
32-bit mode.
ithout AES insns.
n UltraSPARC III.
n ARM Cortex A8.
n PowerPC 970.

dvertises AES-OCB3,
faster. *Quel surprise!*

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do

tomorrow:

AES-GCM is!

096-byte
ryption:
650.
ode.
S insns.
ARC III.
rtex A8.
C 970.

AES-OCB3,
uel surprise!

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do we do af

v:

M is!

3,

se!

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do we do after SHA-3

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do we do after SHA-3?

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Were these AES-GCM speeds
the state of the art?

Not even close. Paper is ignoring

- better AES implementations
  (e.g., 2008 Bernstein–Schwabe);
- faster ciphers than AES-CTR
  (e.g., any eSTREAM finalist);
- faster authenticators
  (e.g., Poly1305, HMAC-???);
- serious redesigns
  (e.g., Phelix, Grain-128a).

Paper is also sloppy with security.
Big trouble near 2^{64} blocks,
avoided by some older schemes.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

ese AES-GCM speeds
e of the art?

n close. Paper is ignoring
AES implementations
2008 Bernstein–Schwabe);
ciphers than AES-CTR
any eSTREAM finalist);
authenticators
Poly1305, HMAC-???);
s redesigns
Phelix, Grain-128a).

also sloppy with security.
ble near 2^{64} blocks,
by some older schemes.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competi
thanks t

GCM speeds
t?

aper is ignoring
ementations
stein–Schwabe);
an AES-CTR
EAM finalist);
ators
HMAC-???);
s
ain-128a).

by with security.
^{64} blocks,
older schemes.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition alrea
thanks to Greg Ro

ds

noring
ns
wabe);
TR
ist);

??);

curity.

mes.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a n
thanks to Greg Rose: eSAF

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Secure
Authenticated
Fast
Encryption

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Slow
Authentication,
Flimsy
Encryption

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Secures
Additional
Funding for
ECRYPT

What do we do after SHA-3?
Let's have a competition
for authenticated encryption!
Much more fun than, e.g.,
cycling back to block ciphers.

Easy: Speed competition.
ECRYPT benchmarking will soon
cover authenticated encryption.

Hard: Security competition.
Needs community to focus.

Potential timing problem:
NIST needs to take a break.
ECRYPT II ends in 2012.
But does this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Secure
Authenticated
Fast
Encryption

o we do after SHA-3?
ve a competition
enticated encryption!
ore fun than, e.g.,
back to block ciphers.

peed competition.
T benchmarking will soon
uthenticated encryption.

ecurity competition.
ommunity to focus.

l timing problem:
eeds to take a break.
T II ends in 2012.
s this really matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Secure
Authenticated
Fast
Encryption

ter SHA-3?
etition
encryption!
an, e.g.,
ock ciphers.

petition.
arking will soon
ed encryption.

mpetition.
to focus.

roblem:
ke a break.
n 2012.
y matter?

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Secure
Authenticated
Fast
Encryption

?

!

s.

soon

ion.

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Secure
Authenticated
Fast
Encryption

Competition already has a name,
thanks to Greg Rose: eSAFE.
(Only 655000 Google hits.)

What does eSAFE stand for?
Not sure yet.

ECRYPT
Secure
Authenticated
Fast
Encryption