

Constant-time  
square-and-multiply

D. J. Bernstein

University of Illinois at Chicago;  
Ruhr University Bochum

---

```
def pow256bit(x,e):  
    y = 1  
    for i in reversed(range(256)):  
        y = y*y  
        if 1&(e>>i):  
            y = y*x  
    return y
```

This code uses 256 squarings,  
plus 1 extra multiplication  
for each bit set in  $e$ .

Problem when  $e$  is secret: time  
leaks number of bits set in  $e$ .

# Constant-time square-and-multiply

D. J. Bernstein

University of Illinois at Chicago;  
Ruhr University Bochum

---

```
def pow256bit(x,e):  
    y = 1  
    for i in reversed(range(256)):  
        y = y*y  
        if 1&(e>>i):  
            y = y*x  
    return y
```

This code uses 256 squarings,  
plus 1 extra multiplication  
for each bit set in  $e$ .

Problem when  $e$  is secret: time  
leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with  
exactly 128 bits set. There are  
enough of these  $e$ , and then  
**there are no more leaks.**”

# Constant-time square-and-multiply

D. J. Bernstein

University of Illinois at Chicago;  
Ruhr University Bochum

---

```
def pow256bit(x,e):  
    y = 1  
    for i in reversed(range(256)):  
        y = y*y  
        if 1&(e>>i):  
            y = y*x  
    return y
```

This code uses 256 squarings,  
plus 1 extra multiplication  
for each bit set in  $e$ .

Problem when  $e$  is secret: time  
leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with  
exactly 128 bits set. There are  
enough of these  $e$ , and then  
**there are no more leaks.**”

— Time still depends on  $e$ ,  
even if each multiplication  
takes time independent of inputs.

t-time  
nd-multiply

ernstein

ty of Illinois at Chicago;  
niversity Bochum

---

256bit(x,e):

```
    in reversed(range(256)):
```

```
        y*y
```

```
        1&(e>>i):
```

```
            = y*x
```

```
    n y
```

1

This code uses 256 squarings,  
plus 1 extra multiplication  
for each bit set in  $e$ .

Problem when  $e$  is secret: time  
leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with  
exactly 128 bits set. There are  
enough of these  $e$ , and then  
**there are no more leaks.**”

— Time still depends on  $e$ ,  
even if each multiplication  
takes time independent of inputs.

2

Hardware  
is inheren  
CPU des

1

This code uses 256 squarings,  
plus 1 extra multiplication  
for each bit set in  $e$ .

Problem when  $e$  is secret: time  
leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with  
exactly 128 bits set. There are  
enough of these  $e$ , and then  
**there are no more leaks.**”

— Time still depends on  $e$ ,  
even if each multiplication  
takes time independent of inputs.

2

Hardware reality:  
is inherently expensive.  
CPU designers try

1

This code uses 256 squarings,  
plus 1 extra multiplication  
for each bit set in  $e$ .

Problem when  $e$  is secret: time  
leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with  
exactly 128 bits set. There are  
enough of these  $e$ , and then  
**there are no more leaks.**”

— Time still depends on  $e$ ,  
even if each multiplication  
takes time independent of inputs.

2

Hardware reality: Accessing  
is inherently expensive.

CPU designers try to reduce

This code uses 256 squarings, plus 1 extra multiplication for each bit set in  $e$ .

Problem when  $e$  is secret: time leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with exactly 128 bits set. There are enough of these  $e$ , and then **there are no more leaks.**”

— Time still depends on  $e$ , even if each multiplication takes time independent of inputs.

Hardware reality: Accessing RAM is inherently expensive.

CPU designers try to reduce cost.

This code uses 256 squarings, plus 1 extra multiplication for each bit set in  $e$ .

Problem when  $e$  is secret: time leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with exactly 128 bits set. There are enough of these  $e$ , and then **there are no more leaks.**”

— Time still depends on  $e$ , even if each multiplication takes time independent of inputs.

Hardware reality: Accessing RAM is inherently expensive.

CPU designers try to reduce cost.

Example: “L1 cache” typically has 32KB of recently used data.

This cache inspects RAM addresses, performs various computations on addresses to try to save time.



This code uses 256 squarings, plus 1 extra multiplication for each bit set in  $e$ .

Problem when  $e$  is secret: time leaks number of bits set in  $e$ .

“I’ll choose secret 256-bit  $e$  with exactly 128 bits set. There are enough of these  $e$ , and then **there are no more leaks.**”

— Time still depends on  $e$ , even if each multiplication takes time independent of inputs.

Hardware reality: Accessing RAM is inherently expensive.

CPU designers try to reduce cost.

Example: “L1 cache” typically has 32KB of recently used data.

This cache inspects RAM addresses, performs various computations on addresses to try to save time.

... so time is a function of RAM addresses. Avoid all data flow from secrets to RAM addresses.

de uses 256 squarings,  
extra multiplication  
bit set in  $e$ .

when  $e$  is secret: time  
number of bits set in  $e$ .

ose secret 256-bit  $e$  with  
128 bits set. There are  
of these  $e$ , and then  
**no more leaks.**"

still depends on  $e$ ,  
each multiplication  
ne independent of inputs.

2

Hardware reality: Accessing RAM  
is inherently expensive.

CPU designers try to reduce cost.

Example: "L1 cache" typically  
has 32KB of recently used data.

This cache inspects RAM  
addresses, performs various  
computations on addresses  
to try to save time.

... so time is a function of RAM  
addresses. Avoid all data flow  
from secrets to RAM addresses.

3

Example  
from sec  
Often de  
for softw  
the same

6 squarings,  
plication

e.

s secret: time  
its set in e.

256-bit e with  
et. There are  
, and then  
**leaks.**"

nds on e,  
plication  
ndent of inputs.

2

Hardware reality: Accessing RAM  
is inherently expensive.

CPU designers try to reduce cost.

Example: "L1 cache" typically  
has 32KB of recently used data.

This cache inspects RAM  
addresses, performs various  
computations on addresses  
to try to save time.

... so time is a function of RAM  
addresses. Avoid all data flow  
from secrets to RAM addresses.

3

Example: Avoid a  
from secrets to bra  
Often described as  
for software, but o  
the same hardware

2

Hardware reality: Accessing RAM is inherently expensive.

CPU designers try to reduce cost.

Example: "L1 cache" typically has 32KB of recently used data.

This cache inspects RAM addresses, performs various computations on addresses to try to save time.

... so time is a function of RAM addresses. Avoid all data flow from secrets to RAM addresses.

3

Example: Avoid all data flow from secrets to branch conditions.

Often described as a separate layer for software, but comes from the same hardware reality.

Hardware reality: Accessing RAM is inherently expensive.

CPU designers try to reduce cost.

Example: “L1 cache” typically has 32KB of recently used data.

This cache inspects RAM addresses, performs various computations on addresses to try to save time.

... so time is a function of RAM addresses. Avoid all data flow from secrets to RAM addresses.

Example: Avoid all data flow from secrets to branch conditions.

Often described as a separate rule for software, but comes from the same hardware reality.

Hardware reality: Accessing RAM is inherently expensive.

CPU designers try to reduce cost.

Example: “L1 cache” typically has 32KB of recently used data.

This cache inspects RAM addresses, performs various computations on addresses to try to save time.

... so time is a function of RAM addresses. Avoid all data flow from secrets to RAM addresses.

Example: Avoid all data flow from secrets to branch conditions.

Often described as a separate rule for software, but comes from the same hardware reality.

How CPU runs a program (example of “code = data”):

```
while True:
```

```
    insn = RAM[state.ip]
```

```
    state = execute(state, insn)
```

ip (“instruction pointer” or “program counter”): address in RAM of next instruction.

the reality: Accessing RAM is  
extremely expensive.

Designers try to reduce cost.

Example: "L1 cache" typically  
is a small amount of recently used data.

The cache inspects RAM  
addresses, performs various  
operations on addresses  
to save time.

Cache hit time is a function of RAM  
accesses. Avoid all data flow  
from secrets to RAM addresses.

3

Example: Avoid all data flow  
from secrets to branch conditions.  
Often described as a separate rule  
for software, but comes from  
the same hardware reality.

How CPU runs a program  
(example of "code = data"):

```
while True:  
    insn = RAM[state.ip]  
    state = execute(state, insn)
```

ip ("instruction pointer" or  
"program counter"): address  
in RAM of next instruction.

4

Standard  
to follow  
Square a

```
def pow(x, n):  
    y = 1  
    for i in range(n):  
        y = y * x  
    return y
```

If bit is  
an unuse



3

Accessing RAM  
nsive.

to reduce cost.

he” typically  
ntly used data.

ts RAM  
ns various  
addresses

e.

nction of RAM

all data flow

AM addresses.

Example: Avoid all data flow  
from secrets to branch conditions.  
Often described as a separate rule  
for software, but comes from  
the same hardware reality.

How CPU runs a program  
(example of “code = data”):

```
while True:
    insn = RAM[state.ip]
    state = execute(state, insn)
```

ip (“instruction pointer” or  
“program counter”): address  
in RAM of next instruction.

4

Standard square-a  
to follow these dat  
Square and always

```
def pow256bit(x,
    y = 1
    for i in rever
        y = y*y
        yx = y*x
        bit = 1&(e>>
        y = y+(yx-y)
    return y
```

If bit is 0 then yx  
an unused “dumm



3

RAM

cost.

lly

ata.

Example: Avoid all data flow from secrets to branch conditions.

Often described as a separate rule for software, but comes from the same hardware reality.

How CPU runs a program (example of “code = data”):

```
while True:
    insn = RAM[state.ip]
    state = execute(state, insn)
```

ip (“instruction pointer” or “program counter”): address in RAM of next instruction.

RAM

ow

ses.

4

Standard square-and-multiply to follow these data-flow rules. Square and always multiply.

```
def pow256bit(x, e):
    y = 1
    for i in reversed(range(256)):
        y = y*y
        yx = y*x
        bit = 1&(e>>i)
        y = y+(yx-y)*bit
    return y
```

If bit is 0 then yx computation is an unused “dummy operation”.

Example: Avoid all data flow from secrets to branch conditions. Often described as a separate rule for software, but comes from the same hardware reality.

How CPU runs a program (example of “code = data”):

```
while True:
    insn = RAM[state.ip]
    state = execute(state, insn)
```

ip (“instruction pointer” or “program counter”): address in RAM of next instruction.

Standard square-and-multiply fix to follow these data-flow rules: Square and always multiply.

```
def pow256bit(x,e):
    y = 1
    for i in reversed(range(256)):
        y = y*y
        yx = y*x
        bit = 1&(e>>i)
        y = y+(yx-y)*bit
    return y
```

If bit is 0 then yx computation is an unused “dummy operation”.

4  
Avoid all data flow  
crets to branch conditions.  
described as a separate rule  
ware, but comes from  
e hardware reality.

U runs a program  
e of “code = data”):

ue:  
= RAM[state.ip]  
= execute(state,insn)

truction pointer” or  
m counter”): address  
of next instruction.

4  
Standard square-and-multiply fix  
to follow these data-flow rules:  
Square and always multiply.

```
def pow256bit(x,e):  
    y = 1  
    for i in reversed(range(256)):  
        y = y*y  
        yx = y*x  
        bit = 1&(e>>i)  
        y = y+(yx-y)*bit  
    return y
```

If bit is 0 then yx computation is  
an unused “dummy operation”.

5  
Another  
def pow  
y,i,j  
while  
if  
y  
i  
e  
else  
y  
i  
return

4

All data flow  
 branch conditions.  
 as a separate rule  
 comes from  
 e reality.  
 program  
 e = data”):  
 te.ip]  
 e(state,insn)  
 ointer” or  
 ’): address  
 struction.

Standard square-and-multiply fix  
 to follow these data-flow rules:  
 Square and always multiply.

```
def pow256bit(x,e):
    y = 1
    for i in reversed(range(256)):
        y = y*y
        yx = y*x
        bit = 1&(e>>i)
        y = y+(yx-y)*bit
    return y
```

If bit is 0 then yx computation is  
 an unused “dummy operation”.

5

Another approach,  
 def pow256bit(x,  
 y,i,j = 1,255,  
 while i >= 0:  
 if j == 0:  
 y = y\*y  
 if 1&(e>>i)  
 j = 1  
 else:  
 i = i-1  
 else:  
 y = y\*x  
 i,j = i-1,  
 return y

4

Standard square-and-multiply fix  
to follow these data-flow rules:  
Square and always multiply.

```
def pow256bit(x,e):
    y = 1
    for i in reversed(range(256)):
        y = y*y
        yx = y*x
        bit = 1&(e>>i)
        y = y+(yx-y)*bit
    return y
```

If bit is 0 then yx computation is  
an unused “dummy operation”.

5

Another approach, not well

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        if j == 0:
            y = y*y
            if 1&(e>>i):
                j = 1
            else:
                i = i-1
        else:
            y = y*x
            i,j = i-1,0
    return y
```

Standard square-and-multiply fix  
to follow these data-flow rules:  
Square and always multiply.

```
def pow256bit(x,e):
    y = 1
    for i in reversed(range(256)):
        y = y*y
        yx = y*x
        bit = 1&(e>>i)
        y = y+(yx-y)*bit
    return y
```

If bit is 0 then yx computation is  
an unused “dummy operation”.

Another approach, not well known:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        if j == 0:
            y = y*y
            if 1&(e>>i):
                j = 1
            else:
                i = i-1
        else:
            y = y*x
            i,j = i-1,0
    return y
```

square-and-multiply fix  
these data-flow rules:  
and always multiply.

```
pow256bit(x,e):
```

```
    for i in reversed(range(256)):
```

```
        y = y*y
```

```
        if 1 & (e >> i):
```

```
            y = y*x
```

```
    return y
```

0 then  $yx$  computation is

called “dummy operation”.

5

Another approach, not well known:

```
def pow256bit(x,e):
```

```
    y,i,j = 1,255,0
```

```
    while i >= 0:
```

```
        if j == 0:
```

```
            y = y*y
```

```
            if 1 & (e >> i):
```

```
                j = 1
```

```
            else:
```

```
                i = i-1
```

```
        else:
```

```
            y = y*x
```

```
            i,j = i-1,0
```

```
    return y
```

6

This is the  
original s

$j$  is “inst

0 if at to

1 if in m

Each “ir

includes

5

and-multiply fix  
 data-flow rules:  
 multiply.

e):

sed(range(256)):

i)

\*bit

x computation is  
 y operation".

Another approach, not well known:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        if j == 0:
            y = y*y
            if 1&(e>>i):
                j = 1
            else:
                i = i-1
        else:
            y = y*x
            i,j = i-1,0
    return y
```

6

This is like CPU's  
 original square-and

*j* is "instruction po

0 if at top of loop

1 if in middle of lo

Each "instruction"

includes exactly on



5

y fix  
es:

(256)):

tion is  
on".

Another approach, not well known:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        if j == 0:
            y = y*y
            if 1&(e>>i):
                j = 1
            else:
                i = i-1
        else:
            y = y*x
            i,j = i-1,0
    return y
```

6

This is like CPU's perspective  
original square-and-multiply.

$j$  is "instruction pointer":

0 if at top of loop,

1 if in middle of loop.

Each "instruction" here

includes exactly one multiply

Another approach, not well known:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        if j == 0:
            y = y*y
            if 1&(e>>i):
                j = 1
            else:
                i = i-1
        else:
            y = y*x
            i,j = i-1,0
    return y
```

This is like CPU's perspective on original square-and-multiply.

$j$  is "instruction pointer":

0 if at top of loop,

1 if in middle of loop.

Each "instruction" here

includes exactly one multiply.

Another approach, not well known:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        if j == 0:
            y = y*y
            if 1&(e>>i):
                j = 1
            else:
                i = i-1
        else:
            y = y*x
            i,j = i-1,0
    return y
```

This is like CPU's perspective on original square-and-multiply.

$j$  is "instruction pointer":

0 if at top of loop,

1 if in middle of loop.

Each "instruction" here includes exactly one multiply.

Try to choose instruction set with big useful operations, avoiding control overhead.

Analogous to designing CPU.

approach, not well known:

```
256bit(x, e):
```

```
    = 1, 255, 0
```

```
    i >= 0:
```

```
        j == 0:
```

```
            = y*y
```

```
        if 1 & (e >> i):
```

```
            j = 1
```

```
        else:
```

```
            i = i-1
```

```
        e:
```

```
            = y*x
```

```
        , j = i-1, 0
```

```
    n y
```

6

This is like CPU's perspective on original square-and-multiply.

*j* is "instruction pointer":

0 if at top of loop,

1 if in middle of loop.

Each "instruction" here includes exactly one multiply.

Try to choose instruction set with big useful operations, avoiding control overhead.

Analogous to designing CPU.

7

Following

assuming

*i* shifts e

assuming

```
def pow2
```

```
    y, i, j
```

```
    while
```

```
        z =
```

```
        y =
```

```
        bit
```

```
        i =
```

```
        j =
```

```
    return
```

6

, not well known:

e):

0

):

0

This is like CPU's perspective on original square-and-multiply.

$j$  is "instruction pointer":

0 if at top of loop,

1 if in middle of loop.

Each "instruction" here includes exactly one multiply.

Try to choose instruction set with big useful operations, avoiding control overhead.

Analogous to designing CPU.

7

Following data-flow assuming all arithr  
( $i$  shifts etc.) is con  
assuming  $e$  weight

```
def pow256bit(x,
    y, i, j = 1, 255,
    while i >= 0:
        z = y+(x-y)*
        y = y*z
        bit = 1&(e>>
        i = i-(j|(1-
        j = bit&(1-j
    return y
```

6

known:

This is like CPU's perspective on original square-and-multiply.

$j$  is "instruction pointer":

0 if at top of loop,

1 if in middle of loop.

Each "instruction" here includes exactly one multiply.

Try to choose instruction set with big useful operations, avoiding control overhead.

Analogous to designing CPU.

7

Following data-flow rules, assuming all arithmetic (including  $i$  shifts etc.) is constant-time assuming  $e$  weight exactly 1

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        z = y+(x-y)*j
        y = y*z
        bit = 1&(e>>i)
        i = i-(j|(1-bit))
        j = bit&(1-j)
    return y
```

This is like CPU's perspective on original square-and-multiply.

$j$  is "instruction pointer":

0 if at top of loop,

1 if in middle of loop.

Each "instruction" here includes exactly one multiply.

Try to choose instruction set with big useful operations, avoiding control overhead.

Analogous to designing CPU.

Following data-flow rules, assuming all arithmetic (including  $i$  shifts etc.) is constant-time, assuming  $e$  weight exactly 128:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        z = y+(x-y)*j
        y = y*z
        bit = 1&(e>>i)
        i = i-(j|(1-bit))
        j = bit&(1-j)
    return y
```

like CPU's perspective on  
square-and-multiply.

instruction pointer":

top of loop,

middle of loop.

instruction" here

exactly one multiply.

choose instruction set

useful operations,

control overhead.

us to designing CPU.

7

Following data-flow rules,  
assuming all arithmetic (including  
*i* shifts etc.) is constant-time,  
assuming *e* weight exactly 128:

```
def pow256bit(x,e):  
    y,i,j = 1,255,0  
    while i >= 0:  
        z = y+(x-y)*j  
        y = y*z  
        bit = 1&(e>>i)  
        i = i-(j|(1-bit))  
        j = bit&(1-j)  
    return y
```

8

Allowing

```
def pow256bit(x,e):
```

```
    y,i,j = 1,255,0
```

```
    for l in range(e):
```

```
        z = y+(x-y)*j
```

```
        y = y*z
```

```
        bit = 1&(e>>l)
```

```
        i = l-(j|(1-bit))
```

```
        j = bit&(1-j)
```

```
    assert i == 0
```

```
    return y
```



7

Following data-flow rules,  
 assuming all arithmetic (including  
 $i$  shifts etc.) is constant-time,  
 assuming  $e$  weight exactly 128:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        z = y+(x-y)*j
        y = y*z
        bit = 1&(e>>i)
        i = i-(j|(1-bit))
        j = bit&(1-j)
    return y
```

8

Allowing any weight

```
def pow256bitwei(x,e):
    y,i,j = 1,255,0
    for loop in range(e):
        z = y+(x-y)*j
        z = z+(1-z)*j
        y = y*z
        bit = 1&(e>>loop)
        i = i-(j|(1-bit))
        j = bit&(1-j)
    assert i < 0
    return y
```

7

ve on

Following data-flow rules,  
 assuming all arithmetic (including  
 $i$  shifts etc.) is constant-time,  
 assuming  $e$  weight exactly 128:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        z = y+(x-y)*j
        y = y*z
        bit = 1&(e>>i)
        i = i-(j|(1-bit))
        j = bit&(1-j)
    return y
```

8

Allowing any weight  $\leq 128$ :

```
def pow256bitweight1e128(
    y,i,j = 1,255,0
    for loop in range(384):
        z = y+(x-y)*j
        z = z+(1-z)*(i<0)
        y = y*z
        bit = 1&(e>>max(i,0))
        i = i-(j|(1-bit))
        j = bit&(1-j)
    assert i < 0
    return y
```

Following data-flow rules,  
 assuming all arithmetic (including  
*i* shifts etc.) is constant-time,  
 assuming *e* weight exactly 128:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        z = y+(x-y)*j
        y = y*z
        bit = 1&(e>>i)
        i = i-(j|(1-bit))
        j = bit&(1-j)
    return y
```

Allowing any weight  $\leq 128$ :

```
def pow256bitweightle128(x,e):
    y,i,j = 1,255,0
    for loop in range(384):
        z = y+(x-y)*j
        z = z+(1-z)*(i<0)
        y = y*z
        bit = 1&(e>>max(i,0))
        i = i-(j|(1-bit))
        j = bit&(1-j)
    assert i < 0
    return y
```

Following data-flow rules,  
 assuming all arithmetic (including  
*i* shifts etc.) is constant-time,  
 assuming *e* weight exactly 128:

```
def pow256bit(x,e):
    y,i,j = 1,255,0
    while i >= 0:
        z = y+(x-y)*j
        y = y*z
        bit = 1&(e>>i)
        i = i-(j|(1-bit))
        j = bit&(1-j)
    return y
```

Allowing any weight  $\leq 128$ :

```
def pow256bitweightle128(x,e):
    y,i,j = 1,255,0
    for loop in range(384):
        z = y+(x-y)*j
        z = z+(1-z)*(i<0)
        y = y*z
        bit = 1&(e>>max(i,0))
        i = i-(j|(1-bit))
        j = bit&(1-j)
    assert i < 0
    return y
```

Exercise: constant-time ECC  
 scalar mult with sliding windows.