

Some challenges in
heavyweight cipher design

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Protocol generates
new AES-128 key k .

Protocol encrypts message block
 m_1 as $\text{AES}_k(1) \oplus m_1$,
 m_2 as $\text{AES}_k(2) \oplus m_2$,
 m_3 as $\text{AES}_k(3) \oplus m_3$,
etc. Also authenticates.

First block m_1 is predictable:

GET / HTTP/1.1\r\n

Attacker learns $\text{AES}_k(1)$.

Can attacker deduce $\text{AES}_k(20)$?

We constantly tell people: “No!
AES is secure! This is all safe!”

challenges in
light cipher design

. Bernstein

ty of Illinois at Chicago &
the Universiteit Eindhoven

Protocol generates
new AES-128 key k .

Protocol encrypts message block
 m_1 as $\text{AES}_k(1) \oplus m_1$,
 m_2 as $\text{AES}_k(2) \oplus m_2$,
 m_3 as $\text{AES}_k(3) \oplus m_3$,
etc. Also authenticates.

First block m_1 is predictable:

GET / HTTP/1.1\r\n

Attacker learns $\text{AES}_k(1)$.

Can attacker deduce $\text{AES}_k(20)$?

We constantly tell people: “No!
AES is secure! This is all safe!”

Attacker
for, say,

Attacker
using fea
Attacker
data for

Is this 2
See 2002

n
r design

n
is at Chicago &
siteit Eindhoven

Protocol generates
new AES-128 key k .

Protocol encrypts message block
 m_1 as $AES_k(1) \oplus m_1$,
 m_2 as $AES_k(2) \oplus m_2$,
 m_3 as $AES_k(3) \oplus m_3$,
etc. Also authenticates.

First block m_1 is predictable:

GET / HTTP/1.1\r\n

Attacker learns $AES_k(1)$.

Can attacker deduce $AES_k(20)$?
We constantly tell people: "No!
AES is secure! This is all safe!"

Attacker learns AES_k
for, say, 2^{40} user keys

Attacker finds *some* data
using feasible 2^{88} operations
Attacker decrypts,
data for that user.

Is this 2^{128} "secure"?
See 2002 Biham "A Chosen-Plaintext Attack on the
AES-128 Cipher"

Protocol generates
new AES-128 key k .

Protocol encrypts message block
 m_1 as $\text{AES}_k(1) \oplus m_1$,
 m_2 as $\text{AES}_k(2) \oplus m_2$,
 m_3 as $\text{AES}_k(3) \oplus m_3$,
etc. Also authenticates.

First block m_1 is predictable:

GET / HTTP/1.1\r\n

Attacker learns $\text{AES}_k(1)$.

Can attacker deduce $\text{AES}_k(20)$?

We constantly tell people: “No!
AES is secure! This is all safe!”

Attacker learns $\text{AES}_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.
Attacker decrypts, maybe for
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collision”

Protocol generates
new AES-128 key k .

Protocol encrypts message block
 m_1 as $\text{AES}_k(1) \oplus m_1$,
 m_2 as $\text{AES}_k(2) \oplus m_2$,
 m_3 as $\text{AES}_k(3) \oplus m_3$,
etc. Also authenticates.

First block m_1 is predictable:

GET / HTTP/1.1\r\n

Attacker learns $\text{AES}_k(1)$.

Can attacker deduce $\text{AES}_k(20)$?

We constantly tell people: “No!
AES is secure! This is all safe!”

Attacker learns $\text{AES}_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.
Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collisions”.

Protocol generates
new AES-128 key k .

Protocol encrypts message block
 m_1 as $\text{AES}_k(1) \oplus m_1$,
 m_2 as $\text{AES}_k(2) \oplus m_2$,
 m_3 as $\text{AES}_k(3) \oplus m_3$,
etc. Also authenticates.

First block m_1 is predictable:

GET / HTTP/1.1\r\n

Attacker learns $\text{AES}_k(1)$.

Can attacker deduce $\text{AES}_k(20)$?

We constantly tell people: “No!
AES is secure! This is all safe!”

Attacker learns $\text{AES}_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.
Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collisions”.

Fragile fix: Complicate protocols
by trying to randomize everything.

Protocol generates
new AES-128 key k .

Protocol encrypts message block
 m_1 as $\text{AES}_k(1) \oplus m_1$,
 m_2 as $\text{AES}_k(2) \oplus m_2$,
 m_3 as $\text{AES}_k(3) \oplus m_3$,
etc. Also authenticates.

First block m_1 is predictable:

GET / HTTP/1.1\r\n

Attacker learns $\text{AES}_k(1)$.

Can attacker deduce $\text{AES}_k(20)$?

We constantly tell people: “No!
AES is secure! This is all safe!”

Attacker learns $\text{AES}_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.
Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collisions”.

Fragile fix: Complicate protocols
by trying to randomize everything.

Much simpler fix: 256-bit keys.
(Side discussion: Is 192 enough?)

generates
128-bit key k .

encrypts message block

$$AES_k(1) \oplus m_1,$$

$$AES_k(2) \oplus m_2,$$

$$AES_k(3) \oplus m_3,$$

and authenticates.

Block m_1 is predictable:

HTTP/1.1\r\n

... learns $AES_k(1)$.

Attacker deduce $AES_k(20)$?

Constantly tell people: "No!

Secure! This is all safe!"

Attacker learns $AES_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key

using feasible 2^{88} computation.

Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} "security"?

See 2002 Biham "key collisions".

Fragile fix: Complicate protocols
by trying to randomize everything.

Much simpler fix: 256-bit keys.

(Side discussion: Is 192 enough?)

Another
about 12
quantum

Grover f
using 2^6
on a sm

s
 k .
message block
 m_1 ,
 m_2 ,
 m_3 ,
cates.
predictable:
 $\backslash n$
 $ES_k(1)$.
ce $AES_k(20)$?
people: "No!
is is all safe!"

Attacker learns $AES_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.

Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} "security"?

See 2002 Biham "key collisions".

Fragile fix: Complicate protocols
by trying to randomize everything.

Much simpler fix: 256-bit keys.
(Side discussion: Is 192 enough?)

Another reason to
about 128-bit ciph
quantum computin
Grover finds k from
using 2^{64} iteration
on a small quantu

Attacker learns $AES_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.

Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collisions”.

Fragile fix: Complicate protocols
by trying to randomize everything.

Much simpler fix: 256-bit keys.
(Side discussion: Is 192 enough?)

Another reason to be concerned
about 128-bit cipher keys:
quantum computing.

Grover finds k from $AES_k(1)$
using 2^{64} iterations
on a small quantum processor.

Attacker learns $AES_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.

Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collisions”.

Fragile fix: Complicate protocols
by trying to randomize everything.

Much simpler fix: 256-bit keys.

(Side discussion: Is 192 enough?)

Another reason to be concerned
about 128-bit cipher keys:
quantum computing.

Grover finds k from $AES_k(1)$
using 2^{64} iterations
on a small quantum processor.

Attacker learns $AES_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.

Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collisions”.

Fragile fix: Complicate protocols
by trying to randomize everything.

Much simpler fix: 256-bit keys.

(Side discussion: Is 192 enough?)

Another reason to be concerned
about 128-bit cipher keys:
quantum computing.

Grover finds k from $AES_k(1)$
using 2^{64} iterations
on a small quantum processor.

Parallelize: N^2 processors,
each running $2^{64}/N$ iterations.
1999 Zalka claims this is optimal.

Attacker learns $AES_k(1)$
for, say, 2^{40} user keys k .

Attacker finds *some* user key
using feasible 2^{88} computation.

Attacker decrypts, maybe forges,
data for that user.

Is this 2^{128} “security”?

See 2002 Biham “key collisions”.

Fragile fix: Complicate protocols
by trying to randomize everything.

Much simpler fix: 256-bit keys.
(Side discussion: Is 192 enough?)

Another reason to be concerned
about 128-bit cipher keys:
quantum computing.

Grover finds k from $AES_k(1)$
using 2^{64} iterations
on a small quantum processor.

Parallelize: N^2 processors,
each running $2^{64}/N$ iterations.
1999 Zalka claims this is optimal.

Multiple targets should allow
much better parallelization.

Related algos: 2009 Bernstein;
2004 Grover–Radhakrishnan.

learns $AES_k(1)$
 2^{40} user keys k .

finds *some* user key

feasible 2^{88} computation.

decrypts, maybe forges,
that user.

128 “security”?

2 Biham “key collisions”.

fix: Complicate protocols
to randomize everything.

simpler fix: 256-bit keys.

discussion: Is 192 enough?)

Another reason to be concerned
about 128-bit cipher keys:
quantum computing.

Grover finds k from $AES_k(1)$
using 2^{64} iterations
on a small quantum processor.

Parallelize: N^2 processors,
each running $2^{64}/N$ iterations.
1999 Zalka claims this is optimal.

Multiple targets should allow
much better parallelization.

Related algos: 2009 Bernstein;
2004 Grover–Radhakrishnan.

Should M

To autho

Computo

different

e.g., $r^4 m$

where r

Generate

$s_n = AE$

Add to c

$r^4 m_1 +$

Widely c

consider

$AES_k(1)$

keys k .

the user key

computation.

maybe forges,

ity"?

"key collisions".

icate protocols

mize everything.

256-bit keys.

s 192 enough?)

Another reason to be concerned about 128-bit cipher keys: quantum computing.

Grover finds k from $AES_k(1)$ using 2^{64} iterations on a small quantum processor.

Parallelize: N^2 processors, each running $2^{64}/N$ iterations. 1999 Zalka claims this is optimal.

Multiple targets should allow much better parallelization.

Related algos: 2009 Bernstein; 2004 Grover–Radhakrishnan.

Should MACs have

To authenticate (m)

Compute function differential probability e.g., $r^4 m_1 + r^3 m_2$ where r is secret.

Generate a **one-time** $s_n = AES_k(n)$ from

Add to obtain MA $r^4 m_1 + r^3 m_2 + r^2$

Widely deployed for consider, e.g., GCM

Another reason to be concerned about 128-bit cipher keys: quantum computing.

Grover finds k from $\text{AES}_k(1)$ using 2^{64} iterations on a small quantum processor.

Parallelize: N^2 processors, each running $2^{64}/N$ iterations. 1999 Zalka claims this is optimal.

Multiple targets should allow much better parallelization.

Related algos: 2009 Bernstein; 2004 Grover–Radhakrishnan.

Should MACs have nonces?

To authenticate (m_1, m_2, m_3)

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 +$
where r is secret.

Generate a **one-time** key
 $s_n = \text{AES}_k(n)$ from master key

Add to obtain MAC:
 $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$

Widely deployed for speed: consider, e.g., GCM.

Another reason to be concerned about 128-bit cipher keys: quantum computing.

Grover finds k from $\text{AES}_k(1)$ using 2^{64} iterations on a small quantum processor.

Parallelize: N^2 processors, each running $2^{64}/N$ iterations. 1999 Zalka claims this is optimal.

Multiple targets should allow much better parallelization.

Related algos: 2009 Bernstein; 2004 Grover–Radhakrishnan.

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$, where r is secret.

Generate a **one-time** key

$s_n = \text{AES}_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed: consider, e.g., GCM.

reason to be concerned
28-bit cipher keys:
n computing.

inds k from $AES_k(1)$
 4 iterations
all quantum processor.

ze: N^2 processors,
ning $2^{64}/N$ iterations.
lka claims this is optimal.

targets should allow
etter parallelization.

algos: 2009 Bernstein;
over–Radhakrishnan.

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small
differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$,
where r is secret.

Generate a **one-time** key

$s_n = AES_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed:

consider, e.g., GCM.

2006 Jor
ntwice in
 \Rightarrow attac
can easil

be concerned

er keys:

ng.

m $AES_k(1)$

s

m processor.

processors,

N iterations.

this is optimal.

ould allow

elization.

09 Bernstein;

makrishnan.

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$,
where r is secret.

Generate a **one-time** key

$s_n = AES_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed:

consider, e.g., GCM.

2006 Joux “forbid

ntwice in GCM \Rightarrow

\Rightarrow attacker figures

can easily forge m

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$,
where r is secret.

Generate a **one-time** key

$s_n = \text{AES}_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed:

consider, e.g., GCM.

2006 Joux “forbidden attack”
repeated twice in GCM \Rightarrow repeated
 \Rightarrow attacker figures out r ,
can easily forge messages.

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$,
where r is secret.

Generate a **one-time** key

$s_n = \text{AES}_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed:

consider, e.g., GCM.

2006 Joux “forbidden attack”:
nonce repeated in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$,
where r is secret.

Generate a **one-time** key

$s_n = \text{AES}_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed:

consider, e.g., GCM.

2006 Joux “forbidden attack”:
nonce twice in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:

$\text{AES}_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$

“seems a safe option”. (Also
suggested and analyzed in, e.g.,
2000 Bernstein; earlier refs?)

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$, where r is secret.

Generate a **one-time** key

$s_n = \text{AES}_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed:

consider, e.g., GCM.

2006 Joux “forbidden attack”:
repeated in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:

$\text{AES}_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$

“seems a safe option”. (Also suggested and analyzed in, e.g., 2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Should MACs have nonces?

To authenticate (m_1, m_2, m_3, m_4) :

Compute function with small differential probabilities.

e.g., $r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4$, where r is secret.

Generate a **one-time** key

$s_n = \text{AES}_k(n)$ from master key k .

Add to obtain MAC:

$r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Widely deployed for speed:

consider, e.g., GCM.

2006 Joux “forbidden attack”:
twice in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:

$\text{AES}_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$

“seems a safe option”. (Also suggested and analyzed in, e.g., 2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Forgery chance $\leq \delta + \epsilon$ where

ϵ is AES PRF insecurity and

$\delta \approx q^2 L / 2^{128}$

for message lengths $\leq L$.

MACs have nonces?

Authenticate (m_1, m_2, m_3, m_4) :

Use function with small
collision probabilities.

$m_1 + r^3 m_2 + r^2 m_3 + r m_4$,
where r is secret.

Use a **one-time** key

$S_k(n)$ from master key k .

How to obtain MAC:

$r^3 m_2 + r^2 m_3 + r m_4 + s_n$.

Not deployed for speed:

Not used, e.g., GCM.

2006 Joux “forbidden attack”:
nonce used twice in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
and can easily forge messages.

Joux’s suggested response:

$AES_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$

“seems a safe option”. (Also
suggested and analyzed in, e.g.,
2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Forgery chance $\leq \delta + \epsilon$ where

ϵ is AES PRF insecurity and

$\delta \approx q^2 L / 2^{128}$

for message lengths $\leq L$.

ϵ is at le

Solution

(2005 B

re nonces?

(m_1, m_2, m_3, m_4) :

with small
probabilities.

$+ r^2 m_3 + r m_4,$

same key

from master key k .

MAC:

$m_3 + r m_4 + s_n.$

for speed:

M.

2006 Joux “forbidden attack”:
nonce twice in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:

$\text{AES}_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$

“seems a safe option”. (Also
suggested and analyzed in, e.g.,
2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Forgery chance $\leq \delta + \epsilon$ where

ϵ is AES PRF insecurity and

$\delta \approx q^2 L / 2^{128}$

for message lengths $\leq L$.

ϵ is at least $q(q - 1)$
Solution: better PRF
(2005 Bernstein),

m_3, m_4):

||

$- r m_4,$

key k .

$+ s_n$.

2006 Joux “forbidden attack”:
 twice in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
 can easily forge messages.

Joux’s suggested response:
 $AES_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$
 “seems a safe option”. (Also
 suggested and analyzed in, e.g.,
 2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?
 Forgery chance $\leq \delta + \epsilon$ where
 ϵ is AES PRF insecurity and
 $\delta \approx q^2 L / 2^{128}$
 for message lengths $\leq L$.

ϵ is at least $q(q - 1) / 2^{129}$.
 Solution: better PRP/PRF
 (2005 Bernstein), ok for $q \approx$

2006 Joux “forbidden attack”:
repeated s_n in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:

$$\text{AES}_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$$

“seems a safe option”. (Also
suggested and analyzed in, e.g.,
2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Forgery chance $\leq \delta + \epsilon$ where

ϵ is AES PRF insecurity and

$$\delta \approx q^2 L / 2^{128}$$

for message lengths $\leq L$.

ϵ is at least $q(q - 1) / 2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

2006 Joux “forbidden attack”:
nonce in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:

$$\text{AES}_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$$

“seems a safe option”. (Also
suggested and analyzed in, e.g.,
2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Forgery chance $\leq \delta + \epsilon$ where

ϵ is AES PRF insecurity and

$$\delta \approx q^2 L / 2^{128}$$

for message lengths $\leq L$.

ϵ is at least $q(q - 1) / 2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

δ is still unacceptably large.

(Show that this is tight? See,
e.g., 2005 Ferguson GCM attack.)

2006 Joux “forbidden attack”:
nonce in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:
 $AES_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$
“seems a safe option”. (Also
suggested and analyzed in, e.g.,
2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Forgery chance $\leq \delta + \epsilon$ where
 ϵ is AES PRF insecurity and
 $\delta \approx q^2 L / 2^{128}$
for message lengths $\leq L$.

ϵ is at least $q(q - 1) / 2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

δ is still unacceptably large.

(Show that this is tight? See,
e.g., 2005 Ferguson GCM attack.)

Fragile solution: “Switch keys!”

2006 Joux “forbidden attack”:
repeated in GCM \Rightarrow repeated s_n
 \Rightarrow attacker figures out r ,
can easily forge messages.

Joux’s suggested response:
 $AES_k(r^4 m_1 + r^3 m_2 + r^2 m_3 + r m_4)$
“seems a safe option”. (Also
suggested and analyzed in, e.g.,
2000 Bernstein; earlier refs?)

Is this 2^{128} “security”?

Forgery chance $\leq \delta + \epsilon$ where
 ϵ is AES PRF insecurity and
 $\delta \approx q^2 L / 2^{128}$
for message lengths $\leq L$.

ϵ is at least $q(q - 1) / 2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

δ is still unacceptably large.

(Show that this is tight? See,
e.g., 2005 Ferguson GCM attack.)

Fragile solution: “Switch keys!”

Much simpler: 256-bit blocks.

2014 Bernstein–Chou “Auth256”:

29 bit ops/message bit for
differential probability $< 2^{-255}$.

Or try EHC from 2013 Nandi?

aux “forbidden attack”:
in GCM \Rightarrow repeated s_n
 attacker figures out r ,
 can forge messages.

suggested response:

$$m_1 + r^3 m_2 + r^2 m_3 + r m_4$$

“a safe option”. (Also

discussed and analyzed in, e.g.,
Bernstein; earlier refs?)

2^{128} “security”?

chance $\leq \delta + \epsilon$ where

δ = PRF insecurity and

$$\epsilon = \frac{L}{2^{128}}$$

message lengths $\leq L$.

ϵ is at least $q(q - 1)/2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

δ is still unacceptably large.

(Show that this is tight? See,
e.g., 2005 Ferguson GCM attack.)

Fragile solution: “Switch keys!”

Much simpler: 256-bit blocks.

2014 Bernstein–Chou “Auth256”:

29 bit ops/message bit for
differential probability $< 2^{-255}$.

Or try EHC from 2013 Nandi?

Improving

Tor wants
easy-to-implement

encumbers

509-byte

(But current

so can be

Also: see

from each

Tor is composed

of AEZ

See, e.g.

from RV

den attack”:

repeated s_n

s out r ,

essages.

response:

$m_2 + r^2 m_3 + r m_4$)

on”. (Also

alyzed in, e.g.,

earlier refs?)

ity”?

$\delta + \epsilon$ where

curity and

$s \leq L$.

ϵ is at least $q(q - 1)/2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

δ is still unacceptably large.

(Show that this is tight? See,
e.g., 2005 Ferguson GCM attack.)

Fragile solution: “Switch keys!”

Much simpler: 256-bit blocks.

2014 Bernstein–Chou “Auth256”:

29 bit ops/message bit for
differential probability $< 2^{-255}$.

Or try EHC from 2013 Nandi?

Improving Tor

Tor wants “fast, p
easy-to-implement

encumbered, side-

509-byte blooock

(But current ciphers

so can consider co

Also: secure chain

from each blooock

Tor is considering

of AEZ or HHFHF

See, e.g., Mathew.

from RWC 2013 a

ϵ is at least $q(q - 1)/2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

δ is still unacceptably large.

(Show that this is tight? See,
e.g., 2005 Ferguson GCM attack.)

Fragile solution: "Switch keys!"

Much simpler: 256-bit blocks.

2014 Bernstein–Chou "Auth256":

29 bit ops/message bit for
differential probability $< 2^{-255}$.

Or try EHC from 2013 Nandi?

Improving Tor

Tor wants "fast, proven, secure,
easy-to-implement, non-patented,
unencumbered, side-channel-free"
509-byte block cipher.

(But current cipher is a disaster,
so can consider compromises)

Also: secure chaining
from each block to the next

Tor is considering deployment
of AEZ or HHFHFH in 2016

See, e.g., Mathewson talks
from RWC 2013 and RWC 2014

ϵ is at least $q(q - 1)/2^{129}$.

Solution: better PRP/PRF switch
(2005 Bernstein), ok for $q \approx 2^{64}$.

δ is still unacceptably large.

(Show that this is tight? See,
e.g., 2005 Ferguson GCM attack.)

Fragile solution: “Switch keys!”

Much simpler: 256-bit blocks.

2014 Bernstein–Chou “Auth256”:
29 bit ops/message bit for
differential probability $< 2^{-255}$.

Or try EHC from 2013 Nandi?

Improving Tor

Tor wants “fast, proven, secure,
easy-to-implement, non-patent-
encumbered, side-channel-free”
509-byte block cipher.

(But current cipher is a disaster,
so can consider compromises.)

Also: secure chaining
from each block to the next.

Tor is considering deployment
of AEZ or HHFHFH in 2016.

See, e.g., Mathewson talks
from RWC 2013 and RWC 2016.

at least $q(q - 1)/2^{129}$.

: better PRP/PRF switch
(Bernstein), ok for $q \approx 2^{64}$.

unacceptably large.

What is this tight? See,

(2015 Ferguson GCM attack.)

Solution: "Switch keys!"

Block cipher: 256-bit blocks.

Bernstein–Chou "Auth256":

128 bits/message bit for

failure probability $< 2^{-255}$.

Auth256 from 2013 Nandi?

Improving Tor

Tor wants "fast, proven, secure, easy-to-implement, non-patent-encumbered, side-channel-free" 509-byte block cipher.

(But current cipher is a disaster, so can consider compromises.)

Also: secure chaining from each block to the next.

Tor is considering deployment of AEZ or HHFHFH in 2016.

See, e.g., Mathewson talks from RWC 2013 and RWC 2016.

Feis

stream
(strong

SCTE
HHFHFH

$1)/2^{129}$.
PRP/PRF switch
ok for $q \approx 2^{64}$.
ably large.
tight? See,
on GCM attack.)
"Switch keys!"
6-bit blocks.
"Auth256":
e bit for
ility $< 2^{-255}$.
2013 Nandi?

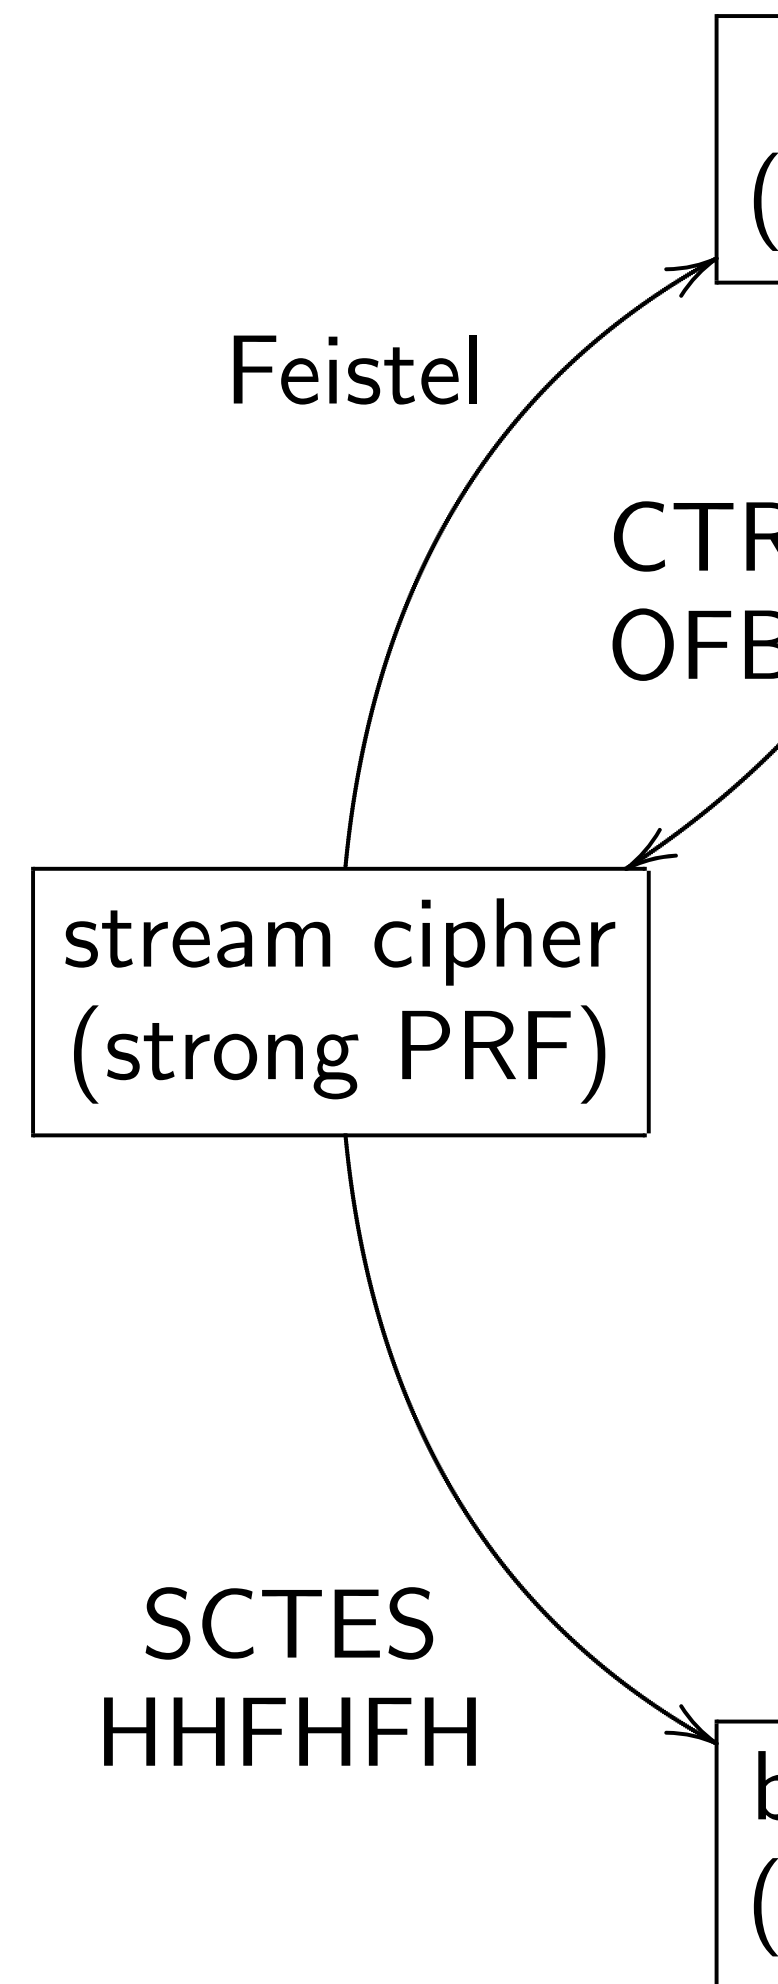
Improving Tor

Tor wants "fast, proven, secure, easy-to-implement, non-patent-encumbered, side-channel-free" 509-byte block cipher.
(But current cipher is a disaster, so can consider compromises.)

Also: secure chaining from each block to the next.

Tor is considering deployment of AEZ or HHFHFH in 2016.

See, e.g., Mathewson talks from RWC 2013 and RWC 2016.



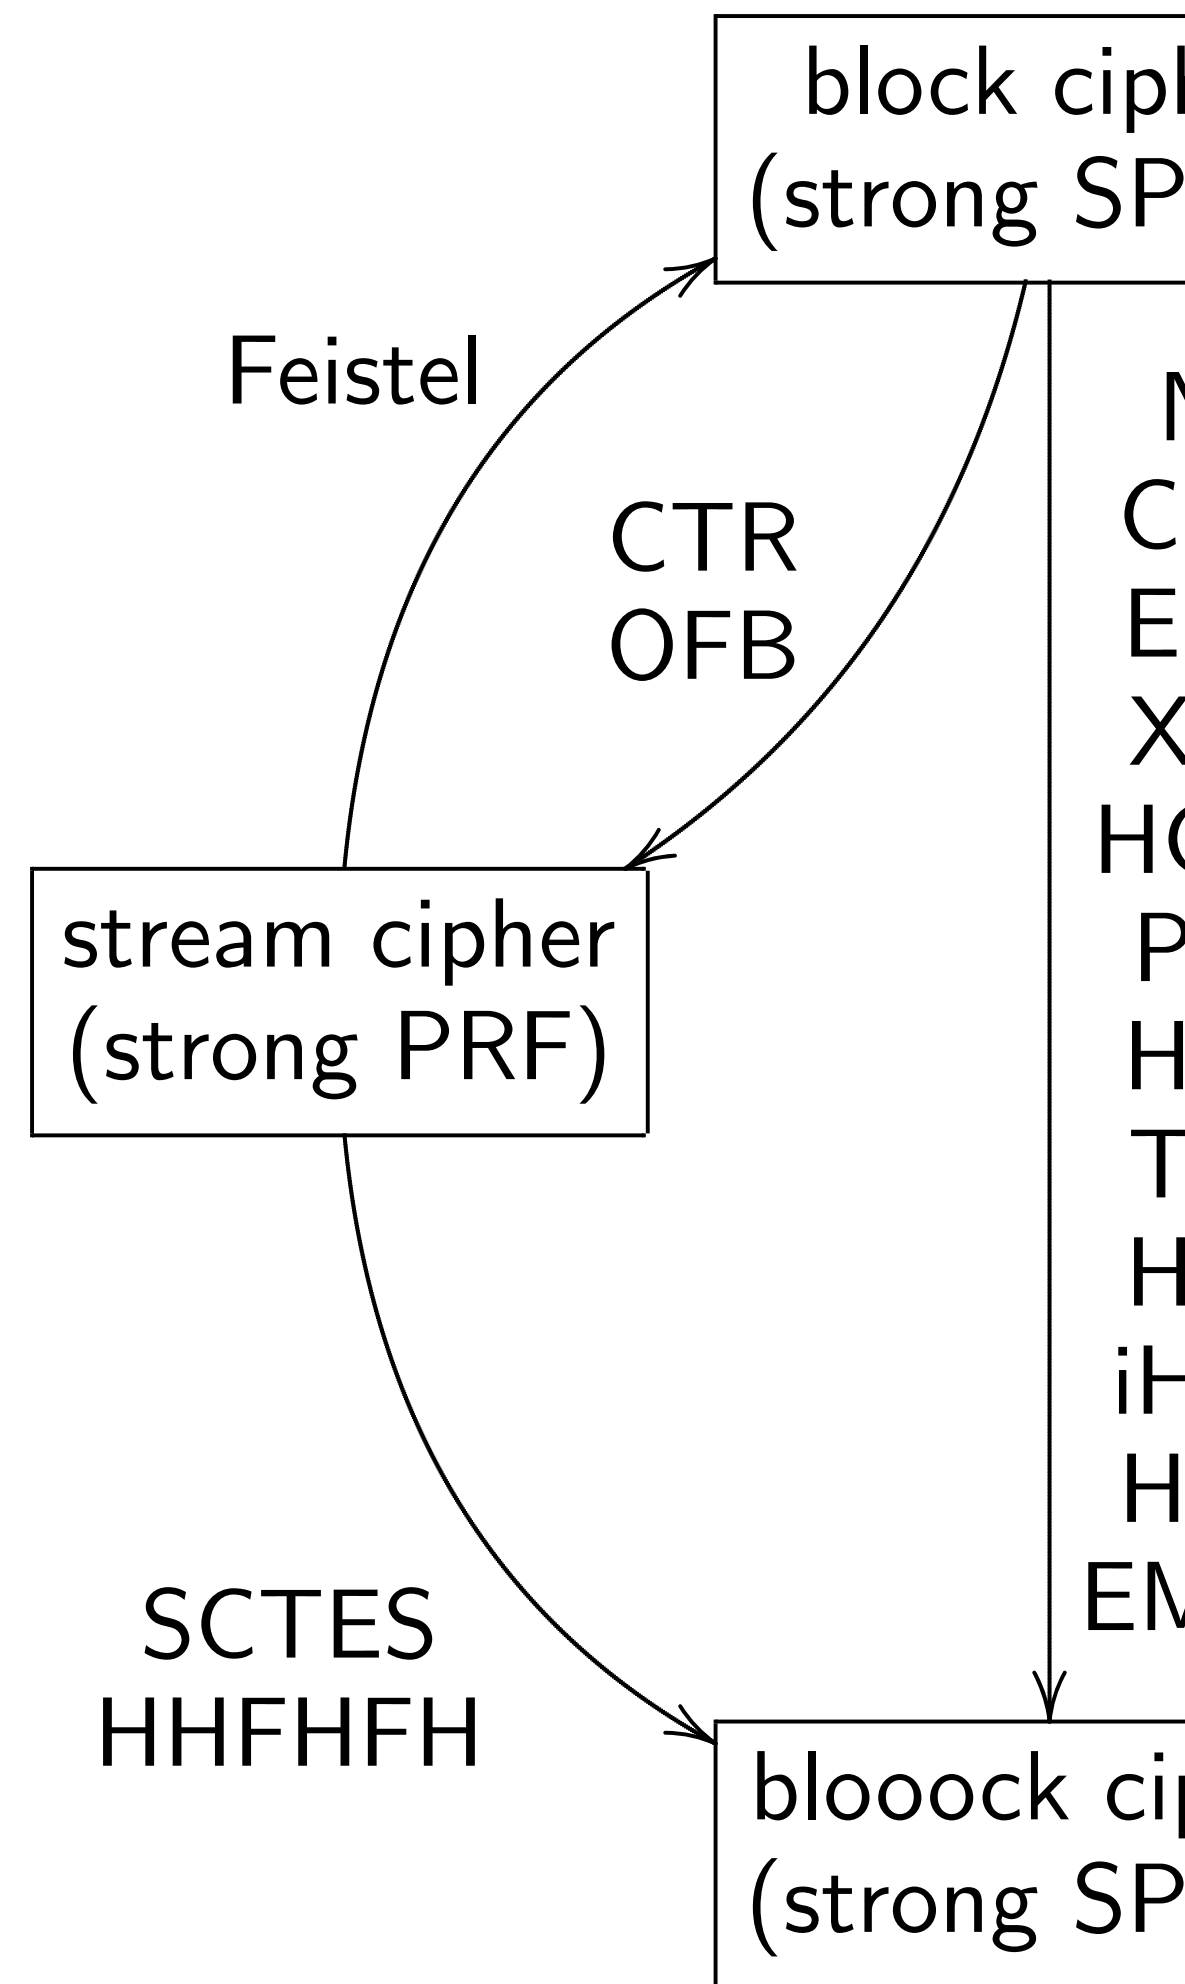
Improving Tor

Tor wants “fast, proven, secure, easy-to-implement, non-patent-encumbered, side-channel-free” 509-byte block cipher. (But current cipher is a disaster, so can consider compromises.)

Also: secure chaining from each block to the next.

Tor is considering deployment of AEZ or HHFHFH in 2016.

See, e.g., Mathewson talks from RWC 2013 and RWC 2016.



Improving Tor

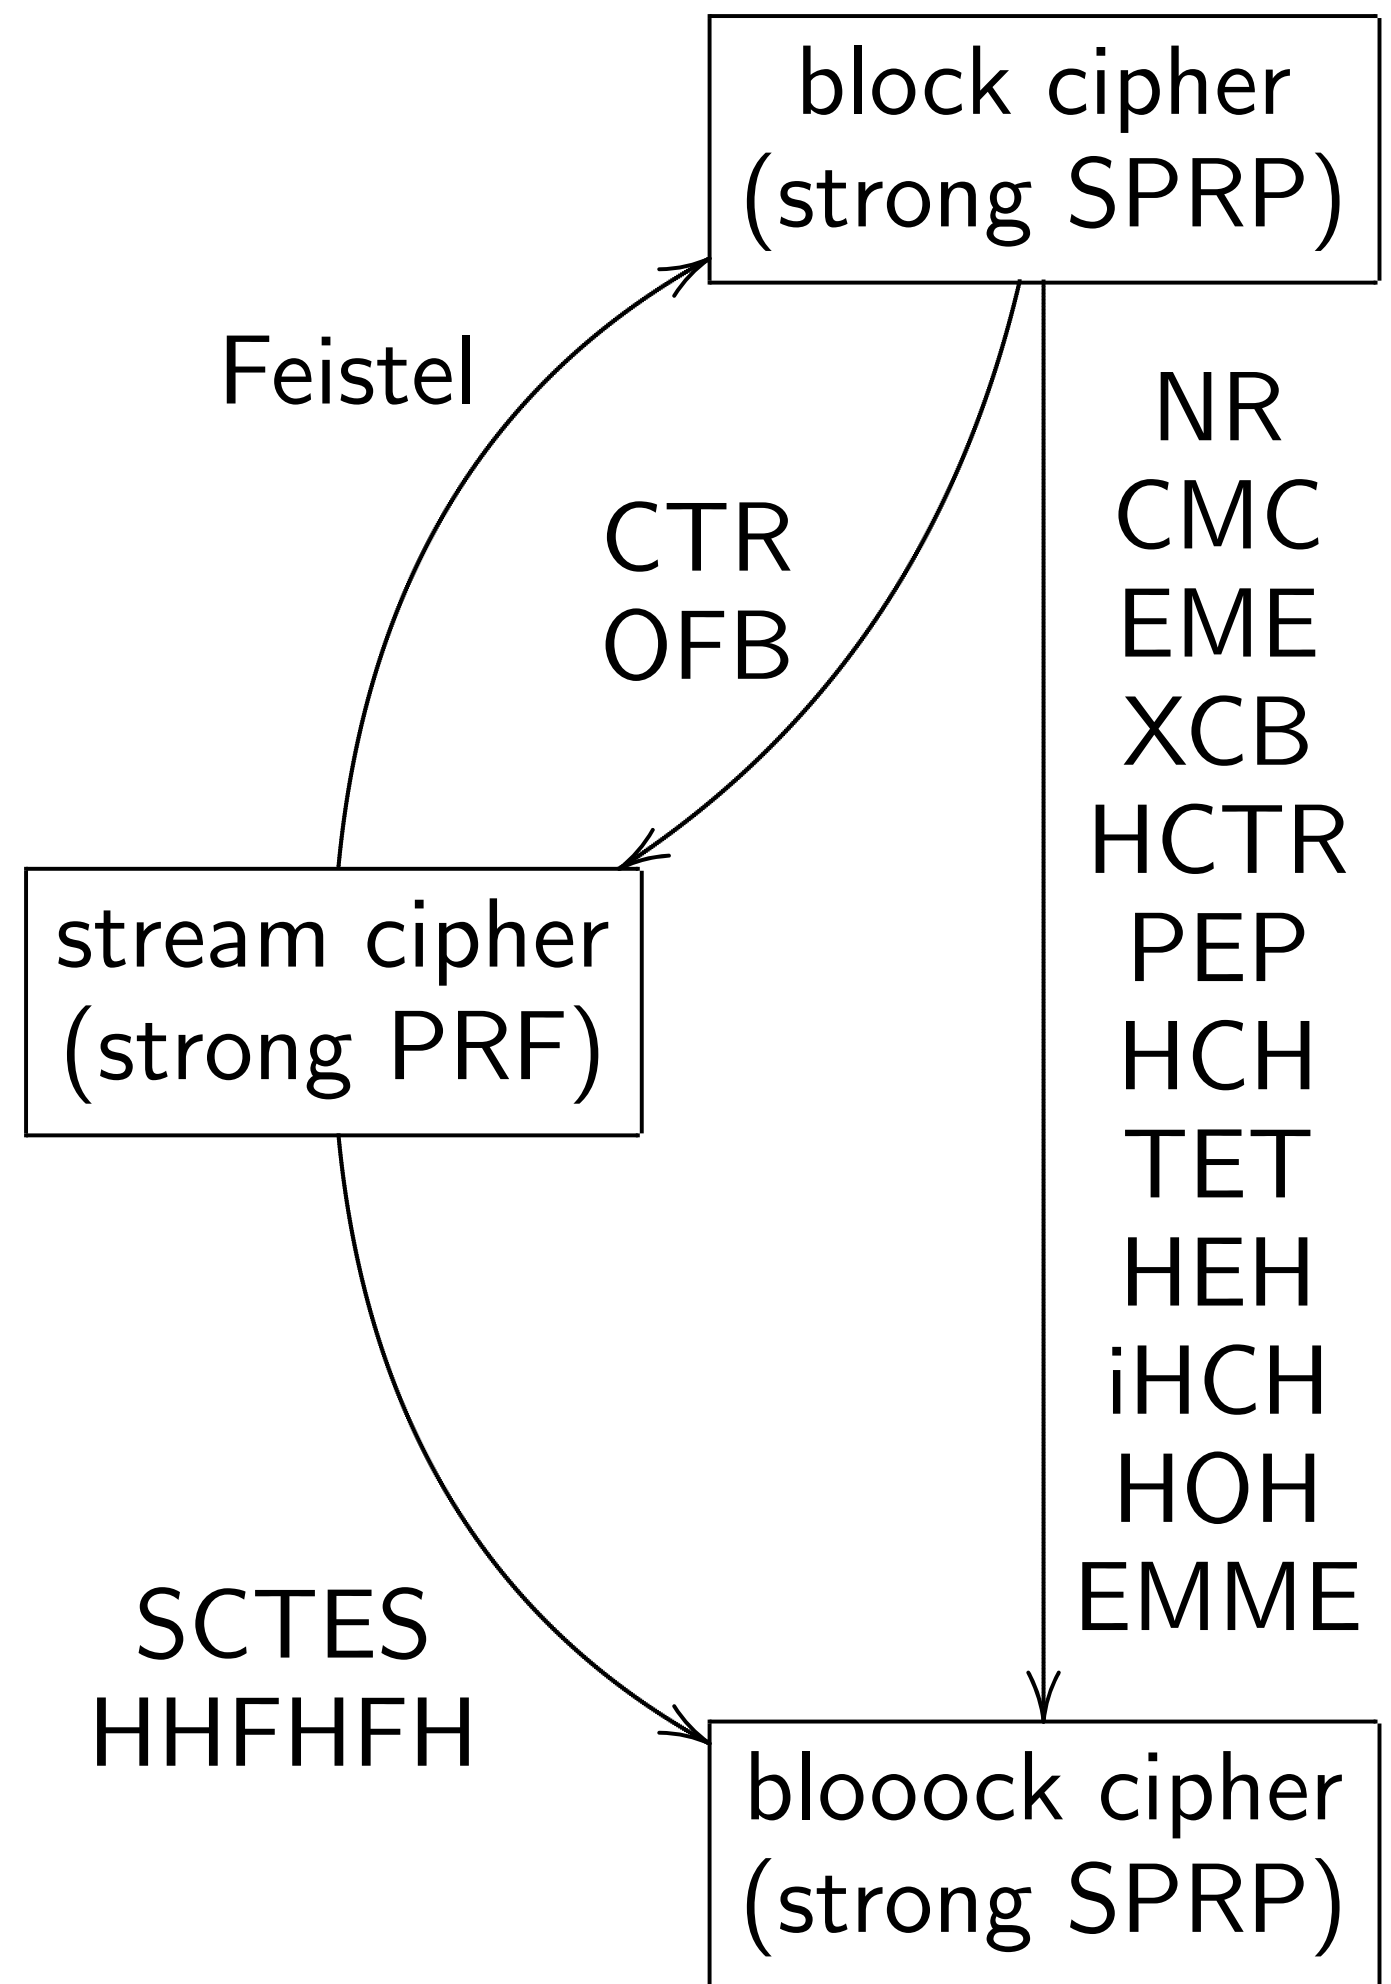
Tor wants “fast, proven, secure, easy-to-implement, non-patent-encumbered, side-channel-free” 509-byte block cipher.

(But current cipher is a disaster, so can consider compromises.)

Also: secure chaining from each block to the next.

Tor is considering deployment of AEZ or HHFHFH in 2016.

See, e.g., Mathewson talks from RWC 2013 and RWC 2016.



ng Tor

ts “fast, proven, secure,
implement, non-patent-
ered, side-channel-free”
e blooock cipher.

urrent cipher is a disaster,
onsider compromises.)

cure chaining



ch blooock to the next.

onsidering deployment

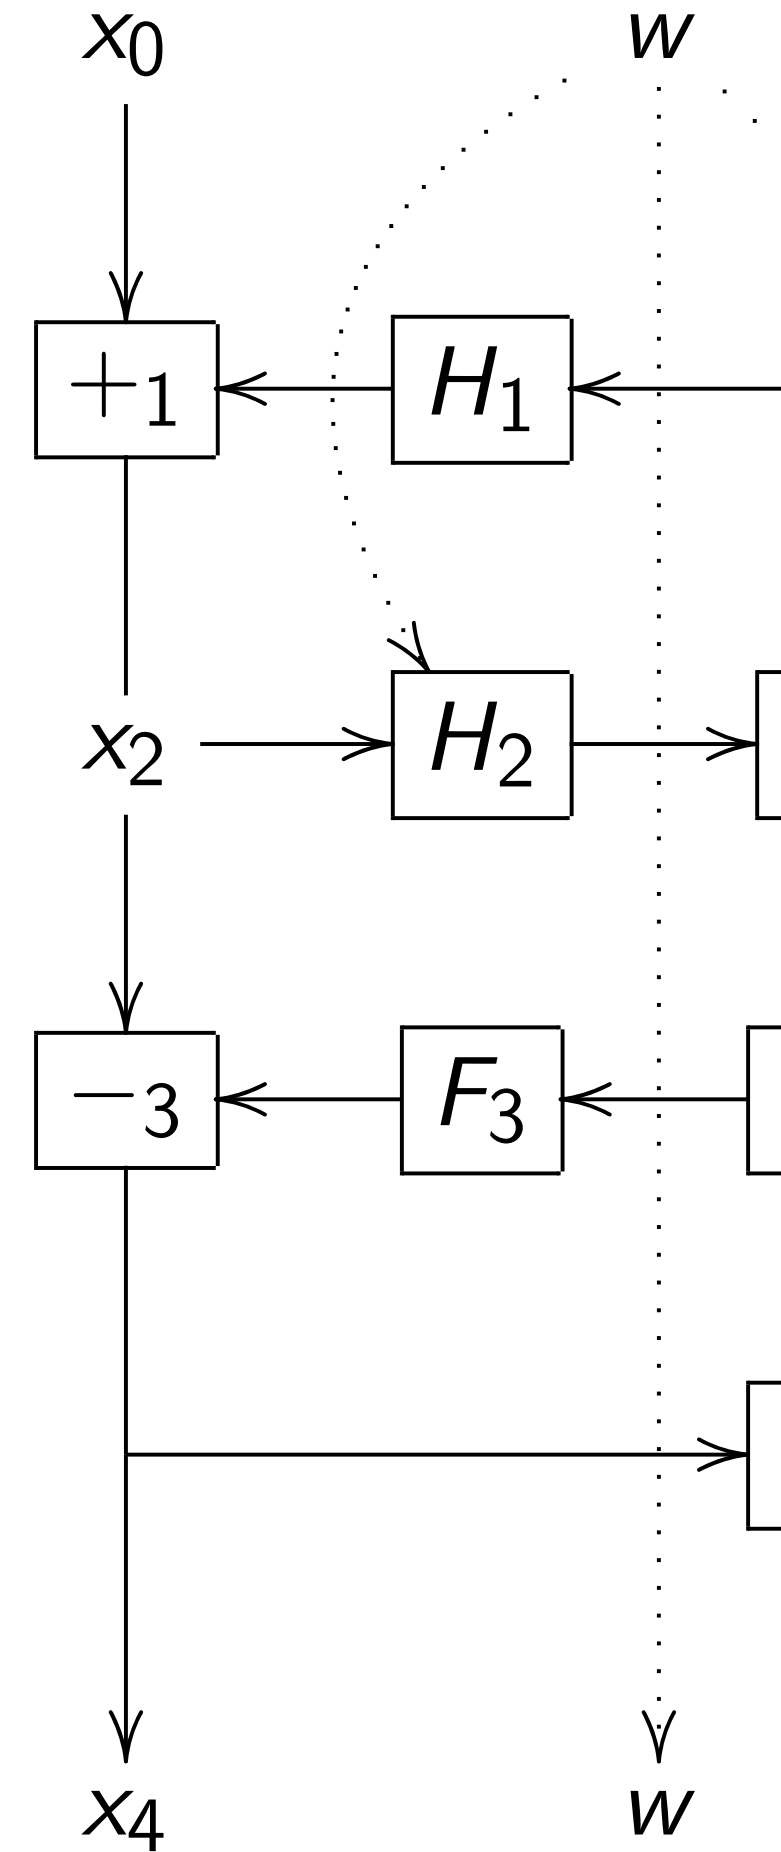

or HHFHFH in 2016.

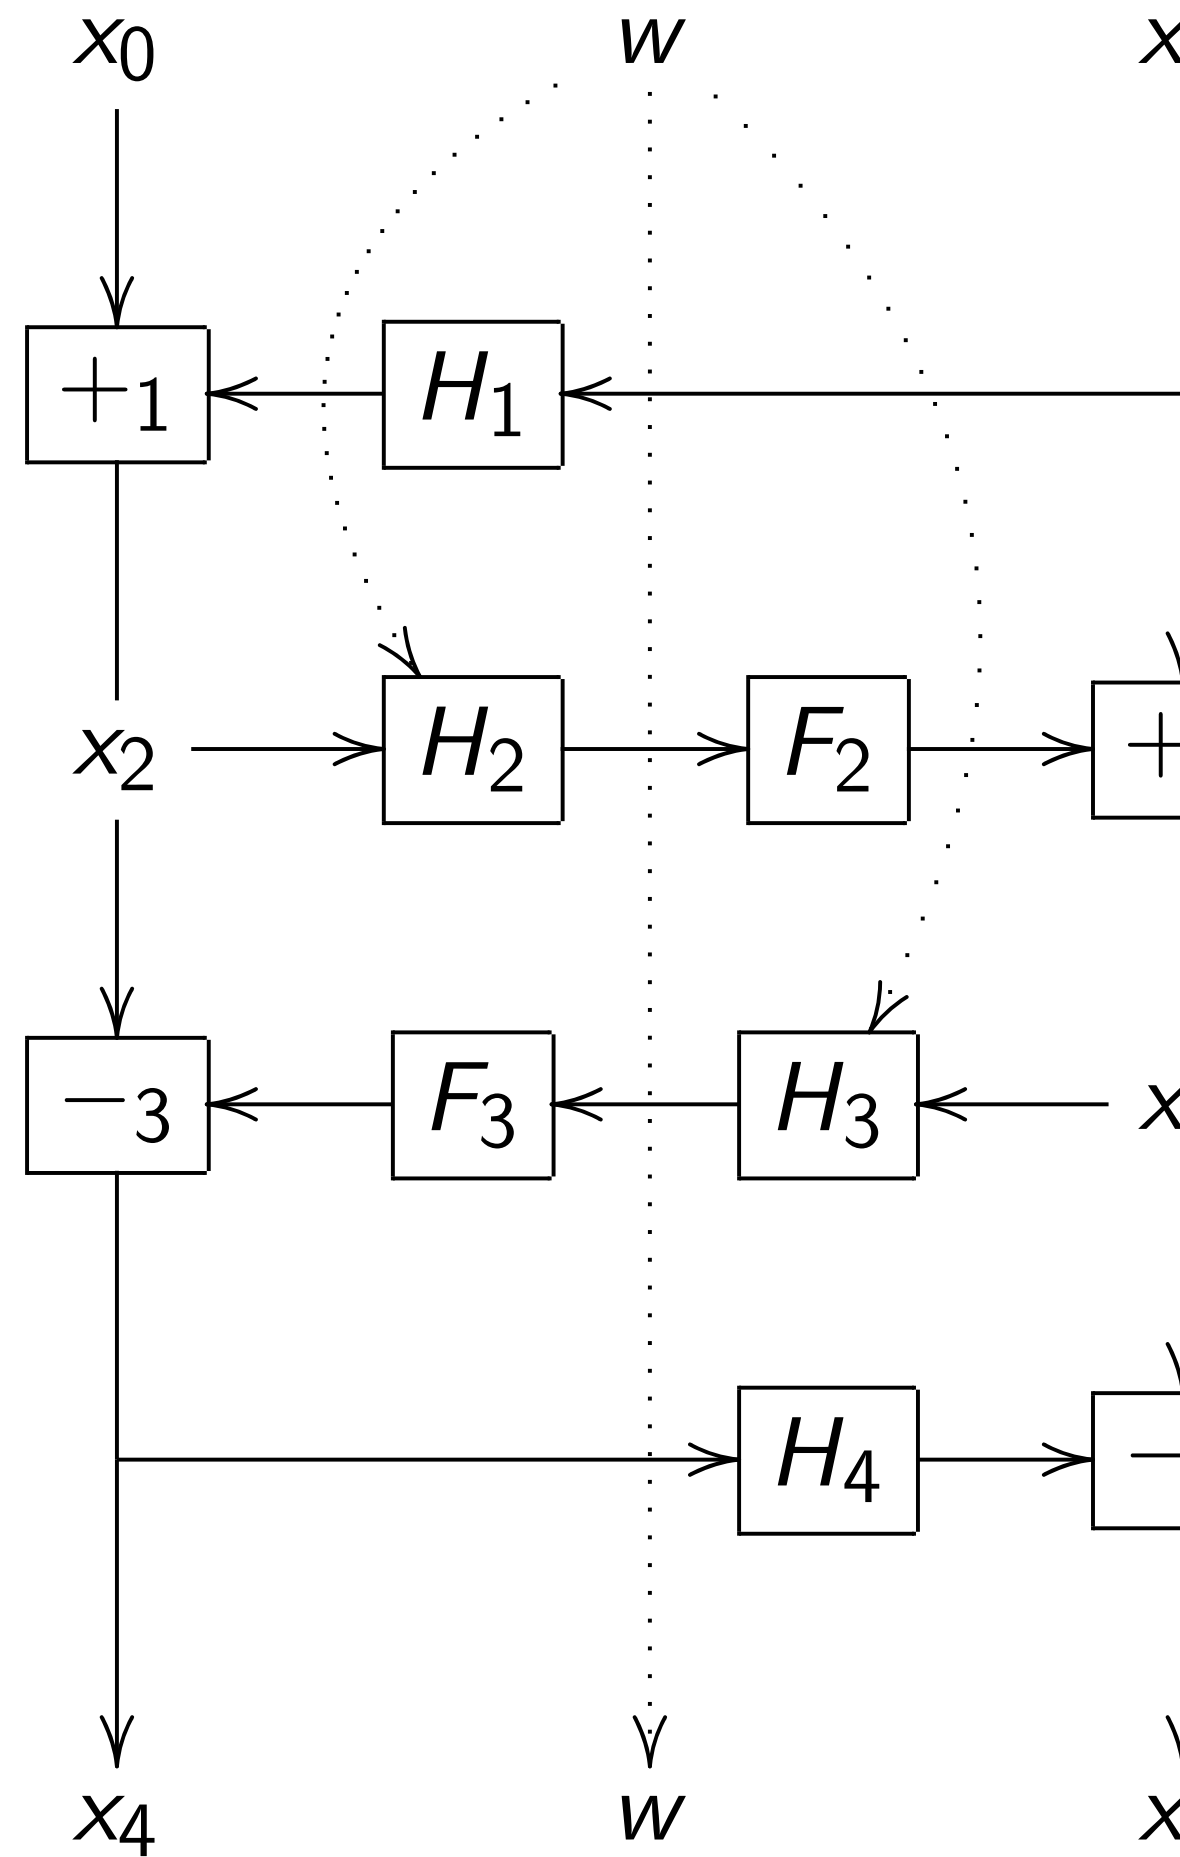
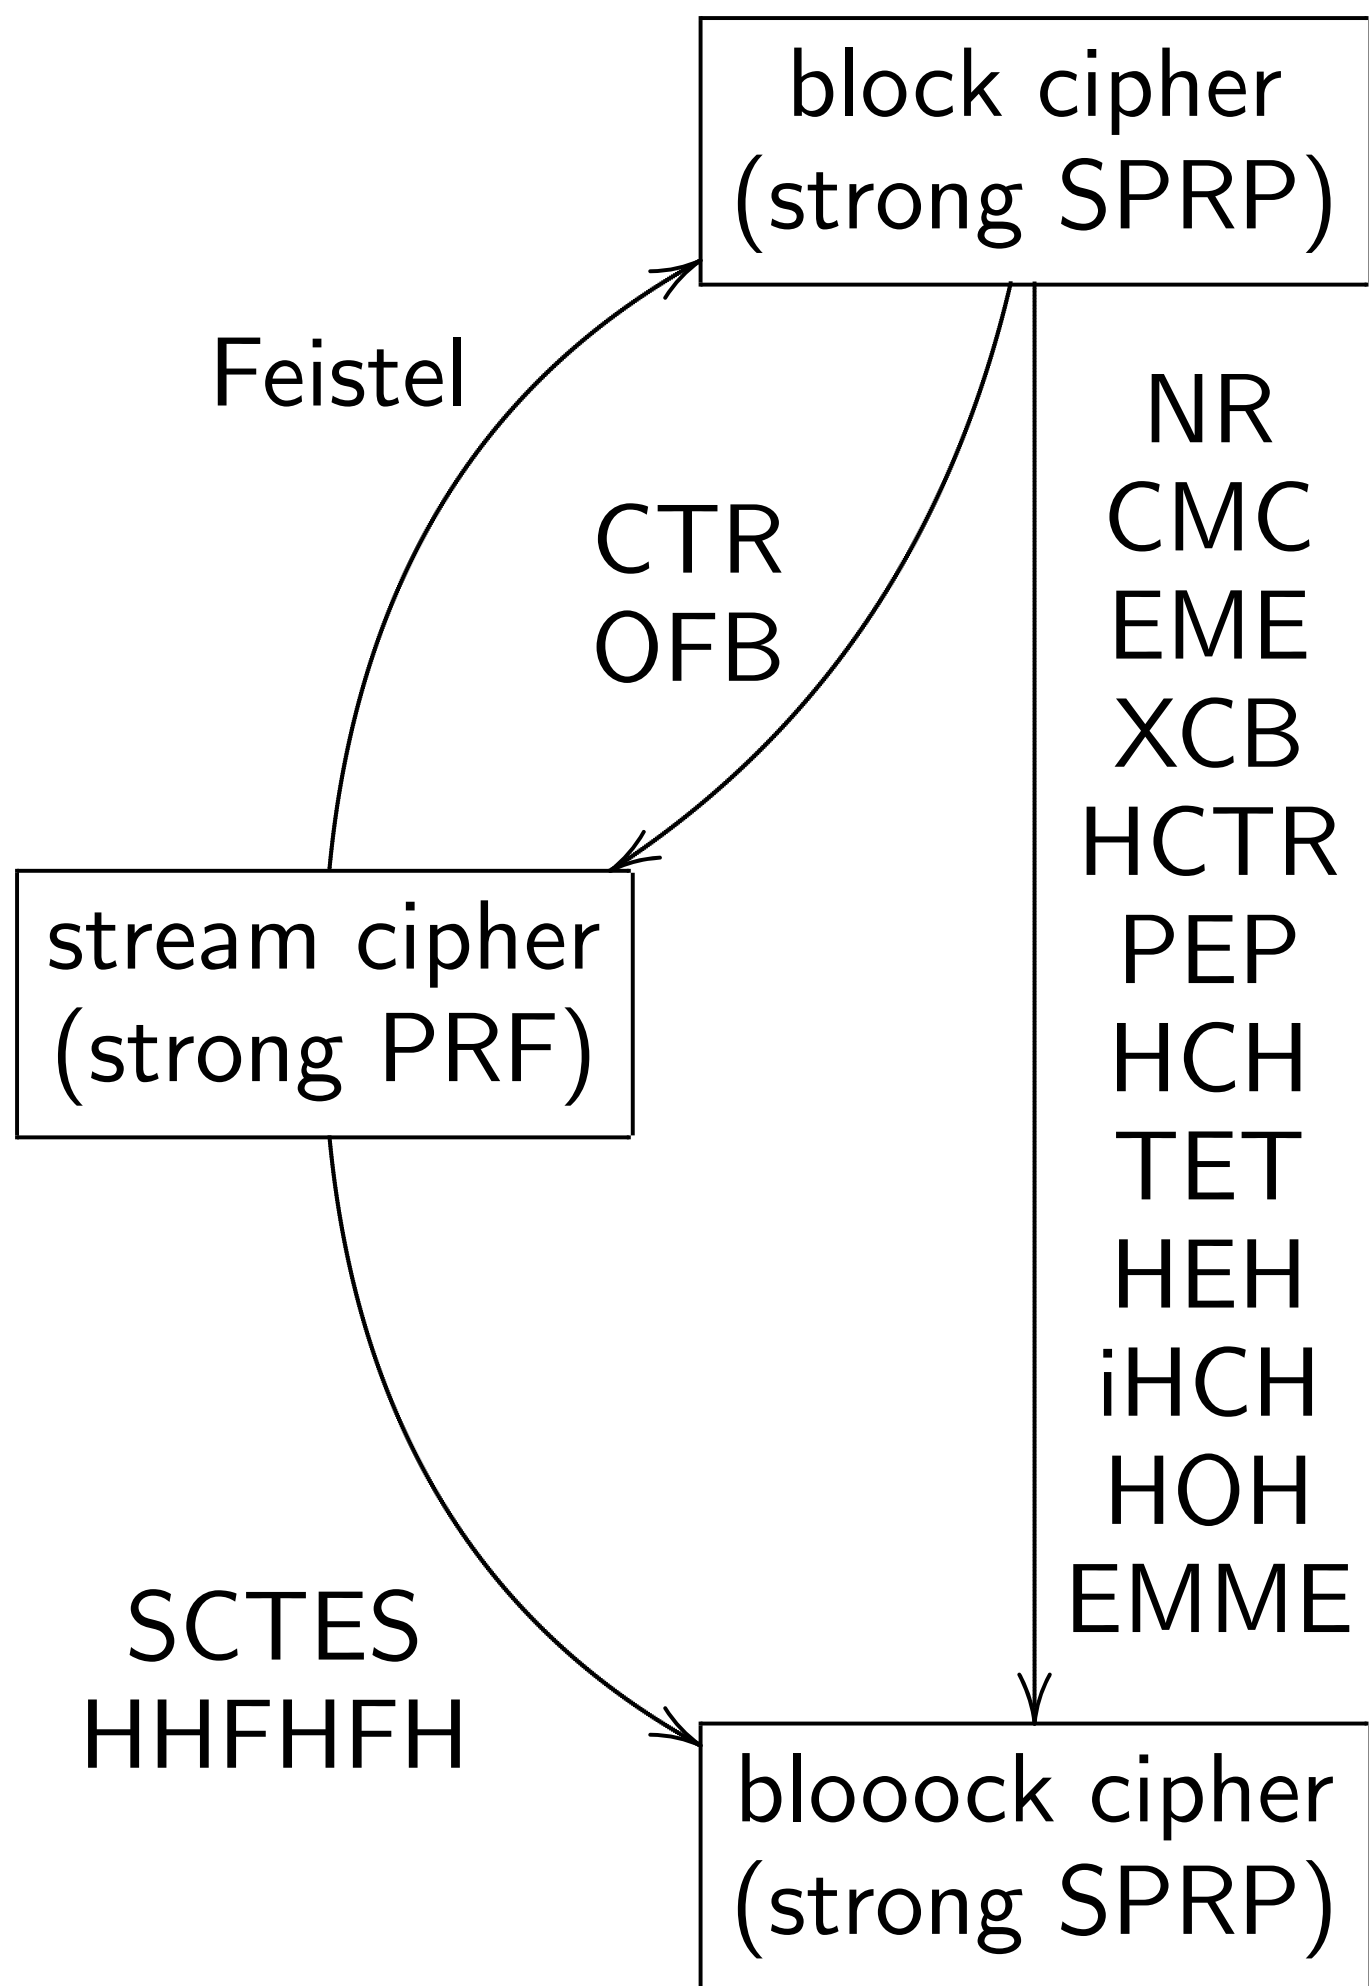
, Mathewson talks

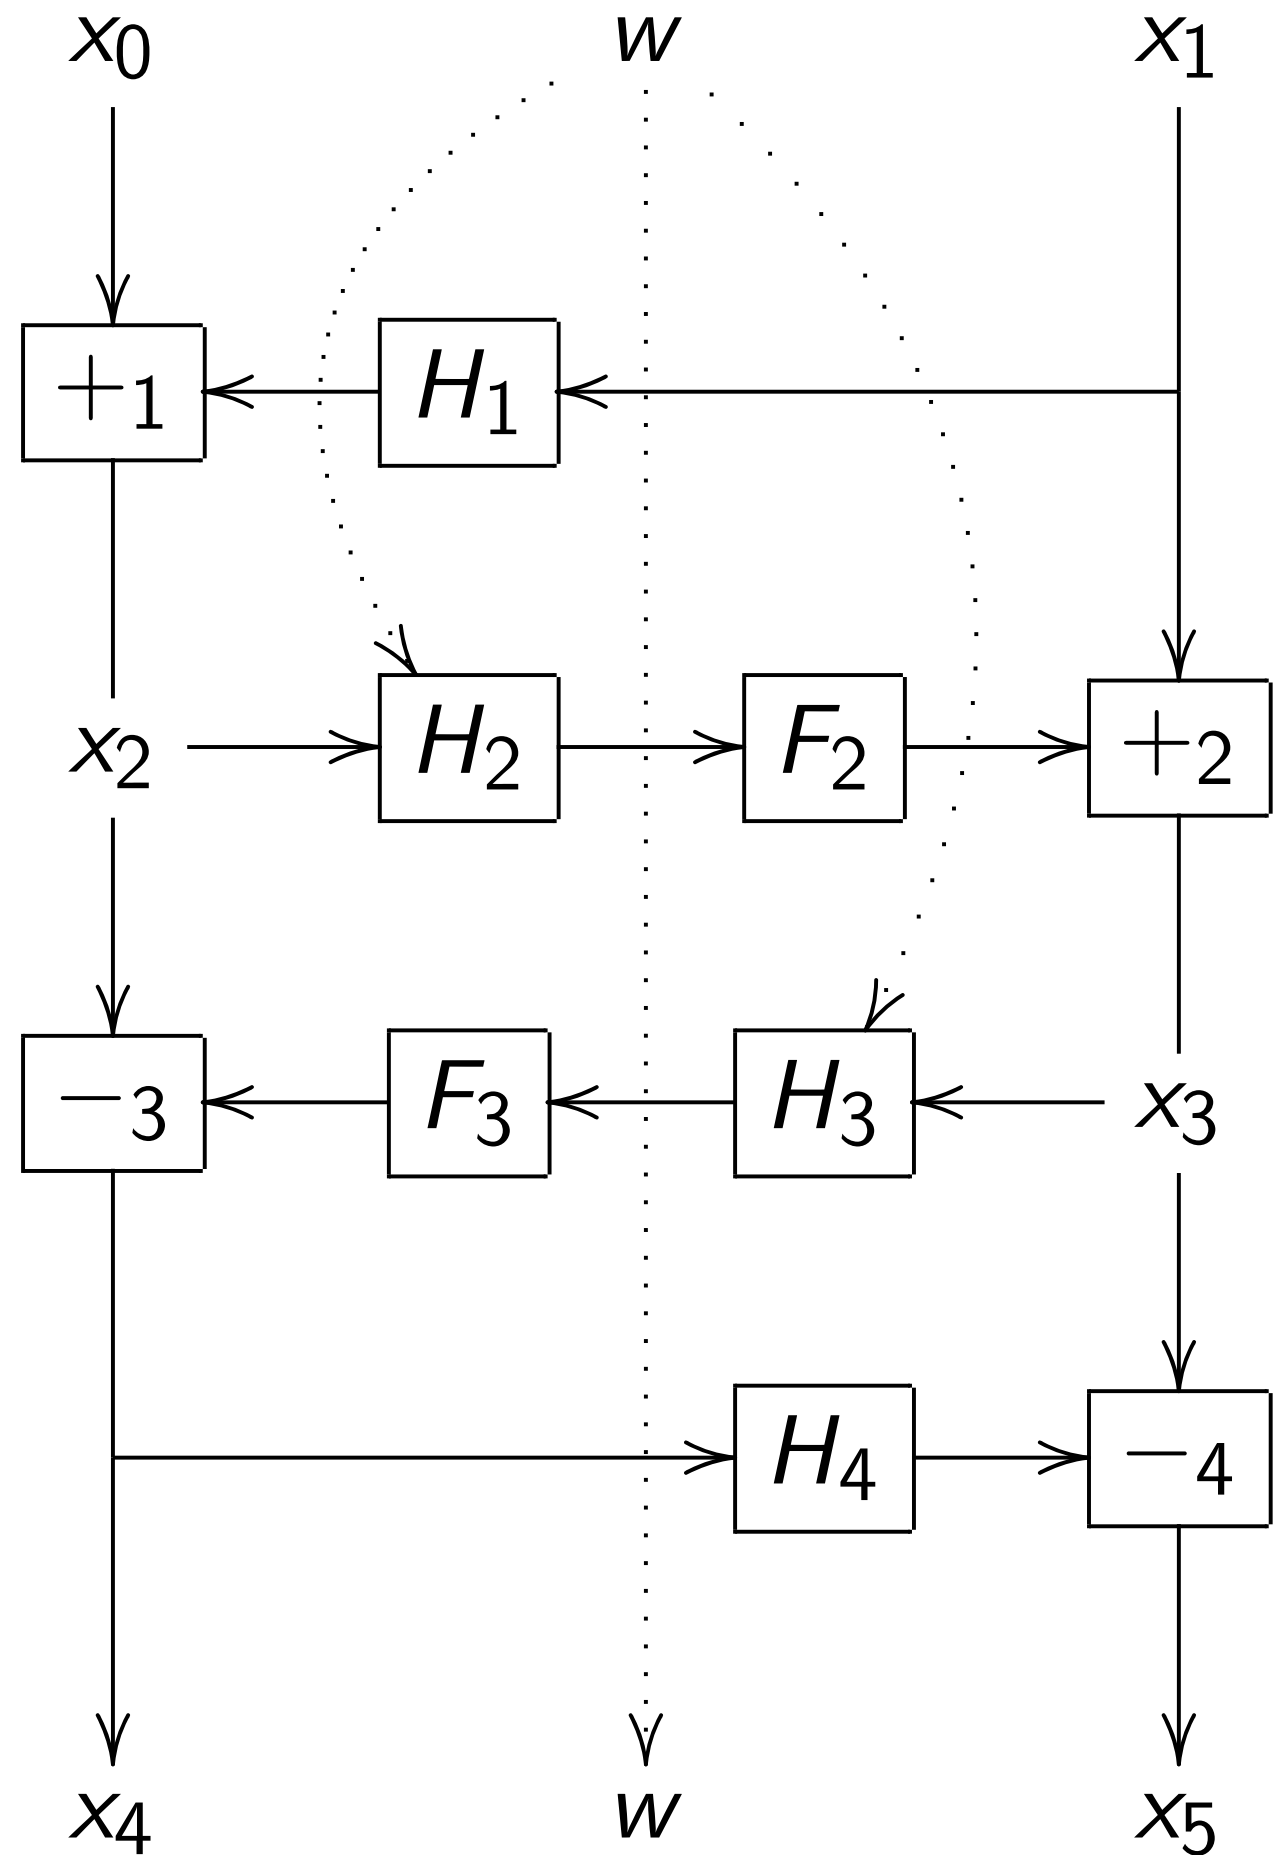
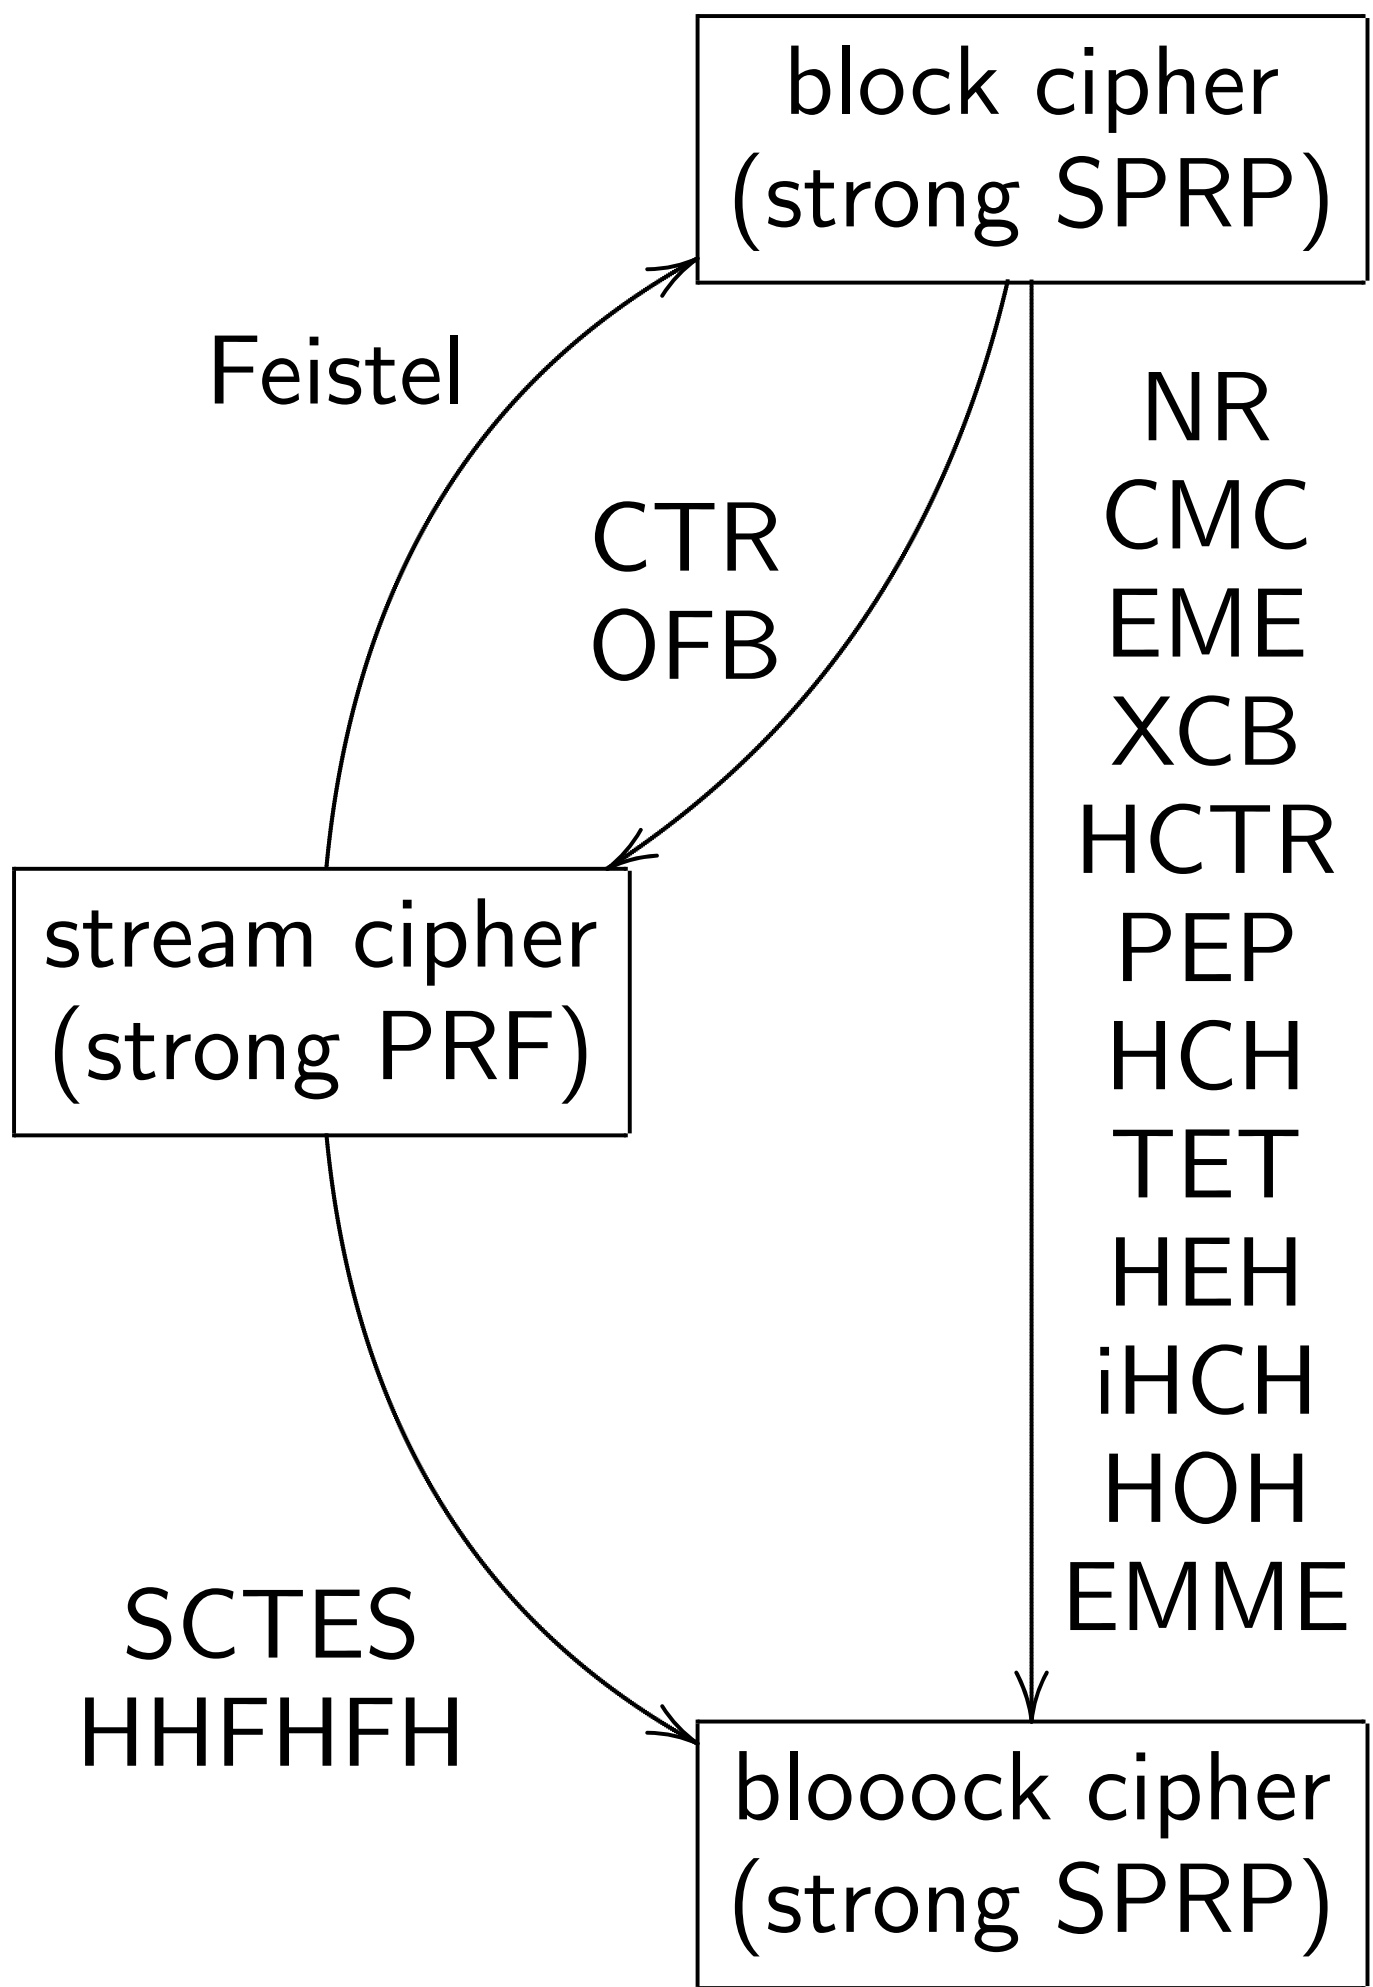
VC 2013 and RWC 2016.

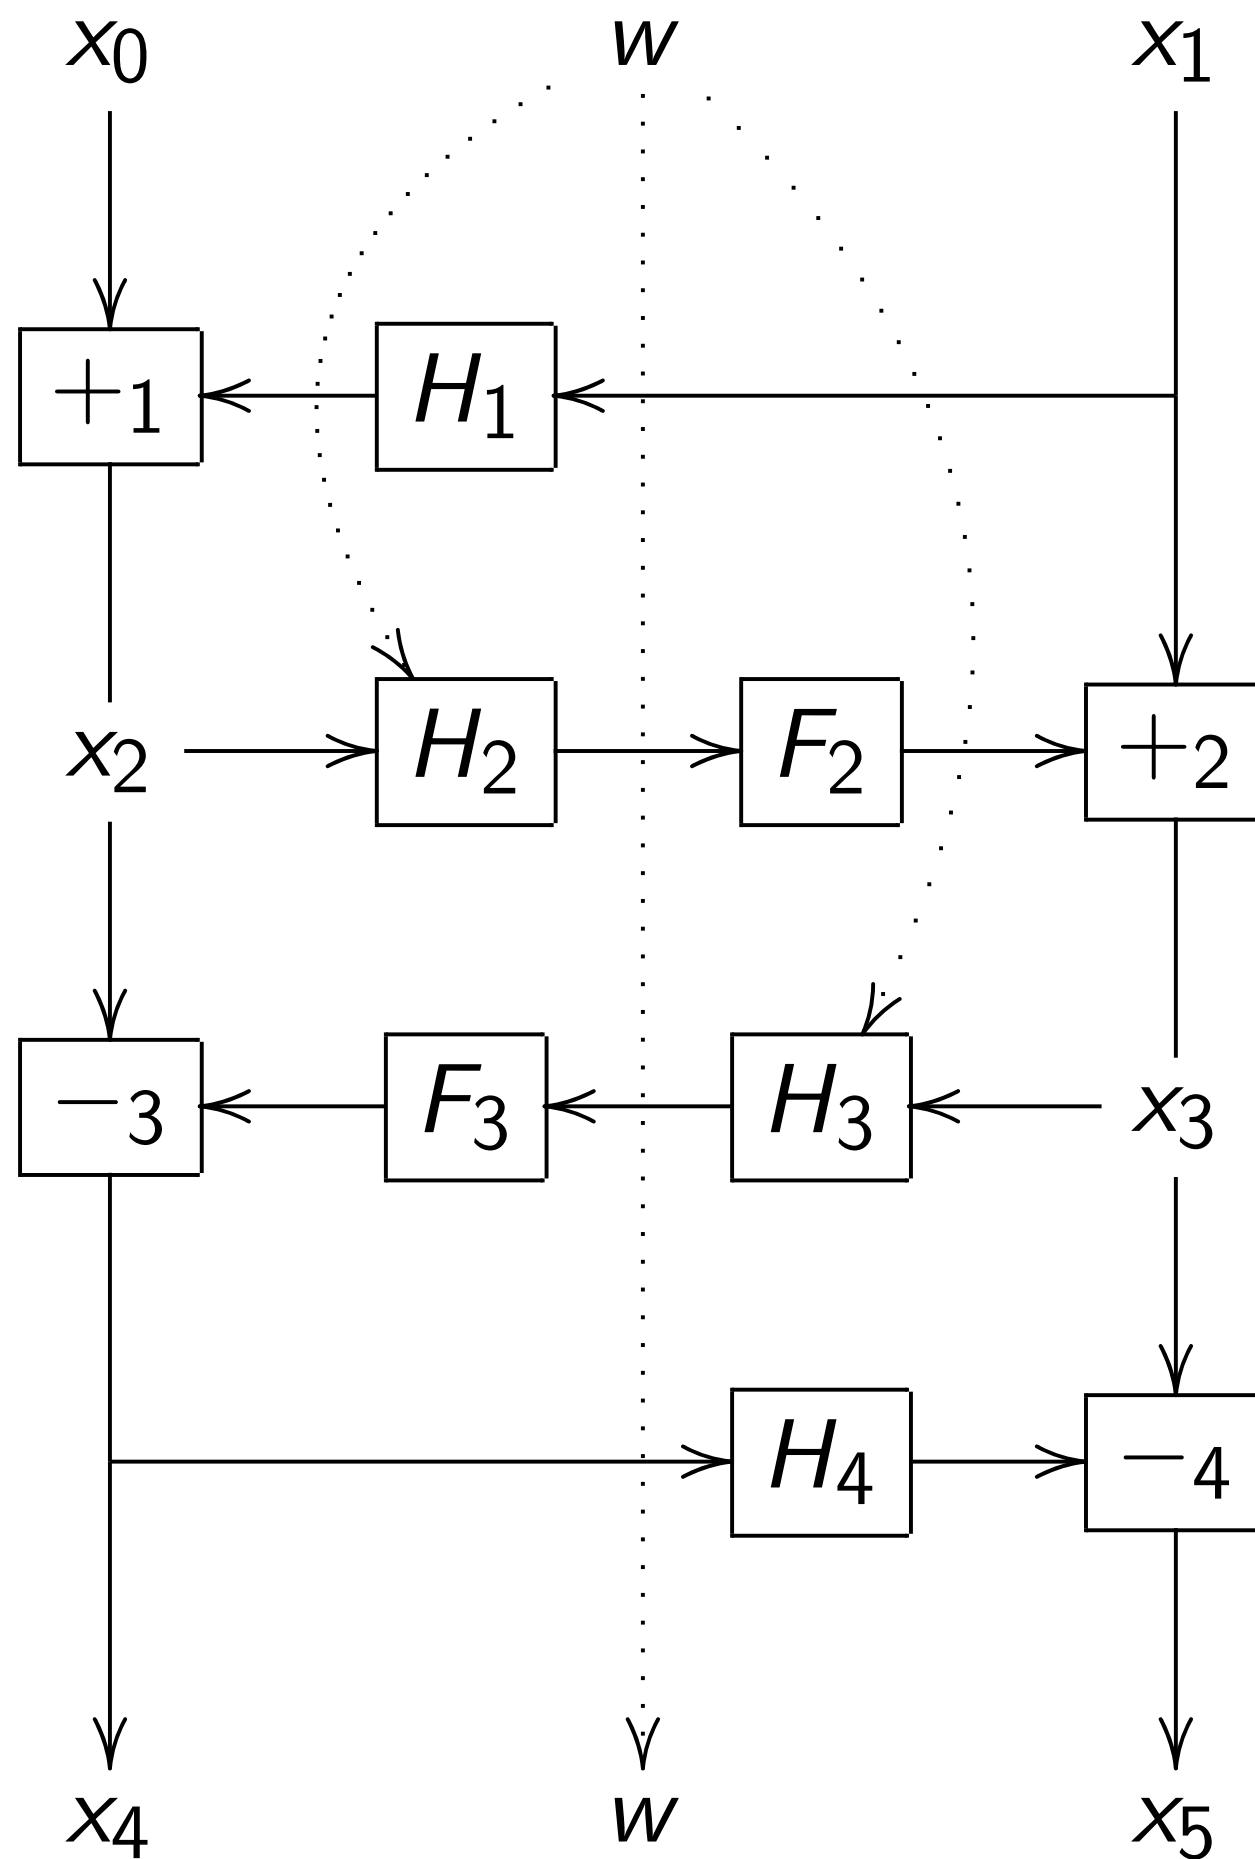
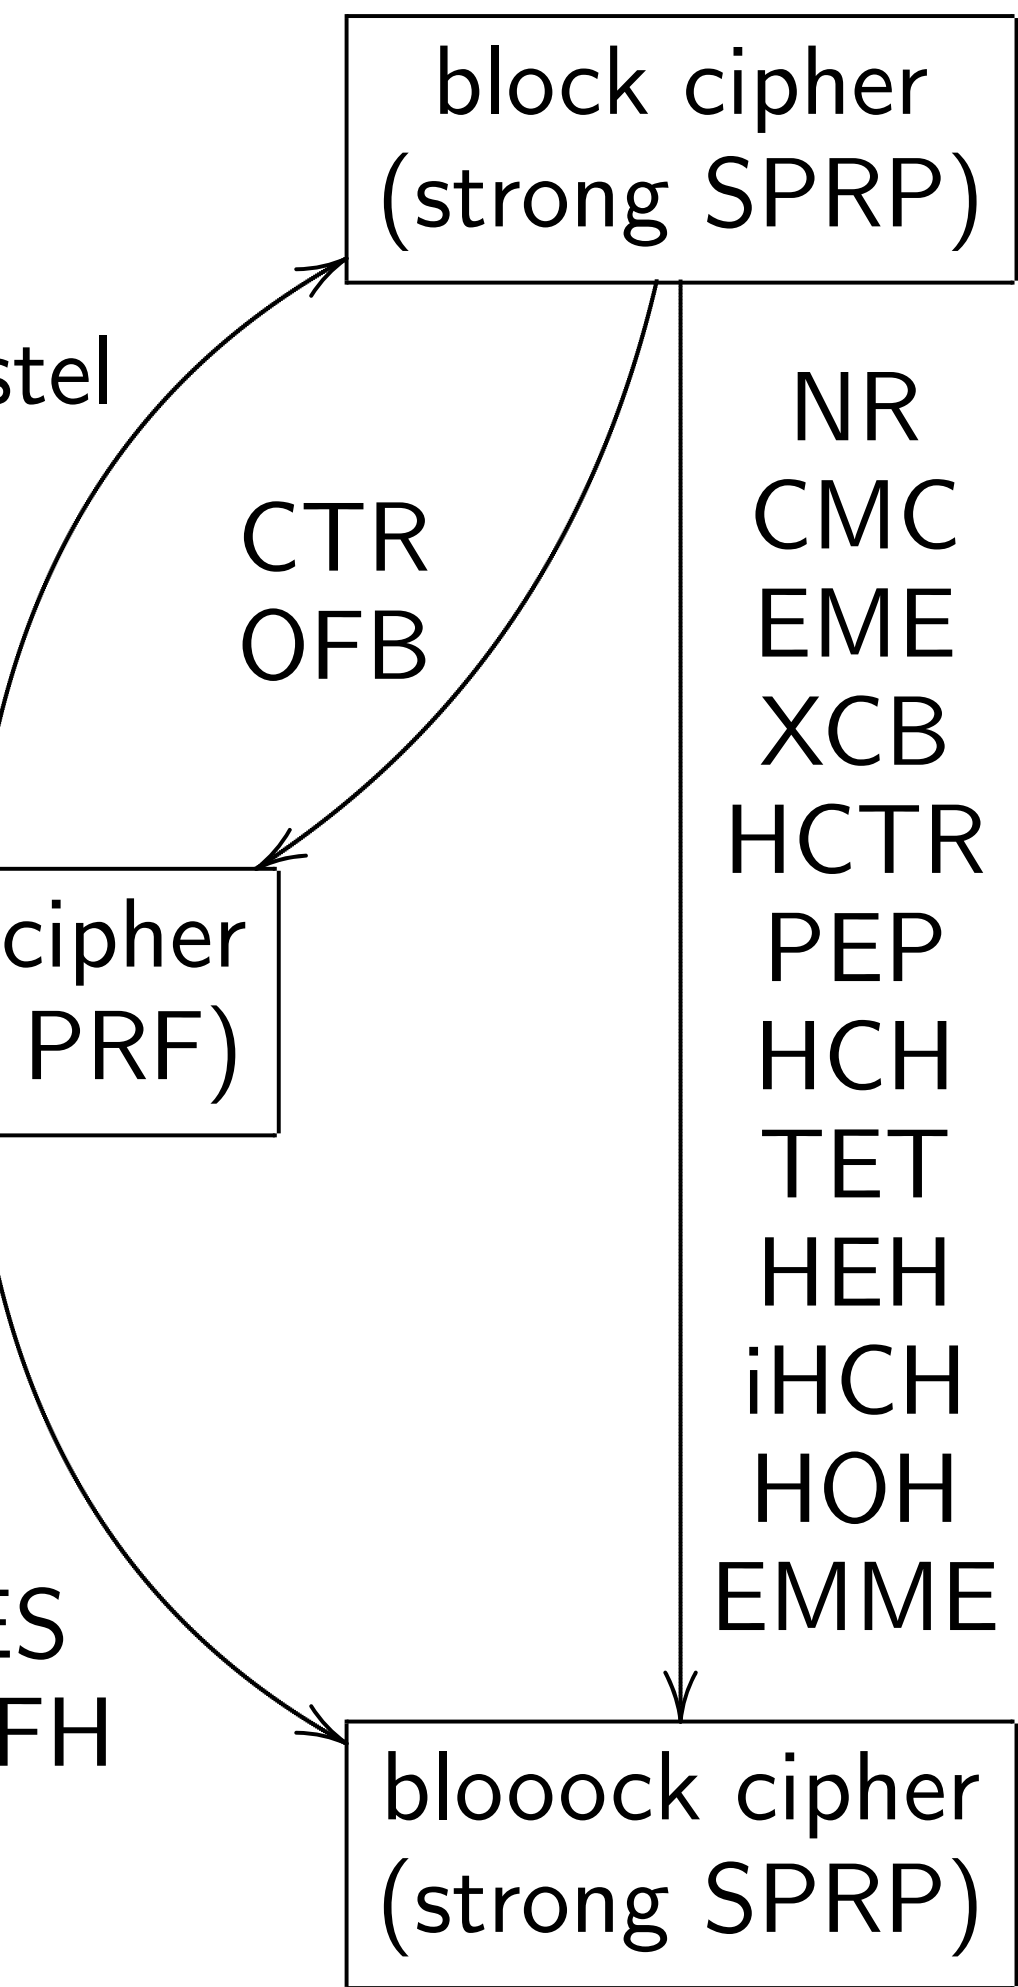


proven, secure,
 , non-patent-
 channel-free”
 cipher.
 er is a disaster,
 compromises.)
 ing
 k to the next.
 deployment
 FH in 2016.
 son talks
 nd RWC 2016.







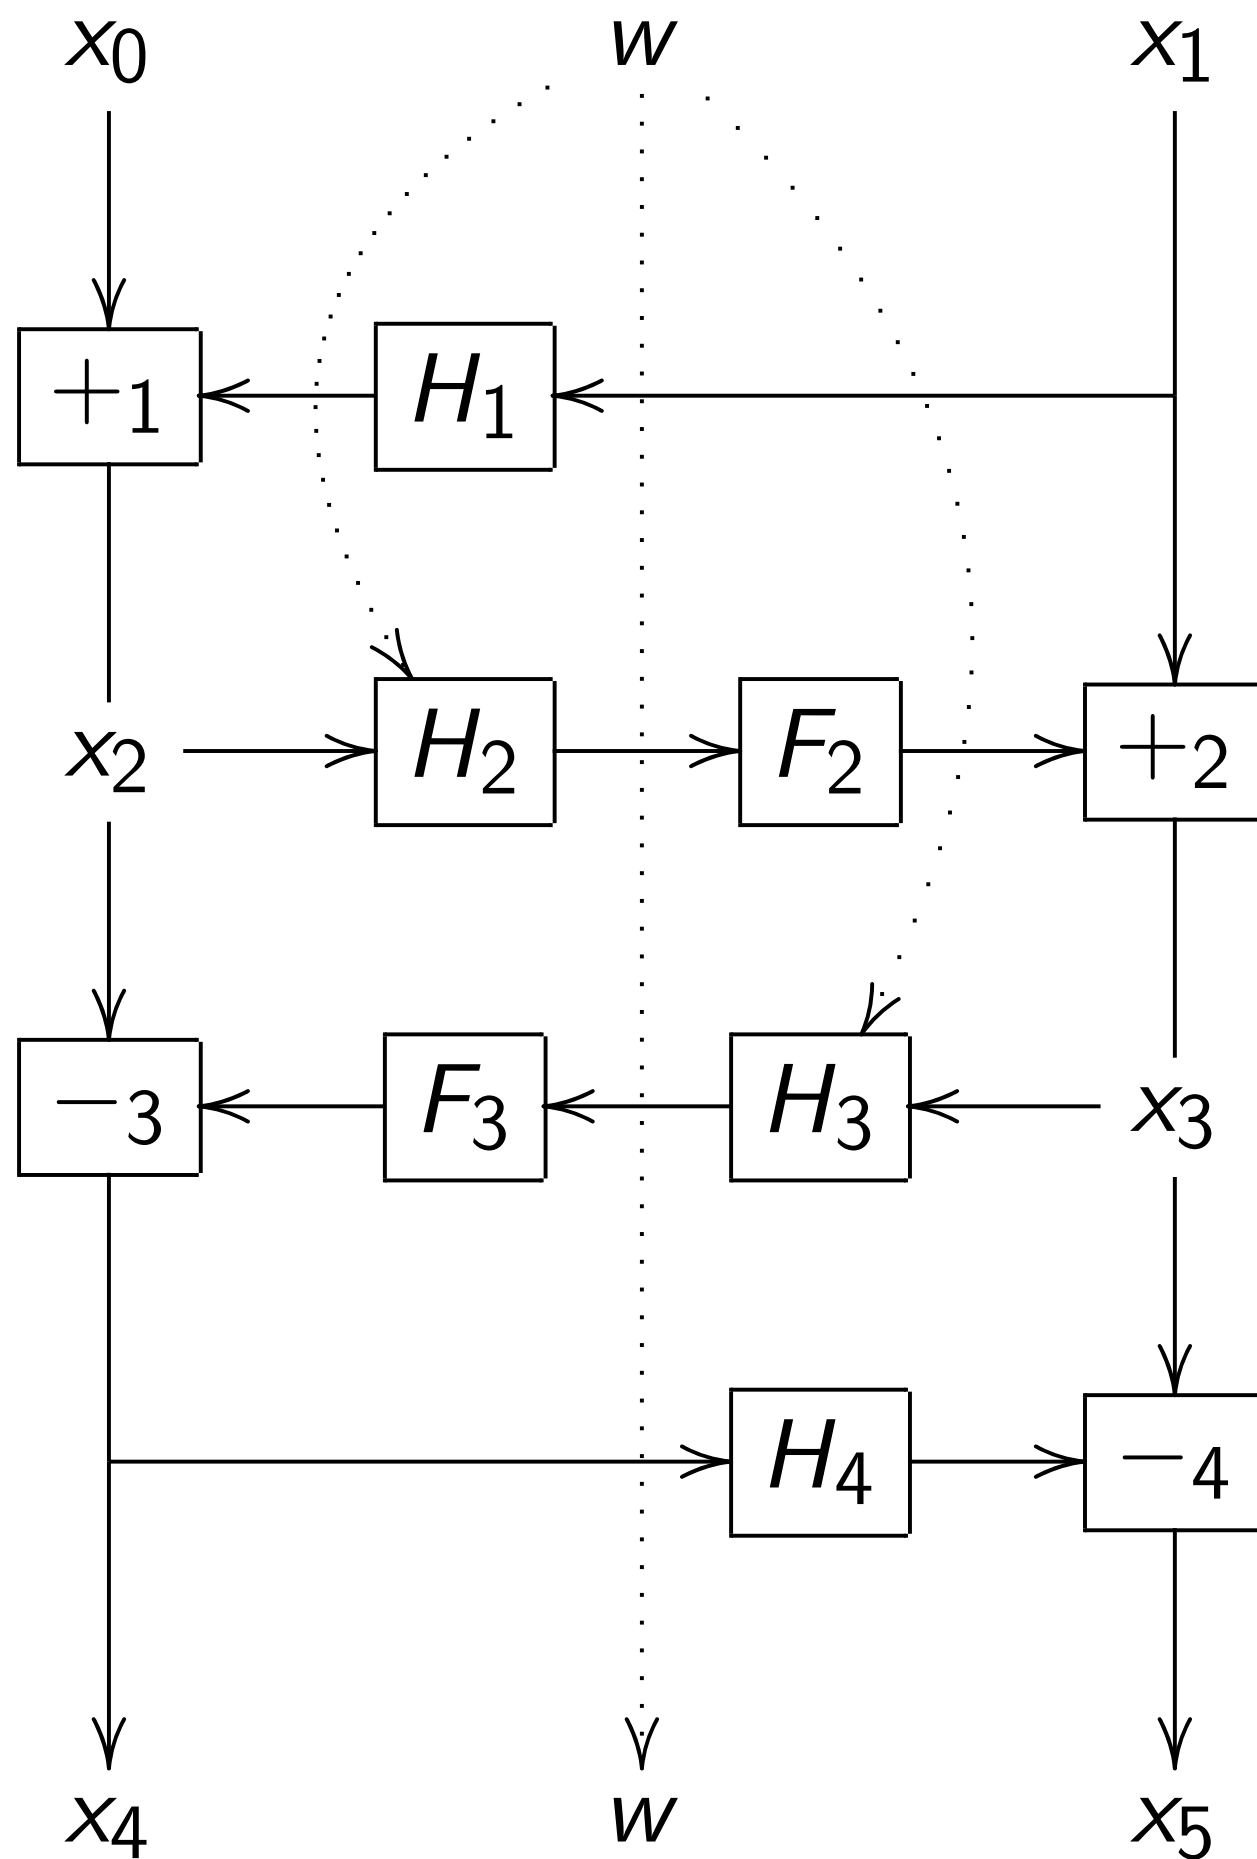


Previous
(Bernste
 H is pur
 F is a st
Ingredient
 H at top
bottom
 H_2, H_3 a
 H_1, H_4 a
XCB/HO
than 200
Allow or
unify H_1
unify H_3

block cipher
strong SPRP)

- NR
- CMC
- EME
- XCB
- HCTR
- PEP
- HCH
- TET
- HEH
- iHCH
- HOH
- EMME

block cipher
strong SPRP)



Previous slide: H
(Bernstein–Nandi–
 H is purely combin
 F is a stream ciph

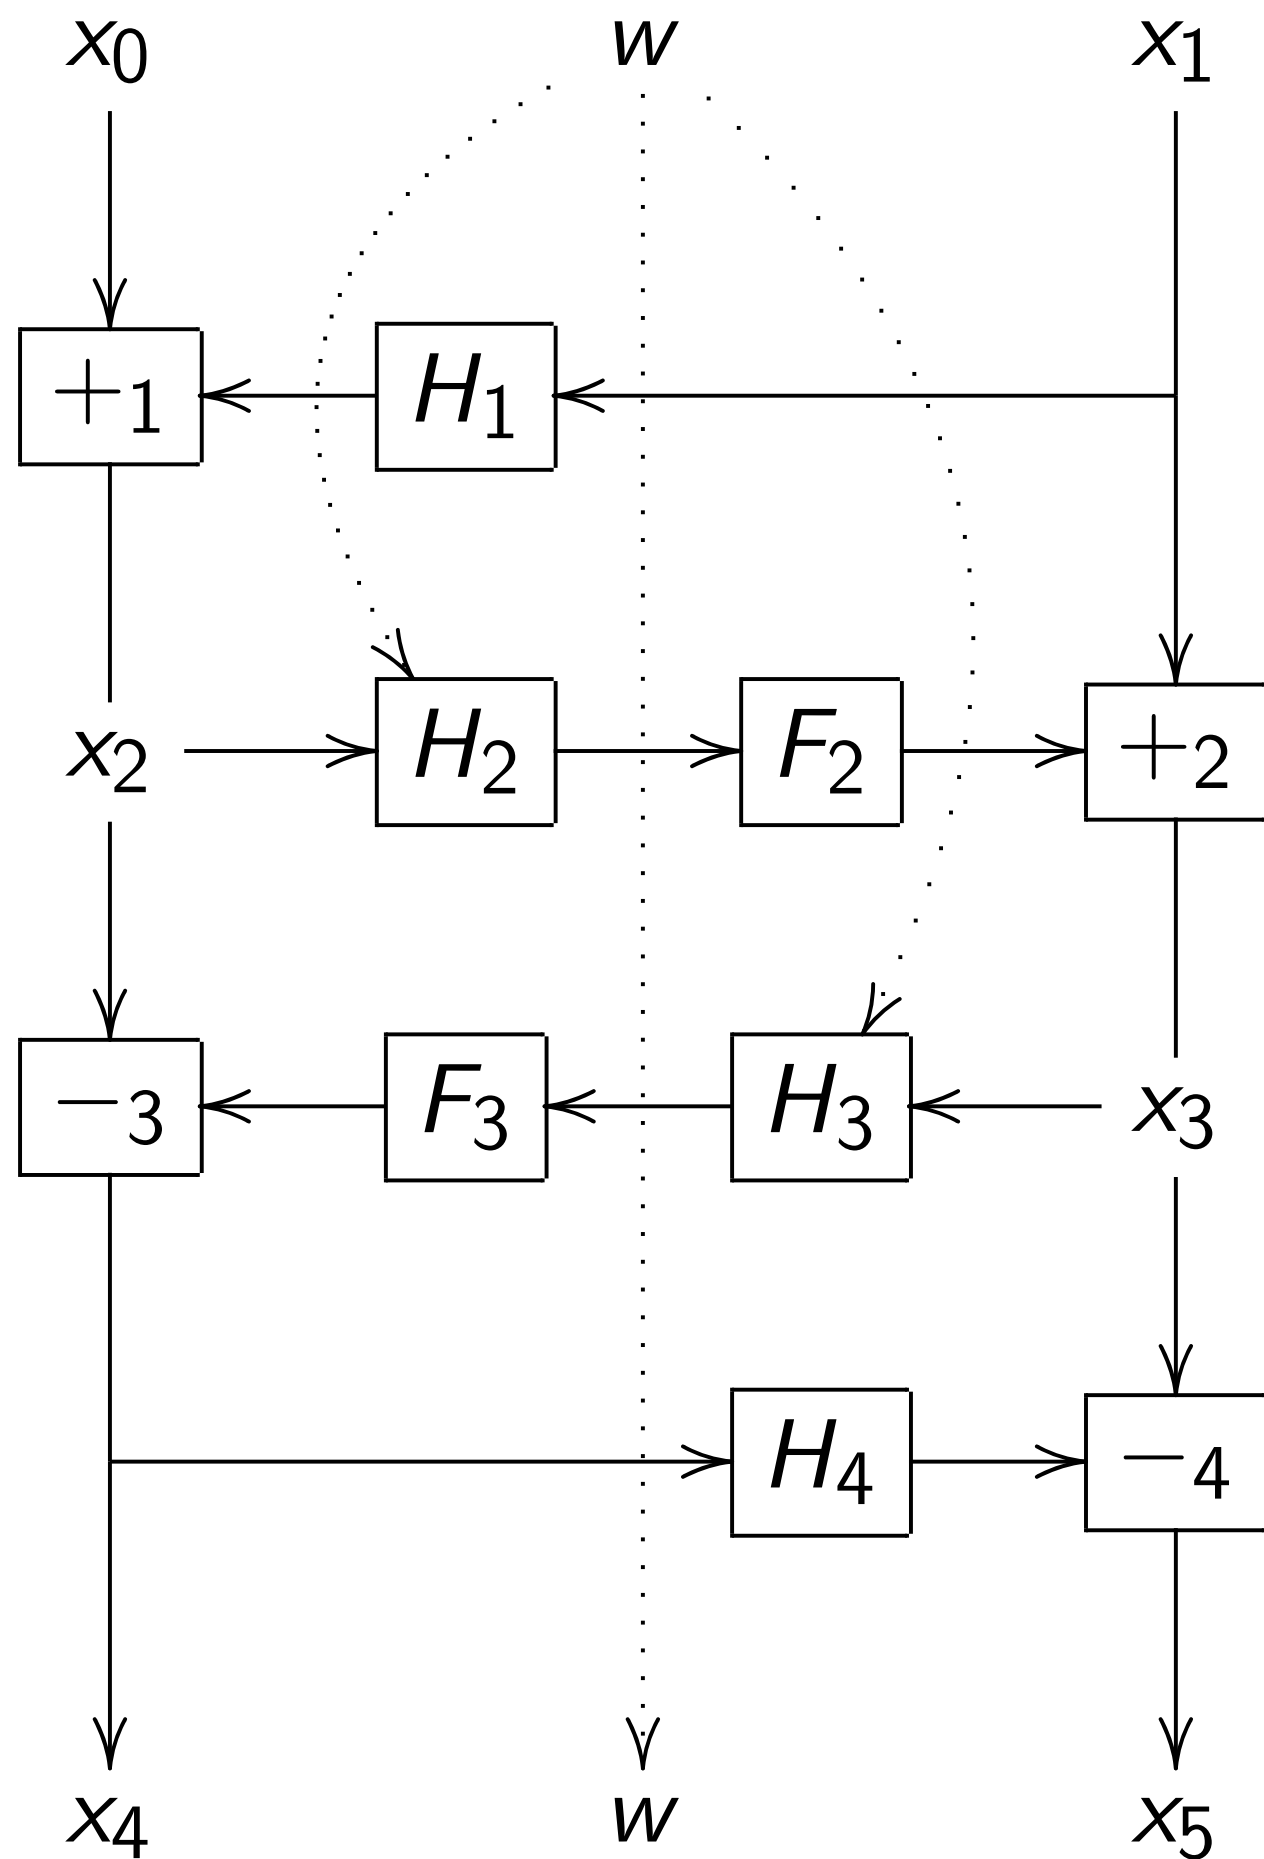
Ingredients: 4-rou
 H at top (1996 Lu
bottom (1997 Nac
 H_2, H_3 allow one-b
 H_1, H_4 are stretch
XCB/HCTR-style
than 2002 Liskov–

Allow one $H_1, H_2,$
unify H_1, H_2 hypo
unify H_3, H_4 hypo

mer
(RP)

NR
MC
ME
CB
CTR
EP
CH
ET
EH
ICH
OH
MME


pher
(RP)



Previous slide: HHFHFH
(Bernstein–Nandi–Sarkar).
 H is purely combinatorial;
 F is a stream cipher.

Ingredients: 4-round Feistel;
 H at top (1996 Lucks),
bottom (1997 Naor–Reingold);
 H_2, H_3 allow one-block nonc
 H_1, H_4 are stretched by 0-pa
XCB/HCTR-style tweak, fas
than 2002 Liskov–Rivest–Wa

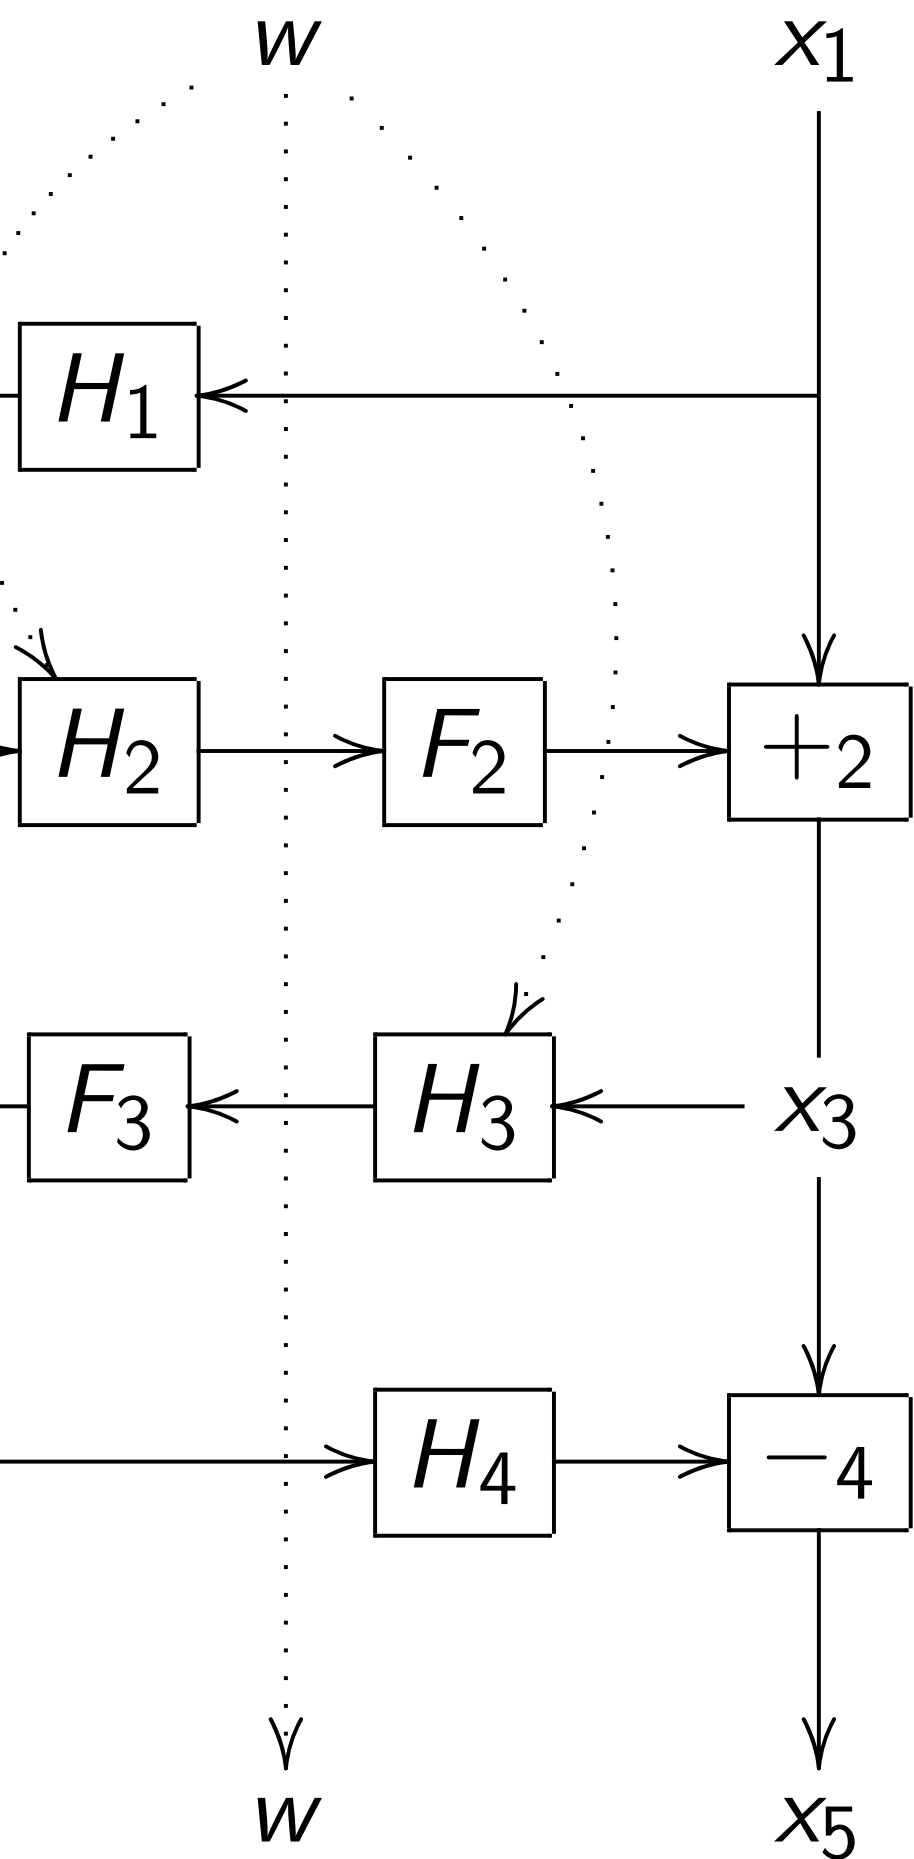
Allow one H_1, H_2, H_3, H_4 ke
unify H_1, H_2 hypotheses;
unify H_3, H_4 hypotheses.



Previous slide: HHFHFH
 (Bernstein–Nandi–Sarkar).
 H is purely combinatorial;
 F is a stream cipher.

Ingredients: 4-round Feistel;
 H at top (1996 Lucks),
 bottom (1997 Naor–Reingold);
 H_2, H_3 allow one-block nonces;
 H_1, H_4 are stretched by 0-pad;
 XCB/HCTR-style tweak, faster
 than 2002 Liskov–Rivest–Wagner.

Allow one H_1, H_2, H_3, H_4 key;
 unify H_1, H_2 hypotheses;
 unify H_3, H_4 hypotheses.



Previous slide: HHFHFH
(Bernstein–Nandi–Sarkar).

H is purely combinatorial;
 F is a stream cipher.

Ingredients: 4-round Feistel;

H at top (1996 Lucks),

bottom (1997 Naor–Reingold);

H_2, H_3 allow one-block nonces;

H_1, H_4 are stretched by 0-pad;

XCB/HCTR-style tweak, faster

than 2002 Liskov–Rivest–Wagner.

Allow one H_1, H_2, H_3, H_4 key;

unify H_1, H_2 hypotheses;

unify H_3, H_4 hypotheses.

One pos
permuta

Full-widt
beats sq

and CTR

Also cho

We're st

Use sing

“chopTC


w as tru

HHFHFH

twice, w

Some thi

more loc



Previous slide: HHFHFH
(Bernstein–Nandi–Sarkar).

H is purely combinatorial;
 F is a stream cipher.

Ingredients: 4-round Feistel;

H at top (1996 Lucks),

bottom (1997 Naor–Reingold);

H_2, H_3 allow one-block nonces;

H_1, H_4 are stretched by 0-pad;

XCB/HCTR-style tweak, faster

than 2002 Liskov–Rivest–Wagner.

Allow one H_1, H_2, H_3, H_4 key;

unify H_1, H_2 hypotheses;

unify H_3, H_4 hypotheses.

One possibility for
permutation in EM

Full-width permut

beats squeezing fo

and CTR is highly

Also choose highly

We're still optimiz

Use single-block tw

“chopTC”: chain

w as truncation of

HHFHFH reads ea

twice, writes each

Something I'm wo

more locality insid

1
Previous slide: HHFHFH
(Bernstein–Nandi–Sarkar).

H is purely combinatorial;
 F is a stream cipher.

2
Ingredients: 4-round Feistel;
 H at top (1996 Lucks),
bottom (1997 Naor–Reingold);
 H_2, H_3 allow one-block nonces;
 H_1, H_4 are stretched by 0-pad;
XCB/HCTR-style tweak, faster
than 2002 Liskov–Rivest–Wagner.

3
4
5
Allow one H_1, H_2, H_3, H_4 key;
unify H_1, H_2 hypotheses;
unify H_3, H_4 hypotheses.

One possibility for F :
permutation in EM in CTR.

Full-width permutation output
beats squeezing for long output
and CTR is highly parallel.

Also choose highly parallel H .
We're still optimizing choice

Use single-block tweak w .

“chopTC”: chain by choosing
 w as truncation of $P \oplus C$.

HHFHFH reads each bit in a
twice, writes each bit once.

Something I'm working on n
more locality inside permuta

Previous slide: HHFHFH
(Bernstein–Nandi–Sarkar).

H is purely combinatorial;
 F is a stream cipher.

Ingredients: 4-round Feistel;
 H at top (1996 Lucks),
bottom (1997 Naor–Reingold);
 H_2, H_3 allow one-block nonces;
 H_1, H_4 are stretched by 0-pad;
XCB/HCTR-style tweak, faster
than 2002 Liskov–Rivest–Wagner.

Allow one H_1, H_2, H_3, H_4 key;
unify H_1, H_2 hypotheses;
unify H_3, H_4 hypotheses.

One possibility for F :
permutation in EM in CTR.

Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.

Also choose highly parallel H .
We're still optimizing choices.

Use single-block tweak w .

“chopTC”: chain by choosing
 w as truncation of $P \oplus C$.

HHFHFH reads each bit in array
twice, writes each bit once.

Something I'm working on now:
more locality inside permutation.

slide: HHFHFH
ein–Nandi–Sarkar).
ely combinatorial;
ream cipher.

nts: 4-round Feistel;
o (1996 Lucks),
(1997 Naor–Reingold);
allow one-block nonces;
are stretched by 0-pad;
CTR-style tweak, faster
02 Liskov–Rivest–Wagner.

ne H_1, H_2, H_3, H_4 key;
, H_2 hypotheses;
, H_4 hypotheses.

One possibility for F :
permutation in EM in CTR.
Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.

Also choose highly parallel H .
We're still optimizing choices.

Use single-block tweak w .
“chopTC”: chain by choosing
 w as truncation of $P \oplus C$.

HHFHFH reads each bit in array
twice, writes each bit once.
Something I'm working on now:
more locality inside permutation.

Security
compare
basically
assuming
and typi
Is this 2

HHFHH
-Sarkar).
natorial;
er.
nd Feistel;
ucks),
or-Reingold);
block nonces;
ed by 0-pad;
tweak, faster
-Rivest-Wagner.
 H_3, H_4 key;
theses;
theses.

One possibility for F :
permutation in EM in CTR.
Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.
Also choose highly parallel H .
We're still optimizing choices.
Use single-block tweak w .
"chopTC": chain by choosing
 w as truncation of $P \oplus C$.
HHFHH reads each bit in array
twice, writes each bit once.
Something I'm working on now:
more locality inside permutation.

Security loss of mo
compared to secur
basically $q^2/2^{128}$,
assuming 128-bit
and typical choice
Is this 2^{128} "secur

One possibility for F :
permutation in EM in CTR.

Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.

Also choose highly parallel H .
We're still optimizing choices.

Use single-block tweak w .
"chopTC": chain by choosing
 w as truncation of $P \oplus C$.

HHFHH reads each bit in array
twice, writes each bit once.

Something I'm working on now:
more locality inside permutation.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .

Is this 2^{128} "security"?

One possibility for F :
permutation in EM in CTR.

Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.

Also choose highly parallel H .
We're still optimizing choices.

Use single-block tweak w .

“chopTC”: chain by choosing
 w as truncation of $P \oplus C$.

HHFHH reads each bit in array
twice, writes each bit once.

Something I'm working on now:
more locality inside permutation.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .

Is this 2^{128} “security”?

One possibility for F :
permutation in EM in CTR.
Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.
Also choose highly parallel H .
We're still optimizing choices.
Use single-block tweak w .
"chopTC": chain by choosing
 w as truncation of $P \oplus C$.
HHFHH reads each bit in array
twice, writes each bit once.
Something I'm working on now:
more locality inside permutation.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .
Is this 2^{128} "security"?
Fragile fix: "beyond-birthday-
bound security." Complicates
implementation, security analysis.

One possibility for F :
permutation in EM in CTR.
Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.
Also choose highly parallel H .
We're still optimizing choices.
Use single-block tweak w .
"chopTC": chain by choosing
 w as truncation of $P \oplus C$.
HHFHH reads each bit in array
twice, writes each bit once.
Something I'm working on now:
more locality inside permutation.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .
Is this 2^{128} "security"?
Fragile fix: "beyond-birthday-
bound security." Complicates
implementation, security analysis.
Simpler fix: "bigger-birthday-
bound security." Use 256-bit
blocks, security $q^2/2^{256}$.

One possibility for F :
permutation in EM in CTR.
Full-width permutation output
beats squeezing for long output;
and CTR is highly parallel.
Also choose highly parallel H .
We're still optimizing choices.
Use single-block tweak w .
"chopTC": chain by choosing
 w as truncation of $P \oplus C$.
HHFHH reads each bit in array
twice, writes each bit once.
Something I'm working on now:
more locality inside permutation.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .
Is this 2^{128} "security"?
Fragile fix: "beyond-birthday-
bound security." Complicates
implementation, security analysis.
Simpler fix: "bigger-birthday-
bound security." Use 256-bit
blocks, security $q^2/2^{256}$.
Is 256-bit n safe in ChaCha?

possibility for F :
operation in EM in CTR.
with permutation output
streaming for long output;
CTR is highly parallel.
use highly parallel H .
still optimizing choices.
single-block tweak w .
"C": chain by choosing
function of $P \oplus C$.
 H reads each bit in array
writes each bit once.
thing I'm working on now:
locality inside permutation.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .

Is this 2^{128} "security"?

Fragile fix: "beyond-birthday-
bound security." Complicates
implementation, security analysis.

Simpler fix: "bigger-birthday-
bound security." Use 256-bit
blocks, security $q^2/2^{256}$.

Is 256-bit n safe in ChaCha?

Heavyweight

Interesting

≥ 256 bits

≥ 256 -bit

≥ 256 -bit

F :
 M in CTR.
ation output
or long output;
parallel.
y parallel H .
ing choices.
weak w .
by choosing
f $P \oplus C$.
ch bit in array
bit once.
rking on now:
e permutation.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .
Is this 2^{128} “security”?
Fragile fix: “beyond-birthday-
bound security.” Complicates
implementation, security analysis.
Simpler fix: “bigger-birthday-
bound security.” Use 256-bit
blocks, security $q^2/2^{256}$.
Is 256-bit n safe in ChaCha?

Heavyweight cipher
Interesting **cipher**-
 ≥ 256 bits for all p
 ≥ 256 -bit keys, ≥ 2
 ≥ 256 -bit subkeys,

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .

Is this 2^{128} “security”?

Fragile fix: “beyond-birthday-
bound security.” Complicates
implementation, security analysis.

Simpler fix: “bigger-birthday-
bound security.” Use 256-bit
blocks, security $q^2/2^{256}$.

Is 256-bit n safe in ChaCha?

Heavyweight ciphers

Interesting **cipher**-design space
 ≥ 256 bits for all pipes.
 ≥ 256 -bit keys, ≥ 256 -bit outputs
 ≥ 256 -bit subkeys, etc.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .

Is this 2^{128} “security”?

Fragile fix: “beyond-birthday-
bound security.” Complicates
implementation, security analysis.

Simpler fix: “bigger-birthday-
bound security.” Use 256-bit
blocks, security $q^2/2^{256}$.

Is 256-bit n safe in ChaCha?

Heavyweight ciphers

Interesting **cipher**-design space:
 ≥ 256 bits for all pipes.
 ≥ 256 -bit keys, ≥ 256 -bit outputs,
 ≥ 256 -bit subkeys, etc.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .

Is this 2^{128} “security”?

Fragile fix: “beyond-birthday-
bound security.” Complicates
implementation, security analysis.

Simpler fix: “bigger-birthday-
bound security.” Use 256-bit
blocks, security $q^2/2^{256}$.

Is 256-bit n safe in ChaCha?

Heavyweight ciphers

Interesting **cipher**-design space:
 ≥ 256 bits for all pipes.
 ≥ 256 -bit keys, ≥ 256 -bit outputs,
 ≥ 256 -bit subkeys, etc.

Occasional designs: Rijndael,
OMD (SHA-2), Keccak, BLAKE2,
NORX, Simpira, This needs
far more attention, optimization.

Hash designs are usually overkill.

Security loss of mode
compared to security of F :
basically $q^2/2^{128}$,
assuming 128-bit blocks
and typical choice of H .

Is this 2^{128} “security”?

Fragile fix: “beyond-birthday-
bound security.” Complicates
implementation, security analysis.

Simpler fix: “bigger-birthday-
bound security.” Use 256-bit
blocks, security $q^2/2^{256}$.

Is 256-bit n safe in ChaCha?

Heavyweight ciphers

Interesting **cipher**-design space:
 ≥ 256 bits for all pipes.
 ≥ 256 -bit keys, ≥ 256 -bit outputs,
 ≥ 256 -bit subkeys, etc.

Occasional designs: Rijndael,
OMD (SHA-2), Keccak, BLAKE2,
NORX, Simpira, This needs
far more attention, optimization.

Hash designs are usually overkill.

Is 256 fundamentally much slower,
or much less energy-efficient,
than 128? My guess: No!

loss of mode

reduced to security of F :

$$q^2 / 2^{128},$$

using 128-bit blocks

arbitrary choice of H .

2^{128} “security”?

fix: “beyond-birthday-

security.” Complicates

implementation, security analysis.

fix: “bigger-birthday-

security.” Use 256-bit

$$\text{security } q^2 / 2^{256}.$$

is n safe in ChaCha?

Heavyweight ciphers

Interesting **cipher**-design space:

≥ 256 bits for all pipes.

≥ 256 -bit keys, ≥ 256 -bit outputs,

≥ 256 -bit subkeys, etc.

Occasional designs: Rijndael,

OMD (SHA-2), Keccak, BLAKE2,

NORX, Simpira, This needs

far more attention, optimization.

Hash designs are usually overkill.

Is 256 fundamentally much slower,

or much less energy-efficient,

than 128? My guess: No!

Another

PRF inside

EdDSA

random

truncate

H is SHA

2015 Be

truncate

high-sec

Even with

reusing p

surely ca

in both s

ode

ity of F :

blocks

of H .

ity"?

nd-birthday-

Complicates

security analysis.

er-birthday-

Use 256-bit

$2/2^{256}$.

n ChaCha?

Heavyweight ciphers

Interesting **cipher**-design space:

≥ 256 bits for all pipes.

≥ 256 -bit keys, ≥ 256 -bit outputs,

≥ 256 -bit subkeys, etc.

Occasional designs: Rijndael,

OMD (SHA-2), Keccak, BLAKE2,

NORX, Simpira, This needs

far more attention, optimization.

Hash designs are usually overkill.

Is 256 fundamentally much slower,

or much less energy-efficient,

than 128? My guess: No!

Another optimization

PRF inside EdDSA

EdDSA generates

random number m

truncated hash: H

H is SHA-512; s is

2015 Bellare–Bern

truncated prefixed

high-security mult

Even with the con

reusing preimage-r

surely can build be

in both software a

Heavyweight ciphers

Interesting **cipher**-design space:

≥ 256 bits for all pipes.

≥ 256 -bit keys, ≥ 256 -bit outputs,

≥ 256 -bit subkeys, etc.

Occasional designs: Rijndael, OMD (SHA-2), Keccak, BLAKE2, NORX, Simpira, This needs far more attention, optimization.

Hash designs are usually overkill.

Is 256 fundamentally much slower, or much less energy-efficient, than 128? My guess: No!

Another optimization target: PRF inside EdDSA signature

EdDSA generates per-signature random number mod 256-bit truncated hash: $H(s, m)$ mod H is SHA-512; s is subkey.

2015 Bellare–Bernstein–Tessier truncated prefixed MD hash high-security multi-user MAC

Even with the constraint of reusing preimage-resistant hash surely can build better design in both software and hardware

Heavyweight ciphers

Interesting **cipher**-design space:

≥ 256 bits for all pipes.

≥ 256 -bit keys, ≥ 256 -bit outputs,

≥ 256 -bit subkeys, etc.

Occasional designs: Rijndael, OMD (SHA-2), Keccak, BLAKE2, NORX, Simpira, This needs far more attention, optimization.

Hash designs are usually overkill.

Is 256 fundamentally much slower, or much less energy-efficient, than 128? My guess: No!

Another optimization target:

PRF inside EdDSA signatures.

EdDSA generates per-signature random number mod 256-bit ℓ as truncated hash: $H(s, m) \bmod \ell$. H is SHA-512; s is subkey.

2015 Bellare–Bernstein–Tessaro: truncated prefixed MD hash is a high-security multi-user MAC.

Even with the constraint of reusing preimage-resistant hash, surely can build better design in both software and hardware.