

Non-uniform

cracks in the concrete:

the power of free precomputation

D. J. Bernstein

University of Illinois at Chicago &

Technische Universiteit Eindhoven

Tanja Lange

Technische Universiteit Eindhoven

Full 53-page paper,

including progress towards

formalizing collision resistance:

eprint.iacr.org/2012/318

Concrete security: an example

What is the best NIST P-256
discrete-log attack algorithm?

ECDL input: P-256 points P, Q ,
where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best”:
minimize “time”.

Non-uniform

cracks in the concrete:

the power of free precomputation

D. J. Bernstein

University of Illinois at Chicago &

Technische Universiteit Eindhoven

Tanja Lange

Technische Universiteit Eindhoven

Full 53-page paper,

including progress towards

formalizing collision resistance:

eprint.iacr.org/2012/318

Concrete security: an example

What is the best NIST P-256
discrete-log attack algorithm?

ECDL input: P-256 points P, Q ,
where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best”:
minimize “time”.

More generally, allow attacks with
<100% success probability;

analyze tradeoffs between

“time” and success probability.

This talk focuses on high prob.

form

in the concrete:

number of free precomputation

Bernstein

University of Illinois at Chicago &

Radboud University Eindhoven

range

Radboud University Eindhoven

page paper,

making progress towards

improving collision resistance:

ia.cr.org/2012/318

Concrete security: an example

What is the best NIST P-256
discrete-log attack algorithm?

ECDL input: P-256 points P, Q ,
where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best”:
minimize “time”.

More generally, allow attacks with
<100% success probability;

analyze tradeoffs between

“time” and success probability.

This talk focuses on high prob.

P-256 di

total TL

Should

Concrete security: an example

What is the best NIST P-256 discrete-log attack algorithm?

ECDL input: P-256 points P, Q , where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best” :
minimize “time” .

More generally, allow attacks with
<100% success probability;
analyze tradeoffs between
“time” and success probability.
This talk focuses on high prob.

P-256 discrete-log
total TLS-ECDHE
Should TLS users

Concrete security: an example

What is the best NIST P-256 discrete-log attack algorithm?

ECDL input: P-256 points P, Q , where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best”:
minimize “time”.

More generally, allow attacks with $<100\%$ success probability;
analyze tradeoffs between
“time” and success probability.
This talk focuses on high prob.

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 bro
Should TLS users worry?

cation

ago &
hoven

hoven

ce:

18

Concrete security: an example

What is the best NIST P-256 discrete-log attack algorithm?

ECDL input: P-256 points P, Q , where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best”:
minimize “time”.

More generally, allow attacks with $<100\%$ success probability;
analyze tradeoffs between
“time” and success probability.
This talk focuses on high prob.

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 break!
Should TLS users worry?

Concrete security: an example

What is the best NIST P-256 discrete-log attack algorithm?

ECDL input: P-256 points P, Q , where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best”:
minimize “time”.

More generally, allow attacks with $<100\%$ success probability;
analyze tradeoffs between
“time” and success probability.
This talk focuses on high prob.

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 break!
Should TLS users worry?

No. Many researchers have
tried and failed to find good
P-256 discrete-log attacks.

Concrete security: an example

What is the best NIST P-256 discrete-log attack algorithm?

ECDL input: P-256 points P, Q , where P is a standard generator.

ECDL output: $\log_P Q$.

Standard definition of “best”:
minimize “time”.

More generally, allow attacks with $<100\%$ success probability;

analyze tradeoffs between
“time” and success probability.

This talk focuses on high prob.

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 break!
Should TLS users worry?

No. Many researchers have
tried and failed to find good
P-256 discrete-log attacks.

Standard conjecture:

For each $p \in [0, 1]$,
each P-256 ECDL algorithm
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Similar conjectures for AES-128,
RSA-3072, etc.: see, e.g.,
2005 Bellare–Rogaway.

Security: an example

the best NIST P-256
log attack algorithm?

Input: P-256 points P, Q ,
 g is a standard generator.

Output: $\log_P Q$.

and definition of “best”:
the “time”.

generally, allow attacks with
success probability;

tradeoffs between
and success probability.
focuses on high prob.

P-256 discrete-log attack \Rightarrow

total TLS-ECDHE-P-256 break!
Should TLS users worry?

No. Many researchers have
tried and failed to find good
P-256 discrete-log attacks.

Standard conjecture:

For each $p \in [0, 1]$,
each P-256 ECDL algorithm
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Similar conjectures for AES-128,
RSA-3072, etc.: see, e.g.,
2005 Bellare–Rogaway.

Concrete

Another
Each TL
with suc
takes “ti

an example

NIST P-256
algorithm?

56 points $P, Q,$
standard generator.

$P, Q.$

of “best”:

ow attacks with
probability;

between

s probability.

on high prob.

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 break!
Should TLS users worry?

No. Many researchers have
tried and failed to find good
P-256 discrete-log attacks.

Standard conjecture:

For each $p \in [0, 1]$,
each P-256 ECDL algorithm
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Similar conjectures for AES-128,
RSA-3072, etc.: see, e.g.,
2005 Bellare–Rogaway.

Concrete reduction

Another conjecture:
Each TLS-ECDHE
with success probab
takes “time” $\geq 2^{128}$

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 break!
Should TLS users worry?

No. Many researchers have
tried and failed to find good
P-256 discrete-log attacks.

Standard conjecture:

For each $p \in [0, 1]$,
each P-256 ECDL algorithm
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Similar conjectures for AES-128,
RSA-3072, etc.: see, e.g.,
2005 Bellare–Rogaway.

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 at
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 break!
Should TLS users worry?

No. Many researchers have
tried and failed to find good
P-256 discrete-log attacks.

Standard conjecture:

For each $p \in [0, 1]$,
each P-256 ECDL algorithm
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Similar conjectures for AES-128,
RSA-3072, etc.: see, e.g.,
2005 Bellare–Rogaway.

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 attack
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

P-256 discrete-log attack \Rightarrow
total TLS-ECDHE-P-256 break!
Should TLS users worry?

No. Many researchers have
tried and failed to find good
P-256 discrete-log attacks.

Standard conjecture:

For each $p \in [0, 1]$,
each P-256 ECDL algorithm
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Similar conjectures for AES-128,
RSA-3072, etc.: see, e.g.,
2005 Bellare–Rogaway.

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 attack
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Why should users have any
confidence in this conjecture?

How many researchers
have really tried to break
ECDHE-P-256? ECDSA-P-256?
ECIES-P-256? ECMQV-P-256?
Other P-256-based protocols?
Far less attention than for ECDL.

discrete-log attack \Rightarrow

TLS-ECDHE-P-256 break!

TLS users worry?

Many researchers have
tried and failed to find good
discrete-log attacks.

Standard conjecture:

For any $p \in [0, 1]$,

any P-256 ECDL algorithm

with success probability $\geq p$

takes “time” $\geq 2^{128} p^{1/2}$.

Similar conjectures for AES-128,

AES-256, etc.: see, e.g.,

Shoup–Rogaway.

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 attack
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Why should users have any
confidence in this conjecture?

How many researchers

have really tried to break

ECDHE-P-256? ECDSA-P-256?

ECIES-P-256? ECMQV-P-256?

Other P-256-based protocols?

Far less attention than for ECDL.

Provable

Prove: if

there is a TLS-E

with success

probability $\geq p$

then there is

a P-256 ECDL

algorithm with

success probab-

ility $\geq p$

and time \leq

$2^{128} p^{1/2}$.

Conjecture:

if there is a

discrete-log

attack \Rightarrow
P-256 break!
worry?

others have
find good
attacks.

re:

algorithm
probability $\geq p$
 $2^{128} p^{1/2}$.

s for AES-128,
see, e.g.,
away.

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 attack
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Why should users have any
confidence in this conjecture?

How many researchers
have really tried to break
ECDHE-P-256? ECDSA-P-256?
ECIES-P-256? ECMQV-P-256?
Other P-256-based protocols?
Far less attention than for ECDL.

Provable security t

Prove: if there is
a TLS-ECDHE-P-
then there is
a P-256 discrete-l
with similar “time”
and success proba

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 attack
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Why should users have any
confidence in this conjecture?

How many researchers
have really tried to break
ECDHE-P-256? ECDSA-P-256?
ECIES-P-256? ECMQV-P-256?
Other P-256-based protocols?
Far less attention than for ECDL.

Provable security to the rescue

Prove: if there is
a TLS-ECDHE-P-256 attack
then there is
a P-256 discrete-log attack
with similar “time”
and success probability.

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 attack with success probability $\geq p$ takes “time” $\geq 2^{128} p^{1/2}$.

Why should users have any confidence in this conjecture?

How many researchers have really tried to break ECDHE-P-256? ECDSA-P-256? ECIES-P-256? ECMQV-P-256? Other P-256-based protocols? Far less attention than for ECDL.

Provable security to the rescue!

Prove: if there is a TLS-ECDHE-P-256 attack then there is a P-256 discrete-log attack with similar “time” and success probability.

Concrete reductions

Another conjecture:

Each TLS-ECDHE-P-256 attack
with success probability $\geq p$
takes “time” $\geq 2^{128} p^{1/2}$.

Why should users have any
confidence in this conjecture?

How many researchers
have really tried to break
ECDHE-P-256? ECDSA-P-256?
ECIES-P-256? ECMQV-P-256?
Other P-256-based protocols?
Far less attention than for ECDL.

Provable security to the rescue!

Prove: if there is
a TLS-ECDHE-P-256 attack
then there is
a P-256 discrete-log attack
with similar “time”
and success probability.

Oops: This turns out to be hard.
But changing DL to DDH
+ adding more assumptions
allows a proof: Crypto 2012
Jager–Kohlar–Schäge–Schwenk
“On the security of TLS-DHE
in the standard model”.

reductions

conjecture:

TLS-ECDHE-P-256 attack

success probability $\geq p$

“time” $\geq 2^{128} p^{1/2}$.

Should users have any

concern in this conjecture?

Many researchers

have already tried to break

ECDHE-P-256? ECDSA-P-256?

ECDHE-P-256? ECMQV-P-256?

Other P-256-based protocols?

More attention than for ECDL.

Provable security to the rescue!

Prove: if there is

a TLS-ECDHE-P-256 attack

then there is

a P-256 discrete-log attack

with similar “time”

and success probability.

Oops: This turns out to be hard.

But changing DL to DDH

+ adding more assumptions

allows a proof: Crypto 2012

Jager–Kohlar–Schäge–Schwenk

“On the security of TLS-DHE

in the standard model”.

Similar proof

“provable

Protocol

that hard

(e.g., P-

security

After ex

maybe g

of P , an

ns

e:

E-P-256 attack

probability $\geq p$

$2^{28} p^{1/2}$.

have any

conjecture?

hers

to break

CDSA-P-256?

MQV-P-256?

d protocols?

than for ECDL.

Provable security to the rescue!

Prove: if there is

a TLS-ECDHE-P-256 attack

then there is

a P-256 discrete-log attack

with similar “time”

and success probability.

Oops: This turns out to be hard.

But changing DL to DDH

+ adding more assumptions

allows a proof: Crypto 2012

Jager–Kohlar–Schäge–Schwenk

“On the security of TLS-DHE

in the standard model”.

Similar pattern th

“provable security”

Protocol designers

that hardness of a

(e.g., P-256 DDH)

security of various

After extensive cry

maybe gain confid

of P , and hence in

Provable security to the rescue!

Prove: if there is

a TLS-ECDHE-P-256 attack

then there is

a P-256 discrete-log attack

with similar “time”

and success probability.

Oops: This turns out to be hard.

But changing DL to DDH

+ adding more assumptions

allows a proof: Crypto 2012

Jager–Kohlar–Schäge–Schwenk

“On the security of TLS-DHE

in the standard model”.

Similar pattern throughout the

“provable security” literature

Protocol designers (try to) prove

that hardness of a problem P

(e.g., P-256 DDH) implies

security of various protocols

After extensive cryptanalysis

maybe gain confidence in hardness

of P , and hence in security of

Provable security to the rescue!

Prove: if there is
a TLS-ECDHE-P-256 attack
then there is
a P-256 discrete-log attack
with similar “time”
and success probability.

Oops: This turns out to be hard.

But changing DL to DDH

+ adding more assumptions

allows a proof: Crypto 2012

Jager–Kohlar–Schäge–Schwenk

“On the security of TLS-DHE

in the standard model”.

Similar pattern throughout the
“provable security” literature.

Protocol designers (try to) prove
that hardness of a problem P
(e.g., P-256 DDH) implies
security of various protocols Q .

After extensive cryptanalysis of P ,
maybe gain confidence in hardness
of P , and hence in security of Q .

Provable security to the rescue!

Prove: if there is
a TLS-ECDHE-P-256 attack
then there is
a P-256 discrete-log attack
with similar “time”
and success probability.

Oops: This turns out to be hard.

But changing DL to DDH

+ adding more assumptions

allows a proof: Crypto 2012

Jager–Kohlar–Schäge–Schwenk

“On the security of TLS-DHE
in the standard model”.

Similar pattern throughout the
“provable security” literature.

Protocol designers (try to) prove
that hardness of a problem P
(e.g., P-256 DDH) implies
security of various protocols Q .

After extensive cryptanalysis of P ,
maybe gain confidence in hardness
of P , and hence in security of Q .

Why not directly cryptanalyze Q ?

Cryptanalysis is hard work: have
to focus on *a few* problems P .

Proofs scale to *many* protocols Q .

security to the rescue!

if there is

CDHE-P-256 attack

there is

discrete-log attack

similar “time”

success probability.

This turns out to be hard.

Linking DL to DDH

using more assumptions

proof: Crypto 2012

Bohler–Schäge–Schwenk

“security of TLS-DHE

in the standard model”.

Similar pattern throughout the
“provable security” literature.

Protocol designers (try to) prove
that hardness of a problem P
(e.g., P-256 DDH) implies
security of various protocols Q .

After extensive cryptanalysis of P ,
one may gain confidence in hardness
of P , and hence in security of Q .

Why not directly cryptanalyze Q ?

Cryptanalysis is hard work: have
to focus on *a few* problems P .

Proofs scale to *many* protocols Q .

Interlude

How many

following

```
def pick
```

```
  if n0
```

```
    if
```

```
      if
```

```
        if
```

```
          if
```

```
            ret
```

```
          if n1
```

```
            if
```

```
              ret
```

```
            if n2
```

```
              retur
```


to the rescue!

256 attack

og attack

bility.

out to be hard.

to DDH

assumptions

ypto 2012

äge–Schwenk

of TLS-DHE

odel” .

Similar pattern throughout the
“provable security” literature.

Protocol designers (try to) prove
that hardness of a problem P
(e.g., P-256 DDH) implies
security of various protocols Q .

After extensive cryptanalysis of P ,
maybe gain confidence in hardness
of P , and hence in security of Q .

Why not directly cryptanalyze Q ?

Cryptanalysis is hard work: have
to focus on *a few* problems P .

Proofs scale to *many* protocols Q .

Interlude regarding

How much “time”
following algorithm

```
def pidigit(n0, n1, n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0:  
                return 0  
            if n2 == 0:  
                return 0  
        if n1 == 0:  
            if n2 == 0:  
                return 0  
            if n2 == 0:  
                return 0  
    if n2 == 0:  
        return 0
```

Similar pattern throughout the “provable security” literature.

Protocol designers (try to) prove that hardness of a problem P (e.g., P-256 DDH) implies security of various protocols Q .

After extensive cryptanalysis of P , maybe gain confidence in hardness of P , and hence in security of Q .

Why not directly cryptanalyze Q ?

Cryptanalysis is hard work: have to focus on *a few* problems P .

Proofs scale to *many* protocols Q .

Interlude regarding “time”

How much “time” does the following algorithm take?

```
def pidigit(n0,n1,n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0: return  
            return  
        if n2 == 0: return  
        return  
    if n1 == 0:  
        if n2 == 0: return  
        return  
    if n2 == 0: return  
    return
```

Similar pattern throughout the “provable security” literature.

Protocol designers (try to) prove that hardness of a problem P (e.g., P-256 DDH) implies security of various protocols Q .

After extensive cryptanalysis of P , maybe gain confidence in hardness of P , and hence in security of Q .

Why not directly cryptanalyze Q ?

Cryptanalysis is hard work: have to focus on *a few* problems P .

Proofs scale to *many* protocols Q .

Interlude regarding “time”

How much “time” does the following algorithm take?

```
def pidigit(n0,n1,n2):
    if n0 == 0:
        if n1 == 0:
            if n2 == 0: return 3
            return 1
        if n2 == 0: return 4
        return 1
    if n1 == 0:
        if n2 == 0: return 5
        return 9
    if n2 == 0: return 2
    return 6
```

pattern throughout the
"security" literature.

designers (try to) prove
hardness of a problem P
(256 DDH) implies
security of various protocols Q .

extensive cryptanalysis of P ,
regain confidence in hardness
and hence in security of Q .

not directly cryptanalyze Q ?
analysis is hard work: have
confidence on a few problems P .
scale to many protocols Q .

Interlude regarding "time"

How much "time" does the
following algorithm take?

```
def pidigit(n0,n1,n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0: return 3  
            return 1  
        if n2 == 0: return 4  
        return 1  
    if n1 == 0:  
        if n2 == 0: return 5  
        return 9  
    if n2 == 0: return 2  
    return 6
```

Students
learn to
Skipped
This alg

throughout the

' literature.

(try to) prove

problem P

) implies

protocols Q .

cryptanalysis of P ,

ence in hardness

n security of Q .

cryptanalyze Q ?

ard work: have

problems P .

any protocols Q .

Interlude regarding "time"

How much "time" does the following algorithm take?

```
def pidigit(n0,n1,n2):
    if n0 == 0:
        if n1 == 0:
            if n2 == 0: return 3
            return 1
        if n2 == 0: return 4
        return 1
    if n1 == 0:
        if n2 == 0: return 5
        return 9
    if n2 == 0: return 2
    return 6
```

Students in algorithm

learn to count exe

Skipped branches

This algorithm use

Interlude regarding “time”

How much “time” does the following algorithm take?

```
def pidigit(n0,n1,n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0: return 3  
            return 1  
        if n2 == 0: return 4  
        return 1  
    if n1 == 0:  
        if n2 == 0: return 5  
        return 9  
    if n2 == 0: return 2  
    return 6
```

Students in algorithm courses learn to count executed “steps”. Skipped branches take 0 “steps”. This algorithm uses 4 “steps”.

Interlude regarding “time”

How much “time” does the following algorithm take?

```
def pidigit(n0,n1,n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0: return 3  
            return 1  
        if n2 == 0: return 4  
        return 1  
    if n1 == 0:  
        if n2 == 0: return 5  
        return 9  
    if n2 == 0: return 2  
    return 6
```

Students in algorithm courses learn to count executed “steps”. Skipped branches take 0 “steps”.

This algorithm uses 4 “steps”.

Interlude regarding “time”

How much “time” does the following algorithm take?

```
def pidigit(n0,n1,n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0: return 3  
            return 1  
        if n2 == 0: return 4  
        return 1  
    if n1 == 0:  
        if n2 == 0: return 5  
        return 9  
    if n2 == 0: return 2  
    return 6
```

Students in algorithm courses learn to count executed “steps”. Skipped branches take 0 “steps”.

This algorithm uses 4 “steps”.

Generalization: There exists an algorithm that, given $n < 2^k$, prints the n th digit of π using $k + 1$ “steps”.

Interlude regarding “time”

How much “time” does the following algorithm take?

```
def pidigit(n0,n1,n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0: return 3  
            return 1  
        if n2 == 0: return 4  
        return 1  
    if n1 == 0:  
        if n2 == 0: return 5  
        return 9  
    if n2 == 0: return 2  
    return 6
```

Students in algorithm courses learn to count executed “steps”. Skipped branches take 0 “steps”.

This algorithm uses 4 “steps”.

Generalization: There exists an algorithm that, given $n < 2^k$, prints the n th digit of π using $k + 1$ “steps”.

Variant: There exists a 258-“step” P-256 discrete-log attack (with 100% success probability).

Interlude regarding “time”

How much “time” does the following algorithm take?

```
def pidigit(n0,n1,n2):  
    if n0 == 0:  
        if n1 == 0:  
            if n2 == 0: return 3  
            return 1  
        if n2 == 0: return 4  
        return 1  
    if n1 == 0:  
        if n2 == 0: return 5  
        return 9  
    if n2 == 0: return 2  
    return 6
```

Students in algorithm courses learn to count executed “steps”. Skipped branches take 0 “steps”.

This algorithm uses 4 “steps”.

Generalization: There exists an algorithm that, given $n < 2^k$, prints the n th digit of π using $k + 1$ “steps”.

Variant: There exists a 258-“step” P-256 discrete-log attack (with 100% success probability). If “time” means “steps” then the standard conjectures are wrong.

the regarding "time"

ch "time" does the
g algorithm take?

```
digit(n0,n1,n2):  
  ) == 0:  
  n1 == 0:  
  if n2 == 0: return 3  
  return 1  
  n2 == 0: return 4  
  urn 1  
  l == 0:  
  n2 == 0: return 5  
  urn 9  
  2 == 0: return 2  
  rn 6
```

Students in algorithm courses
learn to count executed "steps".
Skipped branches take 0 "steps".

This algorithm uses 4 "steps".

Generalization: There exists an
algorithm that, given $n < 2^k$,
prints the n th digit of π
using $k + 1$ "steps".

Variant: There exists a 258-
"step" P-256 discrete-log attack
(with 100% success probability).
If "time" means "steps" then the
standard conjectures are wrong.

1994 Be
"We say
A is a (t
A runs i
makes a

g “time”

does the
n take?

, n1, n2) :

0: return 3

1

: return 4

1

: return 5

9

return 2

6

Students in algorithm courses
learn to count executed “steps” .
Skipped branches take 0 “steps” .

This algorithm uses 4 “steps” .

Generalization: There exists an
algorithm that, given $n < 2^k$,
prints the n th digit of π
using $k + 1$ “steps” .

Variant: There exists a 258-
“step” P-256 discrete-log attack
(with 100% success probability).
If “time” means “steps” then the
standard conjectures are wrong.

1994 Bellare–Kilian

*“We say that
A is a (t, q) -adversary
A runs in at most t
makes at most q queries*

Students in algorithm courses
learn to count executed “steps” .
Skipped branches take 0 “steps” .

This algorithm uses 4 “steps” .

Generalization: There exists an
algorithm that, given $n < 2^k$,
prints the n th digit of π
using $k + 1$ “steps” .

Variant: There exists a 258-
“step” P-256 discrete-log attack
(with 100% success probability).
If “time” means “steps” then the
standard conjectures are wrong.

1994 Bellare–Kilian–Rogawa

“We say that

A is a (t, q) -adversary if

A runs in at most t steps and

makes at most q queries to

rn 3
1
4
1
5
9
2
6

Students in algorithm courses learn to count executed “steps”. Skipped branches take 0 “steps”.

This algorithm uses 4 “steps”.

Generalization: There exists an algorithm that, given $n < 2^k$, prints the n th digit of π using $k + 1$ “steps”.

Variant: There exists a 258-“step” P-256 discrete-log attack (with 100% success probability). If “time” means “steps” then the standard conjectures are wrong.

1994 Bellare–Kilian–Rogaway:

“We say that

A is a (t, q) -adversary if

A runs in at most t steps and makes at most q queries to \mathcal{O} .”

Students in algorithm courses learn to count executed “steps”. Skipped branches take 0 “steps”.

This algorithm uses 4 “steps”.

Generalization: There exists an algorithm that, given $n < 2^k$, prints the n th digit of π using $k + 1$ “steps”.

Variant: There exists a 258-“step” P-256 discrete-log attack (with 100% success probability). If “time” means “steps” then the standard conjectures are wrong.

1994 Bellare–Kilian–Rogaway:

“We say that

A is a (t, q) -adversary if

A runs in at most t steps and makes at most q queries to \mathcal{O} .”

Oops: table-lookup attack has very small t .

Paper conjectured “useful” DES security bounds. Any reasonable interpretation of conjecture was false, given paper’s definition. Theorems in paper were vacuous.

s in algorithm courses
count executed “steps” .
branches take 0 “steps” .

algorithm uses 4 “steps” .

ization: There exists an
m that, given $n < 2^k$,
the n th digit of π
+ 1 “steps” .

There exists a 258-
256 discrete-log attack
(100% success probability).
' means “steps” then the
conjectures are wrong.

1994 Bellare–Kilian–Rogaway:

“We say that

A is a (t, q) -adversary if

A runs in at most t steps and

makes at most q queries to \mathcal{O} .”

Oops: table-lookup attack
has very small t .

Paper conjectured “useful” DES
security bounds. Any reasonable
interpretation of conjecture was
false, given paper’s definition.
Theorems in paper were vacuous.

2000 Be

“We fix

Access I

model o

running

executio

of A’s d

convent

caused [

tables . .

chm courses
cuted “steps” .
take 0 “steps” .
es 4 “steps” .
here exists an
ven $n < 2^k$,
it of π
s” .
ists a 258-
rete-log attack
ss probability).
steps” then the
res are wrong.

1994 Bellare–Kilian–Rogaway:

“We say that

A is a (t, q) -adversary if

*A runs in at most t steps and
makes at most q queries to \mathcal{O} .”*

Oops: table-lookup attack
has very small t .

Paper conjectured “useful” DES
security bounds. Any reasonable
interpretation of conjecture was
false, given paper’s definition.
Theorems in paper were vacuous.

2000 Bellare–Kilian

*“We fix some part
Access Machine (M)
model of computa
running time [mea
execution time plu
of A’s description
convention elimina
caused [by] arbitra
tables”*

1994 Bellare–Kilian–Rogaway:

*“We say that
A is a (t, q) -adversary if
A runs in at most t steps and
makes at most q queries to \mathcal{O} .”*

Oops: table-lookup attack
has very small t .

Paper conjectured “useful” DES
security bounds. Any reasonable
interpretation of conjecture was
false, given paper’s definition.
Theorems in paper were vacuous.

2000 Bellare–Kilian–Rogaway:

*“We fix some particular Random Access Machine (RAM) as a
model of computation. . . . A’s
running time [means] A’s actual
execution time plus the length
of A’s description . . . This
convention eliminates pathologies
caused [by] arbitrarily large
tables . . .”*

1994 Bellare–Kilian–Rogaway:

“We say that

A is a (t, q) -adversary if

*A runs in at most t steps and
makes at most q queries to \mathcal{O} .”*

Oops: table-lookup attack
has very small t .

Paper conjectured “useful” DES
security bounds. Any reasonable
interpretation of conjecture was
false, given paper’s definition.

Theorems in paper were vacuous.

2000 Bellare–Kilian–Rogaway:

“We fix some particular Random

Access Machine (RAM) as a

*model of computation. . . . A’s
running time [means] A’s actual*

execution time plus the length

of A’s description . . . This

*convention eliminates pathologies
caused [by] arbitrarily large lookup
tables . . .”*

1994 Bellare–Kilian–Rogaway:

“We say that

A is a (t, q) -adversary if

*A runs in at most t steps and
makes at most q queries to \mathcal{O} .”*

Oops: table-lookup attack
has very small t .

Paper conjectured “useful” DES
security bounds. Any reasonable
interpretation of conjecture was
false, given paper’s definition.
Theorems in paper were vacuous.

2000 Bellare–Kilian–Rogaway:

“We fix some particular Random

Access Machine (RAM) as a

*model of computation. . . . A’s
running time [means] A’s actual*

execution time plus the length

of A’s description . . . This

*convention eliminates pathologies
caused [by] arbitrarily large lookup
tables . . .”*

Main point of our paper:

There are more pathologies!

Illustrative example: ECDL.

Bellare–Kilian–Rogaway:

that

(t, q)-adversary if

runs at most t steps and

makes at most q queries to \mathcal{O} .”

table-lookup attack

small t .

conjectured “useful” DES

bounds. Any reasonable

refutation of conjecture was

given paper’s definition.

Arguments in paper were vacuous.

2000 Bellare–Kilian–Rogaway:

“We fix some particular Random

Access Machine (RAM) as a

model of computation. . . . A’s

running time [means] A’s actual

execution time plus the length

of A’s description . . . This

convention eliminates pathologies

caused [by] arbitrarily large lookup

tables”

Main point of our paper:

There are more pathologies!

Illustrative example: ECDL.

The rho

Simplified

Make a

R_0, R_1, \dots

where c_i

the next

Birthday

Random

elements

after about

P-256: $\approx 2^{128}$

The walk

Cycle-finding

(e.g., Floyd)

n–Rogaway:

sary if

t steps and

queries to \mathcal{O} .”

p attack

“useful” DES

Any reasonable

onjecture was

s definition.

r were vacuous.

2000 Bellare–Kilian–Rogaway:

“We fix some particular Random

Access Machine (RAM) as a

model of computation. . . . A’s

running time [means] A’s actual

execution time plus the length

of A’s description . . . This

convention eliminates pathologies

caused [by] arbitrarily large lookup

tables”

Main point of our paper:

There are more pathologies!

Illustrative example: ECDL.

The rho method

Simplified, non-pa

Make a pseudo-ran

R_0, R_1, R_2, \dots in t

where current poin

the next point: R_i

Birthday paradox:

Randomly choosin

elements picks one

after about $\sqrt{\pi \ell / 2}$

P-256: $\ell \approx 2^{256}$ so

The walk now ent

Cycle-finding algo

(e.g., Floyd) quick

2000 Bellare–Kilian–Rogaway:

“We fix some particular Random Access Machine (RAM) as a model of computation. . . . A’s running time [means] A’s actual execution time plus the length of A’s description . . . This convention eliminates pathologies caused [by] arbitrarily large lookup tables . . .”

Main point of our paper:

There are more pathologies!

Illustrative example: ECDL.

The rho method

Simplified, non-parallel rho:

Make a pseudo-random walk R_0, R_1, R_2, \dots in the group where current point determines the next point: $R_{i+1} = f(R_i)$

Birthday paradox:

Randomly choosing from ℓ elements picks one element after about $\sqrt{\pi\ell/2}$ draws.

P-256: $\ell \approx 2^{256}$ so $\approx 2^{128}$ d

The walk now enters a cycle

Cycle-finding algorithm

(e.g., Floyd) quickly detects

2000 Bellare–Kilian–Rogaway:

“We fix some particular Random Access Machine (RAM) as a model of computation. . . . A’s running time [means] A’s actual execution time plus the length of A’s description . . . This convention eliminates pathologies caused [by] arbitrarily large lookup tables . . .”

Main point of our paper:

There are more pathologies!

Illustrative example: ECDDL.

The rho method

Simplified, non-parallel rho:

Make a pseudo-random walk R_0, R_1, R_2, \dots in the group $\langle P \rangle$, where current point determines the next point: $R_{i+1} = f(R_i)$.

Birthday paradox:

Randomly choosing from ℓ elements picks one element twice after about $\sqrt{\pi\ell/2}$ draws.

P-256: $\ell \approx 2^{256}$ so $\approx 2^{128}$ draws.

The walk now enters a cycle.

Cycle-finding algorithm

(e.g., Floyd) quickly detects this.

Illare–Kilian–Rogaway:

some particular Random Machine (RAM) as a

f computation. . . . A's

time [means] A's actual

n time plus the length

description . . . This

ion eliminates pathologies

[by] arbitrarily large lookup

."

oint of our paper:

re more pathologies!

ve example: ECDL.

The rho method

Simplified, non-parallel rho:

Make a pseudo-random walk

R_0, R_1, R_2, \dots in the group $\langle P \rangle$,

where current point determines

the next point: $R_{i+1} = f(R_i)$.

Birthday paradox:

Randomly choosing from ℓ

elements picks one element twice

after about $\sqrt{\pi\ell/2}$ draws.

P-256: $\ell \approx 2^{256}$ so $\approx 2^{128}$ draws.

The walk now enters a cycle.

Cycle-finding algorithm

(e.g., Floyd) quickly detects this.

Goal: Co

Assume

we know

so that

Then R_i

$y_i P + x_i$

so $(y_i -$

If $x_i \neq x_j$

$\log_P Q =$

n-Rogaway:
Particular Random
RAM) as a
tion. ... A's
ns] A's actual
is the length
... This
ates pathologies
arily large lookup
paper:
athologies!
e: ECDL.

The rho method

Simplified, non-parallel rho:

Make a pseudo-random walk
 R_0, R_1, R_2, \dots in the group $\langle P \rangle$,
where current point determines
the next point: $R_{i+1} = f(R_i)$.

Birthday paradox:

Randomly choosing from ℓ
elements picks one element twice
after about $\sqrt{\pi\ell/2}$ draws.

P-256: $\ell \approx 2^{256}$ so $\approx 2^{128}$ draws.

The walk now enters a cycle.

Cycle-finding algorithm
(e.g., Floyd) quickly detects this.

Goal: Compute $\log_P Q$

Assume that for each i
we know $x_i, y_i \in \mathbb{Z}$
so that $R_i = y_i P + x_i Q$

Then $R_i = R_j$ means
 $y_i P + x_i Q = y_j P + x_j Q$

so $(y_i - y_j)P = (x_j - x_i)Q$
If $x_i \neq x_j$ the DLP is
 $\log_P Q = (y_j - y_i)P / (x_j - x_i)Q$

The rho method

Simplified, non-parallel rho:

Make a pseudo-random walk

R_0, R_1, R_2, \dots in the group $\langle P \rangle$,

where current point determines

the next point: $R_{i+1} = f(R_i)$.

Birthday paradox:

Randomly choosing from ℓ

elements picks one element twice

after about $\sqrt{\pi\ell/2}$ draws.

P-256: $\ell \approx 2^{256}$ so $\approx 2^{128}$ draws.

The walk now enters a cycle.

Cycle-finding algorithm

(e.g., Floyd) quickly detects this.

Goal: Compute $\log_P Q$.

Assume that for each i

we know $x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$

so that $R_i = y_i P + x_i Q$.

Then $R_i = R_j$ means that

$$y_i P + x_i Q = y_j P + x_j Q$$

$$\text{so } (y_i - y_j)P = (x_j - x_i)Q$$

If $x_i \neq x_j$ the DLP is solved

$$\log_P Q = (y_j - y_i)/(x_i - x_j)$$

The rho method

Simplified, non-parallel rho:

Make a pseudo-random walk R_0, R_1, R_2, \dots in the group $\langle P \rangle$, where current point determines the next point: $R_{i+1} = f(R_i)$.

Birthday paradox:

Randomly choosing from ℓ elements picks one element twice after about $\sqrt{\pi\ell/2}$ draws.

P-256: $\ell \approx 2^{256}$ so $\approx 2^{128}$ draws.

The walk now enters a cycle.

Cycle-finding algorithm

(e.g., Floyd) quickly detects this.

Goal: Compute $\log_P Q$.

Assume that for each i we know $x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$ so that $R_i = y_i P + x_i Q$.

Then $R_i = R_j$ means that

$$y_i P + x_i Q = y_j P + x_j Q$$

$$\text{so } (y_i - y_j)P = (x_j - x_i)Q.$$

If $x_i \neq x_j$ the DLP is solved:

$$\log_P Q = (y_j - y_i)/(x_i - x_j).$$

The rho method

Simplified, non-parallel rho:

Make a pseudo-random walk R_0, R_1, R_2, \dots in the group $\langle P \rangle$, where current point determines the next point: $R_{i+1} = f(R_i)$.

Birthday paradox:

Randomly choosing from ℓ elements picks one element twice after about $\sqrt{\pi\ell/2}$ draws.

P-256: $\ell \approx 2^{256}$ so $\approx 2^{128}$ draws.

The walk now enters a cycle.

Cycle-finding algorithm

(e.g., Floyd) quickly detects this.

Goal: Compute $\log_P Q$.

Assume that for each i we know $x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$ so that $R_i = y_i P + x_i Q$.

Then $R_i = R_j$ means that

$$y_i P + x_i Q = y_j P + x_j Q$$

$$\text{so } (y_i - y_j)P = (x_j - x_i)Q.$$

If $x_i \neq x_j$ the DLP is solved:

$$\log_P Q = (y_j - y_i)/(x_i - x_j).$$

e.g. “base- (P, Q) r -adding walk”:

precompute S_1, S_2, \dots, S_r

as random combinations $aP + bQ$;

define $f(R) = R + S_{H(R)}$

where H hashes to $\{1, 2, \dots, r\}$.

method

ed, non-parallel rho:

pseudo-random walk

R_2, \dots in the group $\langle P \rangle$,

current point determines

point: $R_{i+1} = f(R_i)$.

paradox:

ly choosing from ℓ

s picks one element twice

out $\sqrt{\pi\ell/2}$ draws.

$\ell \approx 2^{256}$ so $\approx 2^{128}$ draws.

k now enters a cycle.

nding algorithm

oyd) quickly detects this.

Goal: Compute $\log_P Q$.

Assume that for each i

we know $x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$

so that $R_i = y_i P + x_i Q$.

Then $R_i = R_j$ means that

$$y_i P + x_i Q = y_j P + x_j Q$$

$$\text{so } (y_i - y_j)P = (x_j - x_i)Q.$$

If $x_i \neq x_j$ the DLP is solved:

$$\log_P Q = (y_j - y_i)/(x_i - x_j).$$

e.g. “base- (P, Q) r -adding walk”:

precompute S_1, S_2, \dots, S_r

as random combinations $aP + bQ$;

define $f(R) = R + S_{H(R)}$

where H hashes to $\{1, 2, \dots, r\}$.

Parallel

1994 var

Declare

the set o

e.g., all

bits of r

Perform,

different

but sam

Termina

once it h

Report p

Server re

all distin

parallel rho:

random walk

the group $\langle P \rangle$,

which determines

$$R_{i+1} = f(R_i).$$

g from ℓ

the element twice

$\sqrt{2}$ draws.

to $\approx 2^{128}$ draws.

enters a cycle.

algorithm

quickly detects this.

Goal: Compute $\log_P Q$.

Assume that for each i

we know $x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$

so that $R_i = y_i P + x_i Q$.

Then $R_i = R_j$ means that

$$y_i P + x_i Q = y_j P + x_j Q$$

$$\text{so } (y_i - y_j)P = (x_j - x_i)Q.$$

If $x_i \neq x_j$ the DLP is solved:

$$\log_P Q = (y_j - y_i) / (x_i - x_j).$$

e.g. “base- (P, Q) r -adding walk”:

precompute S_1, S_2, \dots, S_r

as random combinations $aP + bQ$;

define $f(R) = R + S_{H(R)}$

where H hashes to $\{1, 2, \dots, r\}$.

Parallel rho

1994 van Oorschot

Declare some subset

the set of *distinguished*

e.g., all $R \in \langle P \rangle$ with

bits of representation

Perform, in parallel

different starting points

but same update function

Terminate each walk

once it hits a distinguished

Report point to central

Server receives, stores

all distinguished points

Goal: Compute $\log_P Q$.

Assume that for each i
we know $x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$
so that $R_i = y_i P + x_i Q$.

Then $R_i = R_j$ means that
 $y_i P + x_i Q = y_j P + x_j Q$
so $(y_i - y_j)P = (x_j - x_i)Q$.
If $x_i \neq x_j$ the DLP is solved:
 $\log_P Q = (y_j - y_i)/(x_i - x_j)$.

e.g. “base- (P, Q) r -adding walk”:
precompute S_1, S_2, \dots, S_r
as random combinations $aP + bQ$;
define $f(R) = R + S_{H(R)}$
where H hashes to $\{1, 2, \dots, r\}$.

Parallel rho

1994 van Oorschot–Wiener:

Declare some subset of $\langle P \rangle$
the set of *distinguished points*

e.g., all $R \in \langle P \rangle$ where last
bits of representation of R are

Perform, in parallel, walks for
different starting points $Q +$
but same update function f

Terminate each walk
once it hits a distinguished point

Report point to central server

Server receives, stores, and sorts
all distinguished points.

Goal: Compute $\log_P Q$.

Assume that for each i

we know $x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$

so that $R_i = y_i P + x_i Q$.

Then $R_i = R_j$ means that

$$y_i P + x_i Q = y_j P + x_j Q$$

$$\text{so } (y_i - y_j)P = (x_j - x_i)Q.$$

If $x_i \neq x_j$ the DLP is solved:

$$\log_P Q = (y_j - y_i)/(x_i - x_j).$$

e.g. “base- (P, Q) r -adding walk”:

precompute S_1, S_2, \dots, S_r

as random combinations $aP + bQ$;

define $f(R) = R + S_{H(R)}$

where H hashes to $\{1, 2, \dots, r\}$.

Parallel rho

1994 van Oorschot–Wiener:

Declare some subset of $\langle P \rangle$ to be the set of *distinguished points*:

e.g., all $R \in \langle P \rangle$ where last 20 bits of representation of R are 0.

Perform, in parallel, walks for different starting points $Q + yP$ but same update function f .

Terminate each walk

once it hits a distinguished point.

Report point to central server.

Server receives, stores, and sorts all distinguished points.

compute $\log_P Q$.

that for each i

$$x_i, y_i \in \mathbf{Z}/\ell\mathbf{Z}$$

$$R_i = y_i P + x_i Q.$$

$R_i = R_j$ means that

$$y_i Q = y_j P + x_j Q$$

$$(y_j - y_i)P = (x_j - x_i)Q.$$

For x_j the DLP is solved:

$$x_j = (y_j - y_i)/(x_i - x_j).$$

use “ (P, Q) r -adding walk”:

compute S_1, S_2, \dots, S_r

from combinations $aP + bQ$;

$$H(R) = R + S_{H(R)}$$

H hashes to $\{1, 2, \dots, r\}$.

Parallel rho

1994 van Oorschot–Wiener:

Declare some subset of $\langle P \rangle$ to be the set of *distinguished points*:

e.g., all $R \in \langle P \rangle$ where last 20 bits of representation of R are 0.

Perform, in parallel, walks for different starting points $Q + yP$ but same update function f .

Terminate each walk

once it hits a distinguished point.

Report point to central server.

Server receives, stores, and sorts all distinguished points.

State of

Can break

ℓ in $\sqrt{\pi}$

Use nega

factor $\sqrt{\ell}$

Solving

takes ≈ 2

This is t

cryptana

$g_P Q$.

each i

$\mathbf{Z}/\ell\mathbf{Z}$

$+ x_i Q$.

means that

$+ x_j Q$

$(x_j - x_i)Q$.

P is solved:

$)/(x_i - x_j)$.

" r -adding walk":

$2, \dots, S_r$

relations $aP + bQ$;

$+ S_{H(R)}$

$\{1, 2, \dots, r\}$.

Parallel rho

1994 van Oorschot–Wiener:

Declare some subset of $\langle P \rangle$ to be the set of *distinguished points*:

e.g., all $R \in \langle P \rangle$ where last 20 bits of representation of R are 0.

Perform, in parallel, walks for different starting points $Q + yP$ but same update function f .

Terminate each walk once it hits a distinguished point.

Report point to central server.

Server receives, stores, and sorts all distinguished points.

State of the art

Can break DLP in ℓ in $\sqrt{\pi\ell/2}$ group

Use negation map factor $\sqrt{2}$ for elliptic

Solving DLP on N takes $\approx 2^{128}$ group

This is the best all *cryptanalysts have*

Parallel rho

1994 van Oorschot–Wiener:

Declare some subset of $\langle P \rangle$ to be

the set of *distinguished points*:

e.g., all $R \in \langle P \rangle$ where last 20 bits of representation of R are 0.

Perform, in parallel, walks for different starting points $Q + yP$ but same update function f .

Terminate each walk

once it hits a distinguished point.

Report point to central server.

Server receives, stores, and sorts all distinguished points.

State of the art

Can break DLP in group of ℓ in $\sqrt{\pi\ell/2}$ group operation

Use negation map to gain factor $\sqrt{2}$ for elliptic curves.

Solving DLP on NIST P-256 takes $\approx 2^{128}$ group operation

This is the best algorithm that *cryptanalysts have published*

Parallel rho

1994 van Oorschot–Wiener:

Declare some subset of $\langle P \rangle$ to be the set of *distinguished points*:

e.g., all $R \in \langle P \rangle$ where last 20 bits of representation of R are 0.

Perform, in parallel, walks for different starting points $Q + yP$ but same update function f .

Terminate each walk once it hits a distinguished point.

Report point to central server.

Server receives, stores, and sorts all distinguished points.

State of the art

Can break DLP in group of order ℓ in $\sqrt{\pi\ell/2}$ group operations.

Use negation map to gain factor $\sqrt{2}$ for elliptic curves.

Solving DLP on NIST P-256 takes $\approx 2^{128}$ group operations.

This is the best algorithm that *cryptanalysts have published*.

Parallel rho

1994 van Oorschot–Wiener:

Declare some subset of $\langle P \rangle$ to be the set of *distinguished points*:

e.g., all $R \in \langle P \rangle$ where last 20 bits of representation of R are 0.

Perform, in parallel, walks for different starting points $Q + yP$ but same update function f .

Terminate each walk once it hits a distinguished point.

Report point to central server.

Server receives, stores, and sorts all distinguished points.

State of the art

Can break DLP in group of order ℓ in $\sqrt{\pi\ell/2}$ group operations.

Use negation map to gain factor $\sqrt{2}$ for elliptic curves.

Solving DLP on NIST P-256 takes $\approx 2^{128}$ group operations.

This is the best algorithm that *cryptanalysts have published*.

But is it the best algorithm that *exists*?

rho

in Oorschot–Wiener:

some subset of $\langle P \rangle$ to be

of *distinguished points*:

$R \in \langle P \rangle$ where last 20

representation of R are 0.

, in parallel, walks for

starting points $Q + yP$

the update function f .

ate each walk

hits a distinguished point.

point to central server.

receives, stores, and sorts

guished points.

State of the art

Can break DLP in group of order ℓ in $\sqrt{\pi\ell/2}$ group operations.

Use negation map to gain factor $\sqrt{2}$ for elliptic curves.

Solving DLP on NIST P-256 takes $\approx 2^{128}$ group operations.

This is the best algorithm that *cryptanalysts have published*.

But is it the best algorithm that *exists*?

This paper

Assuming
overwhelms
computer

There exists
algorithm
and has

“Time”

Inescapable
standard

P-256 ECDHE

t–Wiener:

set of $\langle P \rangle$ to be

disjoint points:

where last 20

of R are 0.

el, walks for

points $Q+yP$

function f .

alk

nguished point.

entral server.

ores, and sorts

oints.

State of the art

Can break DLP in group of order ℓ in $\sqrt{\pi\ell/2}$ group operations.

Use negation map to gain factor $\sqrt{2}$ for elliptic curves.

Solving DLP on NIST P-256 takes $\approx 2^{128}$ group operations.

This is the best algorithm that *cryptanalysts have published*.

But is it the best algorithm that *exists*?

This paper's ECDL

Assuming plausible overwhelmingly verifiable computer experiments

There exists a P-256 algorithm that takes $\approx 2^{128}$ and has success probability

“Time” includes a

Inescapable conclusion

standard conjecture

P-256 ECDL hardness

ECDHE security, e

State of the art

Can break DLP in group of order ℓ in $\sqrt{\pi\ell/2}$ group operations.

Use negation map to gain factor $\sqrt{2}$ for elliptic curves.

Solving DLP on NIST P-256 takes $\approx 2^{128}$ group operations.

This is the best algorithm that *cryptanalysts have published*.

But is it the best algorithm that *exists*?

This paper's ECDL algorithm

Assuming plausible heuristic overwhelmingly verified by computer experiment:

There exists a P-256 ECDL algorithm that takes “time” and has success probability $\approx 1/2$.

“Time” includes algorithm ℓ .

Inescapable conclusion: **The standard conjectures** (regarding P-256 ECDL hardness, P-256 ECDHE security, etc.) **are false**.

State of the art

Can break DLP in group of order ℓ in $\sqrt{\pi\ell/2}$ group operations.

Use negation map to gain factor $\sqrt{2}$ for elliptic curves.

Solving DLP on NIST P-256 takes $\approx 2^{128}$ group operations.

This is the best algorithm that *cryptanalysts have published*.

But is it the best algorithm that *exists*?

This paper's ECDL algorithms

Assuming plausible heuristics, overwhelmingly verified by computer experiment:

There exists a P-256 ECDL algorithm that takes “time” $\approx 2^{85}$ and has success probability ≈ 1 .

“Time” includes algorithm length.

Inescapable conclusion: **The standard conjectures** (regarding P-256 ECDL hardness, P-256 ECDHE security, etc.) **are false**.

the art

break DLP in group of order ℓ with $\sqrt{\ell}/2$ group operations.

isogeny map to gain $\sqrt{2}$ for elliptic curves.

break DLP on NIST P-256 with 2^{128} group operations.

is the best algorithm that cryptanalysts have published.

is the best algorithm known to exist?

This paper's ECDL algorithms

Assuming plausible heuristics, are overwhelmingly verified by computer experiment:

There exists a P-256 ECDL algorithm that takes "time" $\approx 2^{85}$ and has success probability ≈ 1 .

"Time" includes algorithm length.

Inescapable conclusion: **The standard conjectures** (regarding P-256 ECDL hardness, P-256 ECDHE security, etc.) **are false.**

Should we be worried about P-256 ECDL?

No!

We have a prime p that is prime, but B takes \sqrt{p} to factor.

We conjecture that nobody knows a better algorithm.

group of order
operations.

to gain

tic curves.

IST P-256

operations.

gorithm that

published.

algorithm

This paper's ECDL algorithms

Assuming plausible heuristics,
overwhelmingly verified by
computer experiment:

There exists a P-256 ECDL
algorithm that takes “time” $\approx 2^{85}$
and has success probability ≈ 1 .

“Time” includes algorithm length.

Inescapable conclusion: **The
standard conjectures** (regarding
P-256 ECDL hardness, P-256
ECDHE security, etc.) **are false.**

Should P-256 ECDL
be worried about the
P-256 ECDL algorithm?

No!

We have a program
that prints out A ,
but B takes “time”

We conjecture that
nobody will ever print

This paper's ECDL algorithms

Assuming plausible heuristics,
overwhelmingly verified by
computer experiment:

There exists a P-256 ECDL
algorithm that takes "time" $\approx 2^{85}$
and has success probability ≈ 1 .

"Time" includes algorithm length.

Inescapable conclusion: **The
standard conjectures** (regarding
P-256 ECDL hardness, P-256
ECDHE security, etc.) **are false.**

Should P-256 ECDHE users
be worried about this
P-256 ECDL algorithm A ?

No!

We have a program B
that prints out A ,
but B takes "time" $\approx 2^{170}$.

We conjecture that
nobody will ever print out A

This paper's ECDL algorithms

Assuming plausible heuristics,
overwhelmingly verified by
computer experiment:

There exists a P-256 ECDL
algorithm that takes “time” $\approx 2^{85}$
and has success probability ≈ 1 .

“Time” includes algorithm length.

Inescapable conclusion: **The
standard conjectures** (regarding
P-256 ECDL hardness, P-256
ECDHE security, etc.) **are false.**

Should P-256 ECDHE users
be worried about this
P-256 ECDL algorithm A ?

No!

We have a program B
that prints out A ,
but B takes “time” $\approx 2^{170}$.

We conjecture that
nobody will ever print out A .

This paper's ECDL algorithms

Assuming plausible heuristics,
overwhelmingly verified by
computer experiment:

There exists a P-256 ECDL
algorithm that takes “time” $\approx 2^{85}$
and has success probability ≈ 1 .

“Time” includes algorithm length.

Inescapable conclusion: **The
standard conjectures** (regarding
P-256 ECDL hardness, P-256
ECDHE security, etc.) **are false.**

Should P-256 ECDHE users
be worried about this
P-256 ECDL algorithm A ?

No!

We have a program B
that prints out A ,
but B takes “time” $\approx 2^{170}$.

We conjecture that
nobody will ever print out A .

But A exists, and the standard
conjecture doesn't see the 2^{170} .

Other's ECDL algorithms

g plausible heuristics,
lmingly verified by
er experiment:

exists a P-256 ECDL
m that takes "time" $\approx 2^{85}$
success probability ≈ 1 .

includes algorithm length.

able conclusion: **The
d conjectures** (regarding
CDL hardness, P-256
security, etc.) **are false.**

Should P-256 ECDHE users
be worried about this
P-256 ECDL algorithm A ?

No!

We have a program B
that prints out A ,
but B takes "time" $\approx 2^{170}$.

We conjecture that
nobody will ever print out A .

But A exists, and the standard
conjecture doesn't see the 2^{170} .

Cryptana

Common
a 2^{170} "
(independ
a 2^{85} "n

For cryp
 2^{170} , mu

For the s
definitio
The mai
much be

L algorithms

e heuristics,
rified by

ent:

56 ECDL

es “time” $\approx 2^{85}$

robability ≈ 1 .

Algorithm length.

ision: **The**

ures (regarding

ness, P-256

etc.) **are false.**

Should P-256 ECDHE users
be worried about this
P-256 ECDL algorithm A ?

No!

We have a program B
that prints out A ,
but B takes “time” $\approx 2^{170}$.

We conjecture that
nobody will ever print out A .

But A *exists*, and the standard
conjecture doesn't see the 2^{170} .

Cryptanalysts *do* s

Common parlance
a 2^{170} “precomput
(independent of Q
a 2^{85} “main comp

For cryptanalysts:
 2^{170} , much worse

For the standard s
definitions and cor
The main computa
much better than

Should P-256 ECDHE users
be worried about this
P-256 ECDL algorithm A ?

No!

We have a program B
that prints out A ,
but B takes “time” $\approx 2^{170}$.

We conjecture that
nobody will ever print out A .

But A exists, and the standard
conjecture doesn't see the 2^{170} .

Cryptanalysts *do* see the 2^{170}

Common parlance: We have
a 2^{170} “precomputation”
(independent of Q) followed
a 2^{85} “main computation”.

For cryptanalysts: This costs
 2^{170} , much worse than 2^{128} .

For the standard security
definitions and conjectures:
The main computation costs
much better than 2^{128} .

Should P-256 ECDHE users
be worried about this
P-256 ECDL algorithm A ?

No!

We have a program B
that prints out A ,
but B takes “time” $\approx 2^{170}$.

We conjecture that
nobody will ever print out A .

But A *exists*, and the standard
conjecture doesn't see the 2^{170} .

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have
a 2^{170} “precomputation”
(independent of Q) followed by
a 2^{85} “main computation”.

For cryptanalysts: This costs
 2^{170} , much worse than 2^{128} .

For the standard security
definitions and conjectures:

The main computation costs 2^{85} ,
much better than 2^{128} .

P-256 ECDHE users
wondered about this
CDL algorithm A ?

Is there a program B
that prints out A ,
and takes “time” $\approx 2^{170}$.

Is there a procedure that
will ever print out A .

exists, and the standard
definition of time doesn't see the 2^{170} .

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have
a 2^{170} “precomputation”
(independent of Q) followed by
a 2^{85} “main computation”.

For cryptanalysts: This costs
 2^{170} , much worse than 2^{128} .

For the standard security
definitions and conjectures:

The main computation costs 2^{85} ,
much better than 2^{128} .

Almost all security
definitions
redefine
security
on P on
chosen
 c_i chosen

DHE users

this

algorithm A ?

algorithm B

" $\approx 2^{170}$.

print out A .

the standard
see the 2^{170} .

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have
a 2^{170} "precomputation"
(independent of Q) followed by
a 2^{85} "main computation".

For cryptanalysts: This costs
 2^{170} , much worse than 2^{128} .

For the standard security
definitions and conjectures:

The main computation costs 2^{85} ,
much better than 2^{128} .

Almost standard w
redefine steps S_i t
on P only; i.e., S_i
 c_i chosen uniforml

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have a 2^{170} “precomputation” (independent of Q) followed by a 2^{85} “main computation”.

For cryptanalysts: This costs 2^{170} , much worse than 2^{128} .

For the standard security definitions and conjectures:

The main computation costs 2^{85} , much better than 2^{128} .

Almost standard walk functions
redefine steps S_i to depend
on P only; i.e., $S_i = c_i P$ with
 c_i chosen uniformly at random.

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have a 2^{170} “precomputation” (independent of Q) followed by a 2^{85} “main computation”.

For cryptanalysts: This costs 2^{170} , much worse than 2^{128} .

For the standard security definitions and conjectures:

The main computation costs 2^{85} , much better than 2^{128} .

Almost standard walk function: redefine steps S_i to depend on P only; i.e., $S_i = c_i P$ with c_i chosen uniformly at random.

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have a 2^{170} “precomputation” (independent of Q) followed by a 2^{85} “main computation”.

For cryptanalysts: This costs 2^{170} , much worse than 2^{128} .

For the standard security definitions and conjectures:

The main computation costs 2^{85} , much better than 2^{128} .

Almost standard walk function: redefine steps S_i to depend on P only; i.e., $S_i = c_i P$ with c_i chosen uniformly at random.

Precomputation:

Start some walks at yP for random choices of y .

Build table of distinct distinguished points D along with $\log_P D$.

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have a 2^{170} “precomputation” (independent of Q) followed by a 2^{85} “main computation”.

For cryptanalysts: This costs 2^{170} , much worse than 2^{128} .

For the standard security definitions and conjectures:

The main computation costs 2^{85} , much better than 2^{128} .

Almost standard walk function: redefine steps S_i to depend on P only; i.e., $S_i = c_i P$ with c_i chosen uniformly at random.

Precomputation:

Start some walks at yP for random choices of y .

Build table of distinct distinguished points D along with $\log_P D$.

Main computation:

Starting from Q , walk to distinguished point $Q + yP$.

Check for $Q + yP$ in table.

Cryptanalysts *do* see the 2^{170} .

Common parlance: We have a 2^{170} “precomputation” (independent of Q) followed by a 2^{85} “main computation”.

For cryptanalysts: This costs 2^{170} , much worse than 2^{128} .

For the standard security definitions and conjectures:

The main computation costs 2^{85} , much better than 2^{128} .

Almost standard walk function: redefine steps S_i to depend on P only; i.e., $S_i = c_i P$ with c_i chosen uniformly at random.

Precomputation:

Start some walks at yP for random choices of y .

Build table of distinct distinguished points D along with $\log_P D$.

Main computation:

Starting from Q , walk to distinguished point $Q + yP$.

Check for $Q + yP$ in table.

(If this fails, rerandomize Q .)

analysts *do* see the 2^{170} .

in parlance: We have
"precomputation"
(independent of Q) followed by
"main computation".

analysts: This costs
much worse than 2^{128} .

standard security

assumptions and conjectures:

main computation costs 2^{85} ,
better than 2^{128} .

Almost standard walk function:
redefine steps S_i to depend
on P only; i.e., $S_i = c_i P$ with
 c_i chosen uniformly at random.

Precomputation:

Start some walks at yP
for random choices of y .

Build table of distinct
distinguished points D
along with $\log_p D$.

Main computation:

Starting from Q , walk to
distinguished point $Q + yP$.

Check for $Q + yP$ in table.

(If this fails, rerandomize Q .)

What you

P-256 is

There ex

AES-128

at cost b

e.g., tim

(Assumi

\Rightarrow Very

between

and actu

Also: An

for fixing

[eprint.](#)

see the 2^{170} .

: We have
tation”

) followed by
utation”.

This costs
than 2^{128} .

security

jectures:

ation costs 2^{85} ,
 2^{128} .

Almost standard walk function:
redefine steps S_i to depend
on P only; i.e., $S_i = c_i P$ with
 c_i chosen uniformly at random.

Precomputation:

Start some walks at yP
for random choices of y .

Build table of distinct
distinguished points D
along with $\log_P D$.

Main computation:

Starting from Q , walk to
distinguished point $Q + yP$.

Check for $Q + yP$ in table.

(If this fails, rerandomize Q .)

What you find in t

P-256 isn't the on

There *exist* algorit

AES-128, RSA-307

at cost below 2^{128}

e.g., time 2^{85} to b

(Assuming standar

\Rightarrow Very large separa

between standard

and actual security

Also: Analysis of v

for fixing the defin

eprint.iacr.org

Almost standard walk function:
redefine steps S_i to depend
on P only; i.e., $S_i = c_i P$ with
 c_i chosen uniformly at random.

Precomputation:

Start some walks at yP
for random choices of y .

Build table of distinct
distinguished points D
along with $\log_P D$.

Main computation:

Starting from Q , walk to
distinguished point $Q + yP$.

Check for $Q + yP$ in table.

(If this fails, rerandomize Q .)

What you find in the full pa
P-256 isn't the only problem
There *exist* algorithms break
AES-128, RSA-3072, DSA-3
at cost below 2^{128} ;

e.g., time 2^{85} to break AES.
(Assuming standard heuristics)

\Rightarrow Very large separation
between standard definition
and actual security.

Also: Analysis of various ide
for fixing the definitions.

eprint.iacr.org/2012/3

Almost standard walk function:
redefine steps S_i to depend
on P only; i.e., $S_i = c_i P$ with
 c_i chosen uniformly at random.

Precomputation:

Start some walks at yP
for random choices of y .

Build table of distinct
distinguished points D
along with $\log_P D$.

Main computation:

Starting from Q , walk to
distinguished point $Q + yP$.

Check for $Q + yP$ in table.

(If this fails, rerandomize Q .)

What you find in the full paper:

P-256 isn't the only problem!

There *exist* algorithms breaking
AES-128, RSA-3072, DSA-3072
at cost below 2^{128} ;

e.g., time 2^{85} to break AES.

(Assuming standard heuristics.)

\Rightarrow Very large separation
between standard definition
and actual security.

Also: Analysis of various ideas
for fixing the definitions.

eprint.iacr.org/2012/318