

Cryptographic software engineering, part 1

Daniel J. Bernstein

This is easy, right?

1. Take general principles of software engineering.
2. Apply principles to crypto.

Let's try some examples . . .

1972 Parnas “On the criteria to be used in decomposing systems into modules” :

“We propose instead that one begins with a list of difficult design decisions or design decisions which are likely to change. Each module is then designed to hide such a decision from the others.”

e.g. If number of cipher rounds is properly modularized as

```
#define ROUNDS 20
```

then it is easy to change.

raphic
engineering,

. Bernstein

easy, right?

general principles

software engineering.

y principles to crypto.

some examples . . .

1

1972 Parnas “On the criteria to be used in decomposing systems into modules” :

“We propose instead that one begins with a list of difficult design decisions or design decisions which are likely to change. Each module is then designed to hide such a decision from the others.”

e.g. If number of cipher rounds is properly modularized as

```
#define ROUNDS 20
```

then it is easy to change.

2

Another
of software
Make th
and the

1

1972 Parnas “On the criteria to be used in decomposing systems into modules” :

“We propose instead that one begins with a list of difficult design decisions or design decisions which are likely to change. Each module is then designed to hide such a decision from the others.”

e.g. If number of cipher rounds is properly modularized as

```
#define ROUNDS 20
```

then it is easy to change.

2

Another general principle of software engineering: Make the right thing easy and the wrong thing hard.

1

1972 Parnas “On the criteria to be used in decomposing systems into modules” :

“We propose instead that one begins with a list of difficult design decisions or design decisions which are likely to change. Each module is then designed to hide such a decision from the others.”

e.g. If number of cipher rounds is properly modularized as

```
#define ROUNDS 20
```

then it is easy to change.

2

Another general principle of software engineering:
Make the right thing simple
and the wrong thing complex

1972 Parnas “On the criteria to be used in decomposing systems into modules” :

“We propose instead that one begins with a list of difficult design decisions or design decisions which are likely to change. Each module is then designed to hide such a decision from the others.”

e.g. If number of cipher rounds is properly modularized as

```
#define ROUNDS 20
```

then it is easy to change.

Another general principle of software engineering:
Make the right thing simple and the wrong thing complex.

1972 Parnas “On the criteria to be used in decomposing systems into modules” :

“We propose instead that one begins with a list of difficult design decisions or design decisions which are likely to change. Each module is then designed to hide such a decision from the others.”

e.g. If number of cipher rounds is properly modularized as

```
#define ROUNDS 20
```

then it is easy to change.

Another general principle of software engineering:
Make the right thing simple and the wrong thing complex.

e.g. Make it difficult to ignore invalid authenticators.

1972 Parnas “On the criteria to be used in decomposing systems into modules” :

“We propose instead that one begins with a list of difficult design decisions or design decisions which are likely to change. Each module is then designed to hide such a decision from the others.”

e.g. If number of cipher rounds is properly modularized as

```
#define ROUNDS 20
```

then it is easy to change.

Another general principle of software engineering:
Make the right thing simple and the wrong thing complex.

e.g. Make it difficult to ignore invalid authenticators.

Do not design APIs like this:

“The sample code used in this manual omits the checking of status values for clarity, but when using cryptlib you should check return values, particularly for critical functions”

rnas “On the criteria
ed in decomposing
into modules”:
pose instead that
ins with a list of
design decisions or
ecisions which are
change. Each module
designed to hide such
on from the others.”
umber of cipher rounds
rly modularized as
ROUNDS 20
s easy to change.

2

Another general principle
of software engineering:
Make the right thing simple
and the wrong thing complex.
e.g. Make it difficult to
ignore invalid authenticators.

Do not design APIs like this:
“The sample code used in
this manual omits the checking
of status values for clarity, but
when using cryptlib you should
check return values, particularly
for critical functions ...”

3

Not so e
1970s: -
compare
against s
one char
stopping

- AAAAA
- FAAAA
- FRAAA

2

the criteria
composing
rules” :

read that
list of

decisions or
which are

Each module
to hide such
the others.”

cipher rounds
rized as
0
change.

Another general principle
of software engineering:
Make the right thing simple
and the wrong thing complex.

e.g. Make it difficult to
ignore invalid authenticators.

Do not design APIs like this:
“The sample code used in
this manual omits the checking
of status values for clarity, but
when using cryptlib you should
check return values, particularly
for critical functions”

3

Not so easy: Timing

1970s: TENEX op
compares user-sup
against secret pass
one character at a
stopping at first d

- AAAAAA vs. FRIE
- FAAAAA vs. FRIE
- FRAAAA vs. FRIE

2

Another general principle of software engineering:
Make the right thing simple and the wrong thing complex.

e.g. Make it difficult to ignore invalid authenticators.

Do not design APIs like this:

“The sample code used in this manual omits the checking of status values for clarity, but when using cryptlib you should check return values, particularly for critical functions ...”

3

Not so easy: Timing attacks

1970s: TENEX operating system compares user-supplied string against secret password one character at a time, stopping at first difference:

- AAAAAA vs. FRIEND: stop at 6th character
- FAAAAA vs. FRIEND: stop at 1st character
- FRAAAA vs. FRIEND: stop at 2nd character

Another general principle of software engineering:
Make the right thing simple and the wrong thing complex.
e.g. Make it difficult to ignore invalid authenticators.
Do not design APIs like this:
“The sample code used in this manual omits the checking of status values for clarity, but when using cryptlib you should check return values, particularly for critical functions”

Not so easy: Timing attacks

1970s: TENEX operating system compares user-supplied string against secret password one character at a time, stopping at first difference:

- AAAAAA vs. FRIEND: stop at 1.
- FAAAAA vs. FRIEND: stop at 2.
- FRAAAA vs. FRIEND: stop at 3.

Another general principle of software engineering:
Make the right thing simple and the wrong thing complex.
e.g. Make it difficult to ignore invalid authenticators.
Do not design APIs like this:
“The sample code used in this manual omits the checking of status values for clarity, but when using cryptlib you should check return values, particularly for critical functions”

Not so easy: Timing attacks

1970s: TENEX operating system compares user-supplied string against secret password one character at a time, stopping at first difference:

- AAAAAA vs. FRIEND: stop at 1.
- FAAAAA vs. FRIEND: stop at 2.
- FRAAAA vs. FRIEND: stop at 3.

Attacker sees comparison time, deduces position of difference.
A few hundred tries reveal secret password.

general principle
are engineering:
the right thing simple
wrong thing complex.

make it difficult to
invalid authenticators.

design APIs like this:

simple code used in
manual omits the checking
values for clarity, but
using cryptlib you should
return values, particularly
cal functions ...”

3

Not so easy: Timing attacks

1970s: TENEX operating system
compares user-supplied string
against secret password
one character at a time,
stopping at first difference:

- AAAAAA vs. FRIEND: stop at 1.
- FAAAAA vs. FRIEND: stop at 2.
- FRAAAA vs. FRIEND: stop at 3.

Attacker sees comparison time,
deduces position of difference.

A few hundred tries
reveal secret password.

4

How typ
16-byte
for (
if
return

3

principle
ering:
ng simple
ng complex.
ult to
enticators.
ls like this:
used in
the checking
r clarity, but
b you should
es, particularly
ns ...”

Not so easy: Timing attacks

1970s: TENEX operating system
compares user-supplied string
against secret password
one character at a time,
stopping at first difference:

- AAAAAA vs. FRIEND: stop at 1.
- FAAAAA vs. FRIEND: stop at 2.
- FRAAAA vs. FRIEND: stop at 3.

Attacker sees comparison time,
deduces position of difference.
A few hundred tries
reveal secret password.

4

How typical software
16-byte authentication

```
for (i = 0; i < 16; i++)  
    if (x[i] != y[i])  
        return 1;
```

Not so easy: Timing attacks

1970s: TENEX operating system compares user-supplied string against secret password one character at a time, stopping at first difference:

- AAAAAA vs. FRIEND: stop at 1.
- FAAAAA vs. FRIEND: stop at 2.
- FRAAAA vs. FRIEND: stop at 3.

Attacker sees comparison time, deduces position of difference.

A few hundred tries reveal secret password.

How typical software checks 16-byte authenticator:

```
for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;
```

Not so easy: Timing attacks

1970s: TENEX operating system compares user-supplied string against secret password one character at a time, stopping at first difference:

- AAAAAA vs. FRIEND: stop at 1.
- FAAAAA vs. FRIEND: stop at 2.
- FRAAAA vs. FRIEND: stop at 3.

Attacker sees comparison time, deduces position of difference.

A few hundred tries reveal secret password.

How typical software checks

16-byte authenticator:

```
for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;
```


Not so easy: Timing attacks

1970s: TENEX operating system compares user-supplied string against secret password one character at a time, stopping at first difference:

- AAAAAA vs. FRIEND: stop at 1.
- FAAAAA vs. FRIEND: stop at 2.
- FRAAAA vs. FRIEND: stop at 3.

Attacker sees comparison time, deduces position of difference.

A few hundred tries reveal secret password.

How typical software checks

16-byte authenticator:

```
for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;
```

Fix, eliminating information flow from secrets to timings:

```
diff = 0;
for (i = 0; i < 16; ++i)
    diff |= x[i] ^ y[i];
return 1 & ((diff-1) >> 8);
```

Notice that the language makes the wrong thing simple and the right thing complex.

easy: Timing attacks

TENEX operating system

as user-supplied string

secret password

character at a time,

stopping at first difference:

A vs. FRIEND: stop at 1.

A vs. FRIEND: stop at 2.

A vs. FRIEND: stop at 3.

Attacker sees comparison time,

and deduces position of difference.

After a hundred tries

attacker knows secret password.

How typical software checks

16-byte authenticator:

```
for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;
```

Fix, eliminating information flow
from secrets to timings:

```
diff = 0;
for (i = 0; i < 16; ++i)
    diff |= x[i] ^ y[i];
return 1 & ((diff-1) >> 8);
```

Notice that the language
makes the wrong thing simple
and the right thing complex.

Language

“right” is

So mistake

4

ng attacks

operating system

plied string

sword

time,

ifference:

END: stop at 1.

END: stop at 2.

END: stop at 3.

parison time,

of difference.

es

word.

How typical software checks

16-byte authenticator:

```

for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;

```

Fix, eliminating information flow
from secrets to timings:

```

diff = 0;
for (i = 0; i < 16; ++i)
    diff |= x[i] ^ y[i];
return 1 & ((diff-1) >> 8);

```

Notice that the language
makes the wrong thing simple
and the right thing complex.

5

Language designers

“right” is too weak

So mistakes continue

4

How typical software checks
16-byte authenticator:

```
for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;
```

Fix, eliminating information flow
from secrets to timings:

```
diff = 0;
for (i = 0; i < 16; ++i)
    diff |= x[i] ^ y[i];
return 1 & ((diff-1) >> 8);
```

Notice that the language
makes the wrong thing simple
and the right thing complex.

5

Language designer's notion of
"right" is too weak for security.

So mistakes continue to happen.

How typical software checks
16-byte authenticator:

```
for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;
```

Fix, eliminating information flow
from secrets to timings:

```
diff = 0;
for (i = 0; i < 16; ++i)
    diff |= x[i] ^ y[i];
return 1 & ((diff-1) >> 8);
```

Notice that the language
makes the wrong thing simple
and the right thing complex.

Language designer's notion of
"right" is too weak for security.
So mistakes continue to happen.

How typical software checks
16-byte authenticator:

```
for (i = 0; i < 16; ++i)
    if (x[i] != y[i]) return 0;
return 1;
```

Fix, eliminating information flow
from secrets to timings:

```
diff = 0;
for (i = 0; i < 16; ++i)
    diff |= x[i] ^ y[i];
return 1 & ((diff-1) >> 8);
```

Notice that the language
makes the wrong thing simple
and the right thing complex.

Language designer's notion of
"right" is too weak for security.

So mistakes continue to happen.

One of many examples,
part of the reference software for
CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

ical software checks

authenticator:

```
i = 0; i < 16; ++i)
(x[i] != y[i]) return 0;
n 1;
```

minating information flow

crets to timings:

```
= 0;
i = 0; i < 16; ++i)
f |= x[i] ^ y[i];
n 1 & ((diff-1) >> 8);
```

hat the language

ne wrong thing simple

right thing complex.

5

Language designer's notion of
"right" is too weak for security.

So mistakes continue to happen.

One of many examples,
part of the reference software for
CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

6

Do timing

Objection

5

are checks

ator:

```
16; ++i)
y[i]) return 0;
```

formation flow

nings:

```
16; ++i)
^ y[i];
iff-1) >> 8);
```

nguage

thing simple

g complex.

Language designer's notion of
"right" is too weak for security.

So mistakes continue to happen.

One of many examples,
part of the reference software for
CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

6

Do timing attacks

Objection: "Timin

5

Language designer's notion of "right" is too weak for security.

So mistakes continue to happen.

One of many examples, part of the reference software for CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

6

Do timing attacks really work?

Objection: "Timings are not

Language designer's notion of "right" is too weak for security.

So mistakes continue to happen.

One of many examples, part of the reference software for CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

Do timing attacks really work?

Objection: "Timings are noisy!"

Language designer's notion of "right" is too weak for security.

So mistakes continue to happen.

One of many examples, part of the reference software for CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

Do timing attacks really work?

Objection: "Timings are noisy!"

Answer #1:

Does noise stop *all* attacks?

To guarantee security, defender must block *all* information flow.

Language designer's notion of "right" is too weak for security.

So mistakes continue to happen.

One of many examples, part of the reference software for CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

Do timing attacks really work?

Objection: "Timings are noisy!"

Answer #1:

Does noise stop *all* attacks?

To guarantee security, defender must block *all* information flow.

Answer #2: Attacker uses statistics to eliminate noise.

Language designer's notion of "right" is too weak for security.

So mistakes continue to happen.

One of many examples, part of the reference software for CAESAR candidate CLOC:

```
/* compare the tag */
int i;
for(i = 0; i < CRYPTO_ABYTES; i++)
    if(tag[i] != c[(*mlen) + i]){
        return RETURN_TAG_NO_MATCH;
    }
return RETURN_SUCCESS;
```

Do timing attacks really work?

Objection: "Timings are noisy!"

Answer #1:

Does noise stop *all* attacks?

To guarantee security, defender must block *all* information flow.

Answer #2: Attacker uses statistics to eliminate noise.

Answer #3, what the 1970s attackers actually did: Cross page boundary, inducing page faults, to amplify timing signal.

the designer's notion of
is too weak for security.

akes continue to happen.

many examples,

the reference software for
R candidate CLOC:

```
are the tag */
```

```
0; i < CRYPTO_ABYTES; i++)  
g[i] != c[(*mlen) + i]){  
return RETURN_TAG_NO_MATCH;  
  
RETURN_SUCCESS;
```

6

Do timing attacks really work?

Objection: “Timings are noisy!”

Answer #1:

Does noise stop *all* attacks?

To guarantee security, defender
must block *all* information flow.

Answer #2: Attacker uses
statistics to eliminate noise.

Answer #3, what the
1970s attackers actually did:
Cross page boundary,
inducing page faults,
to amplify timing signal.

7

Defende

Some of

[1996 Ko](#)

attacks o

Briefly m

Kocher a

Schneier

secret an

affect tim

[2002 Pa](#)

Suzaki–S

timing a

6

r's notion of
k for security.
ue to happen.
mples,
ce software for
e CLOC:

```
ag */  
  
YPTO_ABYTES; i++)  
[(*mlen) + i]){  
N_TAG_NO_MATCH;  
  
CESS;
```

Do timing attacks really work?

Objection: “**Timings are noisy!**”

Answer #1:

Does noise stop *all* attacks?

To guarantee security, defender must block *all* information flow.

Answer #2: Attacker uses statistics to eliminate noise.

Answer #3, what the 1970s attackers actually did:
Cross page boundary,
inducing page faults,
to amplify timing signal.

7

Defenders don't le

Some of the litera
[1996](#) Kocher point
attacks on cryptog
Briefly mentioned
Kocher and by [199](#)
Schneier–Wagner–
secret array indice
affect timing via c
[2002](#) Page, 2003 T
Suzaki–Shigeri–Mi
timing attacks on

6

Do timing attacks really work?

Objection: “Timings are noisy!”

Answer #1:

Does noise stop *all* attacks?

To guarantee security, defender must block *all* information flow.

Answer #2: Attacker uses statistics to eliminate noise.

Answer #3, what the 1970s attackers actually did:
Cross page boundary, inducing page faults, to amplify timing signal.

7

Defenders don't learn

Some of the literature:

[1996](#) Kocher pointed out timing attacks on cryptographic keys

Briefly mentioned by Kocher and by [1998](#) Kelsey–

Schneier–Wagner–Hall: secret array indices can

affect timing via cache misses

[2002](#) Page, 2003 Tsunoo–Sa

Suzaki–Shigeri–Miyachi:

timing attacks on DES.

Do timing attacks really work?

Objection: “**Timings are noisy!**”

Answer #1:

Does noise stop *all* attacks?

To guarantee security, defender must block *all* information flow.

Answer #2: Attacker uses statistics to eliminate noise.

Answer #3, what the 1970s attackers actually did:
Cross page boundary,
inducing page faults,
to amplify timing signal.

Defenders don't learn

Some of the literature:

1996 Kocher pointed out timing attacks on cryptographic key bits.

Briefly mentioned by Kocher and by **1998** Kelsey–Schneier–Wagner–Hall:
secret array indices can affect timing via cache misses.

2002 Page, 2003 Tsunoo–Saito–Suzaki–Shigeri–Miyachi:
timing attacks on DES.

Timing attacks really work?

Conclusion: “Timings are noisy!”

#1:

Can we stop *all* attacks?

To guarantee security, defender must block *all* information flow.

#2: Attacker uses side channels to eliminate noise.

#3, what the

attackers actually did:

Cache boundary,

Cache page faults,

Control flow timing signal.

Defenders don't learn

Some of the literature:

[1996](#) Kocher pointed out timing attacks on cryptographic key bits.

Briefly mentioned by Kocher and by [1998](#) Kelsey–Schneier–Wagner–Hall: secret array indices can affect timing via cache misses.

[2002](#) Page, [2003](#) Tsunoo–Saito–Suzaki–Shigeri–Miyachi: timing attacks on DES.

“Guaran
load ent

7

really work?

ings are noisy!”

// attacks?

rity, defender
ormation flow.

cker uses
ate noise.

the
ctually did:

ary,
ts,
signal.

Defenders don't learn

Some of the literature:

1996 Kocher pointed out timing attacks on cryptographic key bits.

Briefly mentioned by Kocher and by 1998 Kelsey–Schneier–Wagner–Hall: secret array indices can affect timing via cache misses.

2002 Page, 2003 Tsunoo–Saito–Suzaki–Shigeri–Miyachi: timing attacks on DES.

8

“Guaranteed” cou
load entire table in

7

Defenders don't learn

Some of the literature:

1996 Kocher pointed out timing attacks on cryptographic key bits.

Briefly mentioned by Kocher and by **1998** Kelsey–Schneier–Wagner–Hall: secret array indices can affect timing via cache misses.

2002 Page, 2003 Tsunoo–Saito–Suzaki–Shigeri–Miyachi: timing attacks on DES.

8

“Guaranteed” countermeasures: load entire table into cache.

Defenders don't learn

Some of the literature:

1996 Kocher pointed out timing attacks on cryptographic key bits.

Briefly mentioned by Kocher and by **1998** Kelsey–Schneier–Wagner–Hall: secret array indices can affect timing via cache misses.

2002 Page, 2003 Tsunoo–Saito–Suzaki–Shigeri–Miyachi: timing attacks on DES.

“Guaranteed” countermeasure: load entire table into cache.

Defenders don't learn

Some of the literature:

1996 Kocher pointed out timing attacks on cryptographic key bits.

Briefly mentioned by Kocher and by **1998** Kelsey–Schneier–Wagner–Hall: secret array indices can affect timing via cache misses.

2002 Page, 2003 Tsunoo–Saito–Suzaki–Shigeri–Miyachi: timing attacks on DES.

“Guaranteed” countermeasure: load entire table into cache.

2004.11/2005.04 Bernstein:

Timing attacks on AES.

Countermeasure isn't safe;

e.g., secret array indices can affect timing via cache-bank collisions.

What *is* safe: kill all data flow from secrets to array indices.

Defenders don't learn

Some of the literature:

1996 Kocher pointed out timing attacks on cryptographic key bits.

Briefly mentioned by Kocher and by **1998** Kelsey–Schneier–Wagner–Hall: secret array indices can affect timing via cache misses.

2002 Page, 2003 Tsunoo–Saito–Suzaki–Shigeri–Miyachi: timing attacks on DES.

“Guaranteed” countermeasure: load entire table into cache.

2004.11/2005.04 Bernstein:

Timing attacks on AES.

Countermeasure isn't safe;

e.g., secret array indices can affect timing via cache-bank collisions.

What *is* safe: kill all data flow from secrets to array indices.

2005 Tromer–Osvik–Shamir:

65ms to steal Linux AES key used for hard-disk encryption.

rs don't learn

the literature:

cher pointed out timing
on cryptographic key bits.

mentioned by

and by [1998](#) Kelsey–

–Wagner–Hall:

rray indices can

ming via cache misses.

ge, 2003 Tsunoo–Saito–

Shigeri–Miyachi:

ttacks on DES.

8

“Guaranteed” countermeasure:
load entire table into cache.

[2004.11/2005.04 Bernstein](#):

Timing attacks on AES.

Countermeasure isn't safe;

e.g., secret array indices can affect
timing via cache-bank collisions.

What *is* safe: kill all data flow
from secrets to array indices.

[2005](#) Tromer–Osvik–Shamir:

65ms to steal Linux AES key
used for hard-disk encryption.

9

Intel rec

OpenSS

countern

from kno

arn

ture:

ted out timing
graphic key bits.

by

98 Kelsey–

-Hall:

s can

ache misses.

Tsunoo–Saito–

yauchi:

DES.

“Guaranteed” countermeasure:
load entire table into cache.

2004.11/2005.04 Bernstein:

Timing attacks on AES.

Countermeasure isn’t safe;

e.g., secret array indices can affect
timing via cache-bank collisions.

What *is* safe: kill all data flow
from secrets to array indices.

2005 Tromer–Osvik–Shamir:

65ms to steal Linux AES key
used for hard-disk encryption.

Intel recommends,
OpenSSL integrated
countermeasure: a
from known *lines*

8

“Guaranteed” countermeasure:
load entire table into cache.

[2004.11/2005.04 Bernstein:](#)

Timing attacks on AES.

Countermeasure isn't safe;

e.g., secret array indices can affect
timing via cache-bank collisions.

What *is* safe: kill all data flow
from secrets to array indices.

[2005 Tromer–Osvik–Shamir:](#)

65ms to steal Linux AES key
used for hard-disk encryption.

9

Intel recommends, and
OpenSSL integrates, cheaper
countermeasure: always load
from known *lines* of cache.

“Guaranteed” countermeasure:
load entire table into cache.

2004.11/2005.04 Bernstein:

Timing attacks on AES.

Countermeasure isn't safe;
e.g., secret array indices can affect
timing via cache-bank collisions.

What *is* safe: kill all data flow
from secrets to array indices.

2005 Tromer–Osvik–Shamir:
65ms to steal Linux AES key
used for hard-disk encryption.

Intel recommends, and
OpenSSL integrates, cheaper
countermeasure: always loading
from known *lines* of cache.

“Guaranteed” countermeasure:
load entire table into cache.

2004.11/2005.04 Bernstein:

Timing attacks on AES.

Countermeasure isn't safe;
e.g., secret array indices can affect
timing via cache-bank collisions.
What *is* safe: kill all data flow
from secrets to array indices.

2005 Tromer–Osvik–Shamir:
65ms to steal Linux AES key
used for hard-disk encryption.

Intel recommends, and
OpenSSL integrates, cheaper
countermeasure: always loading
from known *lines* of cache.

2013 Bernstein–Schwabe

“A word of warning” :
This countermeasure isn't safe.
Variable-time lab experiment.
Same issues described in 2004.

“Guaranteed” countermeasure:
load entire table into cache.

2004.11/2005.04 Bernstein:

Timing attacks on AES.

Countermeasure isn't safe;
e.g., secret array indices can affect
timing via cache-bank collisions.
What *is* safe: kill all data flow
from secrets to array indices.

2005 Tromer–Osvik–Shamir:
65ms to steal Linux AES key
used for hard-disk encryption.

Intel recommends, and
OpenSSL integrates, cheaper
countermeasure: always loading
from known *lines* of cache.

2013 Bernstein–Schwabe

“A word of warning” :
This countermeasure isn't safe.
Variable-time lab experiment.
Same issues described in 2004.

2016 Yarom–Genkin–Hening
“CacheBleed” steals RSA secret
key via timings of OpenSSL.

“steal” countermeasure:

write table into cache.

/2005.04 Bernstein:

attacks on AES.

countermeasure isn't safe;

secret array indices can affect

via cache-bank collisions.

not safe: kill all data flow

depends on array indices.

Barber–Osvik–Shamir:

steal Linux AES key

via hard-disk encryption.

Intel recommends, and

OpenSSL integrates, cheaper

countermeasure: always loading

from known *lines* of cache.

2013 Bernstein–Schwabe

“A word of warning”:

This countermeasure isn't safe.

Variable-time lab experiment.

Same issues described in 2004.

2016 Yarom–Genkin–Heninger

“CacheBleed” steals RSA secret

key via timings of OpenSSL.

2008 RF

Layer Se

Version

small tim

performa

extent o

fragment

be large

due to t

existing

of the ti

countermeasure:
into cache.

Bernstein:
AES.

isn't safe;
indices can affect
bank collisions.
all data flow
array indices.

Wagner–Shamir:
128-bit AES key
encryption.

Intel recommends, and
OpenSSL integrates, cheaper
countermeasure: always loading
from known *lines* of cache.

2013 Bernstein–Schwabe

“A word of warning”:
This countermeasure isn't safe.
Variable-time lab experiment.
Same issues described in 2004.

2016 Yarom–Genkin–Heninger

“CacheBleed” steals RSA secret
key via timings of OpenSSL.

2008 RFC 5246 “
Layer Security (TL
Version 1.2” : “Th
small timing chan
performance deper
extent on the size
fragment, but it is
be large enough to
due to the large bl
existing MACs and
of the timing signa

re:

Intel recommends, and
OpenSSL integrates, cheaper
countermeasure: always loading
from known *lines* of cache.

2013 Bernstein–Schwabe

“A word of warning”:

This countermeasure isn’t safe.

Variable-time lab experiment.

Same issues described in 2004.

2016 Yarom–Genkin–Heninger

“CacheBleed” steals RSA secret
key via timings of OpenSSL.

2008 RFC 5246 “The Transport
Layer Security (TLS) Protocol
Version 1.2”: “This leaves a
small timing channel, since
performance depends to some
extent on the size of the data
fragment, but it is **not believed
to be large enough to be exploited**
due to the large block size of
existing MACs and the small
of the timing signal.”

affect

ons.

ow

.

:

y

n.

Intel recommends, and OpenSSL integrates, cheaper countermeasure: always loading from known *lines* of cache.

2013 Bernstein–Schwabe

“A word of warning”:

This countermeasure isn’t safe.

Variable-time lab experiment.

Same issues described in 2004.

2016 Yarom–Genkin–Hening

“CacheBleed” steals RSA secret key via timings of OpenSSL.

2008 RFC 5246 “The Transport Layer Security (TLS) Protocol, Version 1.2”: “This leaves a small timing channel, since MAC performance depends to some extent on the size of the data fragment, but it is **not believed to be large enough to be exploitable**, due to the large block size of existing MACs and the small size of the timing signal.”

Intel recommends, and OpenSSL integrates, cheaper countermeasure: always loading from known *lines* of cache.

2013 Bernstein–Schwabe

“A word of warning”:

This countermeasure isn’t safe.

Variable-time lab experiment.

Same issues described in 2004.

2016 Yarom–Genkin–Hening

“CacheBleed” steals RSA secret key via timings of OpenSSL.

2008 RFC 5246 “The Transport Layer Security (TLS) Protocol, Version 1.2”: “This leaves a small timing channel, since MAC performance depends to some extent on the size of the data fragment, but it is **not believed to be large enough to be exploitable**, due to the large block size of existing MACs and the small size of the timing signal.”

2013 AlFardan–Paterson “Lucky Thirteen: breaking the TLS and DTLS record protocols”: exploit these timings; steal plaintext.

ommends, and
 L integrates, cheaper
 measure: always loading
 own *lines* of cache.

rnstein–Schwabe
 of warning”:

intermeasure isn’t safe.
 -time lab experiment.
 sues described in 2004.

rom–Genkin–Heninger
 “Bleed” steals RSA secret
 timings of OpenSSL.

2008 RFC 5246 “The Transport
 Layer Security (TLS) Protocol,
 Version 1.2”: “This leaves a
 small timing channel, since MAC
 performance depends to some
 extent on the size of the data
 fragment, but it is **not believed to
 be large enough to be exploitable**,
 due to the large block size of
 existing MACs and the small size
 of the timing signal.”

2013 AlFardan–Paterson “Lucky
 Thirteen: breaking the TLS and
 DTLS record protocols”: exploit
 these timings; steal plaintext.

How to

If possible
 to control

Look for
 identifying

“Division
 when the
 complete
 cycles re
 values o

Measure
 trusting

and
 es, cheaper
 always loading
 of cache.

chwabe
 g”:
 ure isn't safe.
 experiment.
 bed in 2004.

in–Heninger
 als RSA secret
 OpenSSL.

[2008 RFC 5246](#) “The Transport Layer Security (TLS) Protocol, Version 1.2”: “This leaves a small timing channel, since MAC performance depends to some extent on the size of the data fragment, but it is **not believed to be large enough to be exploitable**, due to the large block size of existing MACs and the small size of the timing signal.”

[2013 AlFardan–Paterson](#) “Lucky Thirteen: breaking the TLS and DTLS record protocols”: exploit these timings; steal plaintext.

How to write cons
 If possible, write c
 to control instruct
 Look for documen
 identifying variabil
 “Division operatio
 when the divide op
 completes, with th
 cycles required dep
 values of the input
 Measure cycles rat
 trusting CPU docu

[2008 RFC 5246](#) “The Transport Layer Security (TLS) Protocol, Version 1.2”: “This leaves a small timing channel, since MAC performance depends to some extent on the size of the data fragment, but it is **not believed to be large enough to be exploitable**, due to the large block size of existing MACs and the small size of the timing signal.”

[2013 AlFardan–Paterson](#) “Lucky Thirteen: breaking the TLS and DTLS record protocols”: exploit these timings; steal plaintext.

How to write constant-time

If possible, write code in assembly to control instruction selection

Look for documentation

identifying variability: e.g.,

“Division operations terminate when the divide operation completes, with the number of cycles required dependent on values of the input operands”

Measure cycles rather than trusting CPU documentation

[2008 RFC 5246](#) “The Transport Layer Security (TLS) Protocol, Version 1.2”: “This leaves a small timing channel, since MAC performance depends to some extent on the size of the data fragment, but it is **not believed to be large enough to be exploitable**, due to the large block size of existing MACs and the small size of the timing signal.”

[2013 AlFardan–Paterson](#) “Lucky Thirteen: breaking the TLS and DTLS record protocols”: exploit these timings; steal plaintext.

How to write constant-time code

If possible, write code in asm to control instruction selection.

Look for documentation identifying variability: e.g., “Division operations terminate when the divide operation completes, with the number of cycles required dependent on the values of the input operands.”

Measure cycles rather than trusting CPU documentation.

FC 5246 “The Transport Security (TLS) Protocol, 1.2”: “This leaves a timing channel, since MAC performance depends to some extent on the size of the data being MACed, but it is **not believed to be exploitable**, due to the large block size of MACs and the small size of the timing signal.”

Fardan–Paterson “Lucky breaks: breaking the TLS and record protocols”: exploit timings; steal plaintext.

How to write constant-time code

If possible, write code in asm to control instruction selection.

Look for documentation

identifying variability: e.g.,

“Division operations terminate when the divide operation completes, with the number of cycles required dependent on the values of the input operands.”

Measure cycles rather than trusting CPU documentation.

Cut off a
secrets t

Cut off a
secrets t

Cut off a
secrets t

Prefer lo

Prefer ve

Watch o

variable-

[Cortex-M](#)

The Transport
(S) Protocol,
is leaves a
nel, since MAC
nds to some
of the data
not believed to
be exploitable,
lock size of
d the small size
al.”

nterson “Lucky
g the TLS and
ocols”: exploit
al plaintext.

How to write constant-time code

If possible, write code in asm
to control instruction selection.

Look for documentation

identifying variability: e.g.,

“Division operations terminate
when the divide operation
completes, with the number of
cycles required dependent on the
values of the input operands.”

Measure cycles rather than
trusting CPU documentation.

Cut off all data flow
secrets to branch c

Cut off all data flow
secrets to array ind

Cut off all data flow
secrets to shift/rot

Prefer logic instruc

Prefer vector instr

Watch out for CP

variable-time mult

[Cortex-M3](#) and mo

How to write constant-time code

If possible, write code in asm to control instruction selection.

Look for documentation identifying variability: e.g.,
“Division operations terminate when the divide operation completes, with the number of cycles required dependent on the values of the input operands.”

Measure cycles rather than trusting CPU documentation.

Cut off all data flow from secrets to branch conditions

Cut off all data flow from secrets to array indices.

Cut off all data flow from secrets to shift/rotate distances

Prefer logic instructions.

Prefer vector instructions.

Watch out for CPUs with variable-time multipliers: e.g.

[Cortex-M3](#) and most PowerPC

How to write constant-time code

If possible, write code in asm to control instruction selection.

Look for documentation identifying variability: e.g., “Division operations terminate when the divide operation completes, with the number of cycles required dependent on the values of the input operands.”

Measure cycles rather than trusting CPU documentation.

Cut off all data flow from secrets to branch conditions.

Cut off all data flow from secrets to array indices.

Cut off all data flow from secrets to shift/rotate distances.

Prefer logic instructions.

Prefer vector instructions.

Watch out for CPUs with variable-time multipliers: e.g., [Cortex-M3](#) and most PowerPCs.

write constant-time code

le, write code in asm
 ol instruction selection.

r documentation

ng variability: e.g.,
 n operations terminate
 e divide operation
 es, with the number of
 equired dependent on the
 f the input operands.”

e cycles rather than
 CPU documentation.

Cut off all data flow from
 secrets to branch conditions.

Cut off all data flow from
 secrets to array indices.

Cut off all data flow from
 secrets to shift/rotate distances.

Prefer logic instructions.

Prefer vector instructions.

Watch out for CPUs with
 variable-time multipliers: e.g.,
[Cortex-M3](#) and most PowerPCs.

Suppose
 const-tim

Suppose
 has “sec

Easy for
 that sec
 by const

Proofs o
 (uninitia
 ctgrind,

constant-time code
 code in asm
 branch selection.
 rotation
 safety: e.g.,
 branches terminate
 operation
 the number of
 dependent on the
 of operands.”
 rather than
 documentation.

Cut off all data flow from
 secrets to branch conditions.

Cut off all data flow from
 secrets to array indices.

Cut off all data flow from
 secrets to shift/rotate distances.

Prefer logic instructions.

Prefer vector instructions.

Watch out for CPUs with
 variable-time multipliers: e.g.,
[Cortex-M3](#) and most PowerPCs.

Suppose we know
 const-time machine

Suppose program
 has “secret” type

Easy for compiler
 that secret types
 by const-time inst

Proofs of concept:
 (uninitialized data
 ctgrind, ct-verif, F

code

n

on.

ate

of

n the

s.”

n.

Cut off all data flow from secrets to branch conditions.

Cut off all data flow from secrets to array indices.

Cut off all data flow from secrets to shift/rotate distances.

Prefer logic instructions.

Prefer vector instructions.

Watch out for CPUs with variable-time multipliers: e.g., [Cortex-M3](#) and most PowerPCs.

Suppose we know (some) const-time machine instructions.

Suppose programming language has “secret” types.

Easy for compiler to guarantee that secret types are used by const-time instructions.

Proofs of concept: Valgrind (uninitialized data as secret), ctgrind, ct-verif, FlowTracke

Cut off all data flow from secrets to branch conditions.

Cut off all data flow from secrets to array indices.

Cut off all data flow from secrets to shift/rotate distances.

Prefer logic instructions.

Prefer vector instructions.

Watch out for CPUs with variable-time multipliers: e.g., [Cortex-M3](#) and most PowerPCs.

Suppose we know (some) const-time machine instructions.

Suppose programming language has “secret” types.

Easy for compiler to guarantee that secret types are used only by const-time instructions.

Proofs of concept: Valgrind (uninitialized data as secret), ctgrind, ct-verif, FlowTracker.

Cut off all data flow from secrets to branch conditions.

Cut off all data flow from secrets to array indices.

Cut off all data flow from secrets to shift/rotate distances.

Prefer logic instructions.

Prefer vector instructions.

Watch out for CPUs with variable-time multipliers: e.g., [Cortex-M3](#) and most PowerPCs.

Suppose we know (some) const-time machine instructions.

Suppose programming language has “secret” types.

Easy for compiler to guarantee that secret types are used only by const-time instructions.

Proofs of concept: Valgrind (uninitialized data as secret), ctgrind, ct-verif, FlowTracker.

How can we implement, e.g., sorting of a secret array?

all data flow from
to branch conditions.

all data flow from
to array indices.

all data flow from
to shift/rotate distances.

logic instructions.

vector instructions.

out for CPUs with

time multipliers: e.g.,

M3 and most PowerPCs.

Suppose we know (some)
const-time machine instructions.

Suppose programming language
has “secret” types.

Easy for compiler to guarantee
that secret types are used only
by const-time instructions.

Proofs of concept: Valgrind
(uninitialized data as secret),
ctgrind, ct-verif, FlowTracker.

How can we implement, e.g.,
sorting of a secret array?

Eliminat

Let's try

Assume

ow from
conditions.

ow from
dices.

ow from
tate distances.

ctions.

uctions.

Us with

pliers: e.g.,

ost PowerPCs.

Suppose we know (some)
const-time machine instructions.

Suppose programming language
has “secret” types.

Easy for compiler to guarantee
that secret types are used only
by const-time instructions.

Proofs of concept: Valgrind
(uninitialized data as secret),
ctgrind, ct-verif, FlowTracker.

How can we implement, e.g.,
sorting of a secret array?

Eliminating branches

Let's try sorting 2

Assume `int32` is s

Suppose we know (some) const-time machine instructions. Suppose programming language has “secret” types.

Easy for compiler to guarantee that secret types are used only by const-time instructions.

Proofs of concept: Valgrind (uninitialized data as secret), ctgrind, ct-verif, FlowTracker.

How can we implement, e.g., sorting of a secret array?

Eliminating branches

Let's try sorting 2 integers. Assume `int32` is secret.

Suppose we know (some)
const-time machine instructions.

Suppose programming language
has “secret” types.

Easy for compiler to guarantee
that secret types are used only
by const-time instructions.

Proofs of concept: Valgrind
(uninitialized data as secret),
ctgrind, ct-verif, FlowTracker.

How can we implement, e.g.,
sorting of a secret array?

Eliminating branches

Let's try sorting 2 integers.

Assume `int32` is secret.

Suppose we know (some) const-time machine instructions. Suppose programming language has “secret” types.

Easy for compiler to guarantee that secret types are used only by const-time instructions.

Proofs of concept: Valgrind (uninitialized data as secret), ctgrind, ct-verif, FlowTracker.

How can we implement, e.g., sorting of a secret array?

Eliminating branches

Let's try sorting 2 integers. Assume int32 is secret.

```
void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}
```

Suppose we know (some) const-time machine instructions. Suppose programming language has “secret” types.

Easy for compiler to guarantee that secret types are used only by const-time instructions.

Proofs of concept: Valgrind (uninitialized data as secret), ctgrind, ct-verif, FlowTracker.

How can we implement, e.g., sorting of a secret array?

Eliminating branches

Let's try sorting 2 integers. Assume int32 is secret.

```
void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}
```

Unacceptable: not constant-time.

we know (some)
 some machine instructions.
 programming language
 "secret" types.

compiler to guarantee
 secret types are used only
 time instructions.

of concept: Valgrind
 lized data as secret),
 ct-verif, FlowTracker.

n we implement, e.g.,
 of a secret array?

Eliminating branches

Let's try sorting 2 integers.
 Assume int32 is secret.

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}
```

Unacceptable: not constant-time.

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}
```

14

(some)
 e instructions.
 ming language
 es.

to guarantee
 s are used only
 ructions.

Valgrind
 as secret),
 lowTracker.

ment, e.g.,
 et array?

Eliminating branches

Let's try sorting 2 integers.
 Assume int32 is secret.

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}
```

Unacceptable: not constant-time.

15

```
void sort2(int32
{ int32 x0 = x[0]
  int32 x1 = x[1]
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}
```

Eliminating branches

Let's try sorting 2 integers.

Assume `int32` is secret.

```
void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}
```

Unacceptable: not constant-time.

```
void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}
```


Eliminating branches

Let's try sorting 2 integers.

Assume `int32` is secret.

```

void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}

```

Unacceptable: not constant-time.

```

void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}

```

Eliminating branches

Let's try sorting 2 integers.

Assume `int32` is secret.

```

void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  }
}

```

Unacceptable: not constant-time.

```

void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}

```

Safe compiler won't allow this.
Branch timing leaks secrets.

ing branches

sorting 2 integers.

int32 is secret.

```
void sort2(int32 *x)
```

```
{ int32 x0 = x[0];
```

```
  int32 x1 = x[1];
```

```
  if (x1 < x0) {
```

```
    x[0] = x1;
```

```
    x[1] = x0;
```

table: not constant-time.

```
void sort2(int32 *x)
```

```
{ int32 x0 = x[0];
```

```
  int32 x1 = x[1];
```

```
  if (x1 < x0) {
```

```
    x[0] = x1;
```

```
    x[1] = x0;
```

```
  } else {
```

```
    x[0] = x0;
```

```
    x[1] = x1;
```

```
  }
```

```
}
```

Safe compiler won't allow this.

Branch timing leaks secrets.

```
void sort2(int32 *x)
```

```
{ int32 x0 = x[0];
```

```
  int32 x1 = x[1];
```

```
  if (x1 < x0) {
```

```
    x[0] = x1;
```

```
    x[1] = x0;
```

```
  }
```

15

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}

```

Safe compiler won't allow this.
Branch timing leaks secrets.

16

```

void sort2(int32
{ int32 x0 = x[0]
  int32 x1 = x[1]
  int32 c = (x1
x[0] = (c ? x1
x[1] = (c ? x0
}

```

15

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}
```

Safe compiler won't allow this.
Branch timing leaks secrets.

16

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[0] = (c ? x1 : x0);
  x[1] = (c ? x0 : x1);
}
```

-time.

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}
```

Safe compiler won't allow this.
Branch timing leaks secrets.

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[0] = (c ? x1 : x0);
  x[1] = (c ? x0 : x1);
}
```

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  if (x1 < x0) {
    x[0] = x1;
    x[1] = x0;
  } else {
    x[0] = x0;
    x[1] = x1;
  }
}

```

Safe compiler won't allow this.
Branch timing leaks secrets.

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[0] = (c ? x1 : x0);
  x[1] = (c ? x0 : x1);
}

```

Syntax is different but “?:”
is a branch by definition:

```

if (x1 < x0) x[0] = x1;
else x[0] = x0;
if (x1 < x0) x[1] = x0;
else x[1] = x1;

```

16

```

sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    if (x1 < x0) {
        x[0] = x1;
        x[1] = x0;
    }
    else {
        x[0] = x0;
        x[1] = x1;
    }
}

```

Compiler won't allow this.
 Branch timing leaks secrets.

17

```

void sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    int32 c = (x1 < x0);
    x[0] = (c ? x1 : x0);
    x[1] = (c ? x0 : x1);
}

```

Syntax is different but “?:”
 is a branch by definition:

```

if (x1 < x0) x[0] = x1;
else x[0] = x0;
if (x1 < x0) x[1] = x0;
else x[1] = x1;

```

```

void sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    int32 c = (x1 < x0);
    x[c] = x1;
    x[1 - c] = x0;
}

```


16

```

*x)
];
];

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[0] = (c ? x1 : x0);
  x[1] = (c ? x0 : x1);
}

```

Syntax is different but “?:”
is a branch by definition:

```

if (x1 < x0) x[0] = x1;
else x[0] = x0;
if (x1 < x0) x[1] = x0;
else x[1] = x1;

```

't allow this.
ks secrets.

17

```

void sort2(int32
{ int32 x0 = x[0]
  int32 x1 = x[1]
  int32 c = (x1
x[c] = x0;
x[1 - c] = x1;
}

```

16

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[0] = (c ? x1 : x0);
  x[1] = (c ? x0 : x1);
}
```

Syntax is different but “?:”
is a branch by definition:

```
if (x1 < x0) x[0] = x1;
else x[0] = x0;
if (x1 < x0) x[1] = x0;
else x[1] = x1;
```

17

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}
```

his.

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[0] = (c ? x1 : x0);
  x[1] = (c ? x0 : x1);
}
```

Syntax is different but “?:”
is a branch by definition:

```
if (x1 < x0) x[0] = x1;
else x[0] = x0;
if (x1 < x0) x[1] = x0;
else x[1] = x1;
```

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}
```

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[0] = (c ? x1 : x0);
  x[1] = (c ? x0 : x1);
}

```

Syntax is different but “?:”
is a branch by definition:

```

if (x1 < x0) x[0] = x1;
else x[0] = x0;
if (x1 < x0) x[1] = x0;
else x[1] = x1;

```

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}

```

Safe compiler won't allow this:
won't allow secret data
to be used as an array index.

Cache timing is not constant:
see earlier attack examples.

17

```

rt2(int32 *x)
  x0 = x[0];
  x1 = x[1];
  c = (x1 < x0);
  x[c] = (c ? x1 : x0);
  x[1 - c] = (c ? x0 : x1);

```

is different but “?:”
is not by definition:

```

(x1 < x0) x[0] = x1;
x[0] = x0;
(x1 < x0) x[1] = x0;
x[1] = x1;

```

18

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}

```

Safe compiler won't allow this:
won't allow secret data
to be used as an array index.

Cache timing is not constant:
see earlier attack examples.

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  c *= 1000000;
  x[c] = x0;
  x[1 - c] = x1;
}

```

17

```

*x)
];
];
< x0);
: x0);
: x1);

```

but “?:”

inition:

```

[0] = x1;
;
[1] = x0;
;

```

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}

```

Safe compiler won't allow this:
won't allow secret data
to be used as an array index.

Cache timing is not constant:
see earlier attack examples.

18

```

void sort2(int32
{ int32 x0 = x[0]
  int32 x1 = x[1]
  int32 c = (x1
  c *= x1 - x0;
  x[0] = x0 + c;
  x[1] = x1 - c;
}

```

17

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}
```

Safe compiler won't allow this:
won't allow secret data
to be used as an array index.

Cache timing is not constant:
see earlier attack examples.

18

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  c *= x1 - x0;
  x[0] = x0 + c;
  x[1] = x1 - c;
}
```

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}
```

Safe compiler won't allow this:
won't allow secret data
to be used as an array index.

Cache timing is not constant:
see earlier attack examples.

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  c *= x1 - x0;
  x[0] = x0 + c;
  x[1] = x1 - c;
}
```



```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  x[c] = x0;
  x[1 - c] = x1;
}

```

Safe compiler won't allow this:
won't allow secret data
to be used as an array index.

Cache timing is not constant:
see earlier attack examples.

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  c *= x1 - x0;
  x[0] = x0 + c;
  x[1] = x1 - c;
}

```

Does safe compiler allow
multiplication of secrets?

Recall that multiplication
takes variable time on, e.g.,
Cortex-M3 and most PowerPCs.

```

sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    int32 c = (x1 < x0);
    x[c] = x0;
    x[1-c] = x1;
}

```

Compiler won't allow this:

low secret data
used as an array index.

Timing is not constant:
side channel attack examples.

```

void sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    int32 c = (x1 < x0);
    c *= x1 - x0;
    x[0] = x0 + c;
    x[1] = x1 - c;
}

```

Does safe compiler allow
multiplication of secrets?

Recall that multiplication
takes variable time on, e.g.,
Cortex-M3 and most PowerPCs.

Will want
for fast p
but let's
for this s

```

void sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    int32 c = (x1 < x0);
    c &= x1 - x0;
    x[0] = x0 + c;
    x[1] = x1 - c;
}

```

18

```

*x)
];
];
< x0);

```

't allow this:

et data
array index.

ot constant:
examples.

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  c *= x1 - x0;
  x[0] = x0 + c;
  x[1] = x1 - c;
}

```

Does safe compiler allow
multiplication of secrets?

Recall that multiplication
takes variable time on, e.g.,
Cortex-M3 and most PowerPCs.

19

Will want to hand
for fast prime-field
but let's dodge the
for this sorting code

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}

```

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  c *= x1 - x0;
  x[0] = x0 + c;
  x[1] = x1 - c;
}

```

Does safe compiler allow
multiplication of secrets?

Recall that multiplication
takes variable time on, e.g.,
Cortex-M3 and most PowerPCs.

Will want to handle this issue
for fast prime-field ECC etc.
but let's dodge the issue
for this sorting code:

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}

```

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = (x1 < x0);
  c *= x1 - x0;
  x[0] = x0 + c;
  x[1] = x1 - c;
}

```

Does safe compiler allow multiplication of secrets?

Recall that multiplication takes variable time on, e.g., Cortex-M3 and most PowerPCs.

Will want to handle this issue for fast prime-field ECC etc., but let's dodge the issue for this sorting code:

```

void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}

```

```

void sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    int32 c = (x1 < x0);
    x0 = x0 - c;
    x1 = x1 + c;
}

```

Does compiler allow
constant-time
comparisons of secrets?

What about multiplication

with variable time on, e.g.,

ARMv7-M3 and most PowerPCs.

Will want to handle this issue
for fast prime-field ECC etc.,
but let's dodge the issue
for this sorting code:

```

void sort2(int32 *x)
{
    int32 x0 = x[0];
    int32 x1 = x[1];
    int32 c = -(x1 < x0);
    x0 ^= c;
    x1 ^= c;
}

```

1. Possible
(also for
C standard
int32 and
“undefined
Real CP
but C co

```

*x)
];
];
< x0);

```

r allow
 ecrets?
 lication
 e on, e.g.,
 ost PowerPCs.

Will want to handle this issue
 for fast prime-field ECC etc.,
 but let's dodge the issue
 for this sorting code:

```

void sort2(int32 *x)
{
  int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}

```

1. Possible correct
 (also for previous
 C standard does not
 int32 as twos-com
 “undefined” behav
 Real CPU uses two
 but *C compiler ca*

Will want to handle this issue for fast prime-field ECC etc., but let's dodge the issue for this sorting code:

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];

  int32 c = -(x1 < x0);

  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}
```

PCs.

1. Possible correctness problem (also for previous code): C standard does not define `int32` as twos-complement; “undefined” behavior on overflow. Real CPU uses twos-complement, but *C compiler can screw that*

Will want to handle this issue for fast prime-field ECC etc., but let's dodge the issue for this sorting code:

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}
```

1. Possible correctness problems (also for previous code):
C standard does not define `int32` as twos-complement; says “undefined” behavior on overflow. Real CPU uses twos-complement but *C compiler can screw this up.*

Will want to handle this issue for fast prime-field ECC etc., but let's dodge the issue for this sorting code:

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}
```

1. Possible correctness problems (also for previous code):
C standard does not define `int32` as twos-complement; says “undefined” behavior on overflow. Real CPU uses twos-complement but *C compiler can screw this up*.
Fix: use `gcc -fwrapv`.

Will want to handle this issue for fast prime-field ECC etc., but let's dodge the issue for this sorting code:

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}
```

1. Possible correctness problems (also for previous code):
C standard does not define `int32` as twos-complement; says “undefined” behavior on overflow. Real CPU uses twos-complement but *C compiler can screw this up.*

Fix: use `gcc -fwrapv`.

2. Does safe compiler allow “`x1 < x0`” for secrets?

What do we do if it doesn't?

Will want to handle this issue for fast prime-field ECC etc., but let's dodge the issue for this sorting code:

```
void sort2(int32 *x)
{ int32 x0 = x[0];
  int32 x1 = x[1];
  int32 c = -(x1 < x0);
  c &= x1 ^ x0;
  x[0] = x0 ^ c;
  x[1] = x1 ^ c;
}
```

1. Possible correctness problems (also for previous code):
C standard does not define `int32` as twos-complement; says “undefined” behavior on overflow. Real CPU uses twos-complement but *C compiler can screw this up.*

Fix: use `gcc -fwrapv`.

2. Does safe compiler allow “`x1 < x0`” for secrets?

What do we do if it doesn't?

C compilers *sometimes* use constant-time instructions for this.

nt to handle this issue
 prime-field ECC etc.,
 dodge the issue
 sorting code:

```
rt2(int32 *x)
  x0 = x[0];
  x1 = x[1];
  c = -(x1 < x0);
  x1 ^ x0;
  = x0 ^ c;
  = x1 ^ c;
```

1. Possible correctness problems
 (also for previous code):
 C standard does not define
 int32 as twos-complement; says
 “undefined” behavior on overflow.
 Real CPU uses twos-complement
 but *C compiler can screw this up.*

Fix: use `gcc -fwrapv`.

2. Does safe compiler allow
 “`x1 < x0`” for secrets?

What do we do if it doesn't?

C compilers *sometimes* use
 constant-time instructions for this.

Constant

```
int32 is
{ return
```

Returns

le this issue
l ECC etc.,

e issue

de:

```
*x)
```

```
];
```

```
];
```

```
< x0);
```

1. Possible correctness problems
(also for previous code):
C standard does not define
`int32` as twos-complement; says
“undefined” behavior on overflow.
Real CPU uses twos-complement
but *C compiler can screw this up.*

Fix: use `gcc -fwrapv`.

2. Does safe compiler allow
“`x1 < x0`” for secrets?

What do we do if it doesn't?

C compilers *sometimes* use
constant-time instructions for this.

Constant-time con

```
int32 isnegative
```

```
{ return x >> 31
```

Returns `-1` if `x <`

1. Possible correctness problems (also for previous code):
C standard does not define `int32` as twos-complement; says “undefined” behavior on overflow.
Real CPU uses twos-complement but *C compiler can screw this up.*

Fix: use `gcc -fwrapv`.

2. Does safe compiler allow “`x1 < x0`” for secrets?

What do we do if it doesn't?

C compilers *sometimes* use constant-time instructions for this.

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns `-1` if `x < 0`, otherwise

1. Possible correctness problems (also for previous code):
C standard does not define `int32` as twos-complement; says “undefined” behavior on overflow.
Real CPU uses twos-complement but *C compiler can screw this up.*

Fix: use `gcc -fwrapv`.

2. Does safe compiler allow “`x1 < x0`” for secrets?

What do we do if it doesn't?

C compilers *sometimes* use constant-time instructions for this.

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns `-1` if `x < 0`, otherwise `0`.

1. Possible correctness problems (also for previous code):
 C standard does not define `int32` as twos-complement; says “undefined” behavior on overflow.
 Real CPU uses twos-complement but *C compiler can screw this up.*

Fix: use `gcc -fwrapv`.

2. Does safe compiler allow “`x1 < x0`” for secrets?

What do we do if it doesn't?

C compilers *sometimes* use constant-time instructions for this.

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns `-1` if `x < 0`, otherwise `0`.

Why this works: the bits

$(b_{31}, b_{30}, \dots, b_2, b_1, b_0)$

represent the integer $b_0 + 2b_1 + 4b_2 + \dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“1-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_3, b_2, b_1)$.

“31-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_{31}, b_{31}, b_{31})$.

able correctness problems

(previous code):

ard does not define

s twos-complement; says

ed” behavior on overflow.

U uses twos-complement

compiler can screw this up.

`gcc -fwrapv.`

safe compiler allow

0” for secrets?

o we do if it doesn't?

lers *sometimes* use

t-time instructions for this.

Constant-time comparisons

```
int32 isnegative(int32 x)
```

```
{ return x >> 31; }
```

Returns -1 if $x < 0$, otherwise 0 .

Why this works: the bits

$(b_{31}, b_{30}, \dots, b_2, b_1, b_0)$

represent the integer $b_0 + 2b_1 + 4b_2 + \dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“1-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_3, b_2, b_1)$.

“31-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_{31}, b_{31}, b_{31})$.

```
int32 is
```

```
{ return
```

ness problems
 code):
 ot define
 mplement; says
 vior on overflow.
 os-complement
n screw this up.
 rapv.
 oiler allow
 ets?
 it doesn't?
times use
 ructions for this.

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns -1 if $x < 0$, otherwise 0 .

Why this works: the bits

$(b_{31}, b_{30}, \dots, b_2, b_1, b_0)$

represent the integer $b_0 + 2b_1 + 4b_2 + \dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“1-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_3, b_2, b_1)$.

“31-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_{31}, b_{31}, b_{31})$.

```
int32 ispositive
{ return isnegat
```

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns -1 if $x < 0$, otherwise 0 .

Why this works: the bits

$(b_{31}, b_{30}, \dots, b_2, b_1, b_0)$

represent the integer $b_0 + 2b_1 + 4b_2 + \dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“1-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_3, b_2, b_1)$.

“31-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_{31}, b_{31}, b_{31})$.

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns -1 if $x < 0$, otherwise 0 .

Why this works: the bits

$(b_{31}, b_{30}, \dots, b_2, b_1, b_0)$

represent the integer $b_0 + 2b_1 + 4b_2 + \dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“1-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_3, b_2, b_1)$.

“31-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_{31}, b_{31}, b_{31})$.

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns -1 if $x < 0$, otherwise 0 .

Why this works: the bits

$(b_{31}, b_{30}, \dots, b_2, b_1, b_0)$

represent the integer $b_0 + 2b_1 + 4b_2 + \dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“1-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_3, b_2, b_1)$.

“31-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_{31}, b_{31}, b_{31})$.

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Constant-time comparisons

```
int32 isnegative(int32 x)
{ return x >> 31; }
```

Returns -1 if $x < 0$, otherwise 0 .

Why this works: the bits

$(b_{31}, b_{30}, \dots, b_2, b_1, b_0)$

represent the integer $b_0 + 2b_1 + 4b_2 + \dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“1-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_3, b_2, b_1)$.

“31-bit signed right shift”:

$(b_{31}, b_{31}, \dots, b_{31}, b_{31}, b_{31})$.

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Constant-time comparisons

```
isnegative(int32 x)
```

```
{ return (x >> 31); }
```

-1 if $x < 0$, otherwise 0.

How it works: the bits

$b_3, \dots, b_2, b_1, b_0$

represent the integer $b_0 + 2b_1 +$

$\dots + 2^{30}b_{30} - 2^{31}b_{31}$.

“Arithmetic right shift”:

$b_3, \dots, b_3, b_2, b_1$.

“Logical right shift”:

$b_3, \dots, b_{31}, b_{31}, b_{31}$.

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Side note

```
int32 ispositive(int32 x)
```

```
{ if (x > 0)
```

```
    return 1;
}
```


Comparisons

```
(int32 x)
; }
```

0, otherwise 0.

the bits

b_1, b_0)

ger $b_0 + 2b_1 +$
 $- 2^{31}b_{31}$.

“right shift”:

b_2, b_1).

“left shift”:

b_{31}, b_{31}).

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Side note illustrating

```
int32 ispositive
{ if (x == -x) r
  return isnegat
```

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Side note illustrating `-fwrapv`

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
current gcc can remove the
`x == -x` test, breaking this code.

```
int32 ispositive(int32 x)
{ return isnegative(-x); }
```

This code is incorrect!

Fails for input -2^{31} ,
because “ $-x$ ” produces -2^{31} .

Can catch this bug by testing:

```
int64 x; int32 c;
for (x = INT32_MIN;
     x <= INT32_MAX; ++x) {
    c = ispositive(x);
    assert(c == -(x > 0));
}
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
current gcc can remove the
`x == -x` test, breaking this code.

Incompetent gcc engineering:
source of many security holes.
Incompetent language standard.

```
ispositive(int32 x)
{ return !isnegative(-x); }
```

code is incorrect!

input -2^{31} ,
 “ $-x$ ” produces -2^{31} .

check this bug by testing:

```
int32 c;
c = INT32_MIN;
for (int32 x = INT32_MIN; x <= INT32_MAX; ++x) {
    if (!ispositive(x))
        printf("c == -(x > 0));
}
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
 current gcc can remove the
`x == -x` test, breaking this code.

Incompetent gcc engineering:
 source of many security holes.
 Incompetent language standard.

```
int32 isnegative(int32 x)
{ return x < 0 || ispositive(x); }
```

```
(int32 x)
ive(-x); }
```

irect!

31,

duces -2^{31} .

g by testing:

;

IN;

```
MAX; ++x) {
```

```
(x);
```

```
x > 0));
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
current gcc can remove the
`x == -x` test, breaking this code.

Incompetent gcc engineering:
source of many security holes.
Incompetent language standard.

```
int32 isnonzero(
{ return isnegat
  || isnegative(
```


Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
current gcc can remove the
`x == -x` test, breaking this code.

Incompetent gcc engineering:
source of many security holes.
Incompetent language standard.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
current gcc can remove the
`x == -x` test, breaking this code.

Incompetent gcc engineering:
source of many security holes.
Incompetent language standard.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
current gcc can remove the
`x == -x` test, breaking this code.

Incompetent gcc engineering:
source of many security holes.
Incompetent language standard.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated
only if first part is zero.

Side note illustrating `-fwrapv`:

```
int32 ispositive(int32 x)
{ if (x == -x) return 0;
  return isnegative(-x); }
```

Not constant-time.

Even worse: without `-fwrapv`,
current gcc can remove the
`x == -x` test, breaking this code.

Incompetent gcc engineering:
source of many security holes.
Incompetent language standard.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated
only if first part is zero.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  | isnegative(-x); }
```

Constant-time logic instructions.
Safe compiler will allow this.

e illustrating `-fwrapv`:

```
ispositive(int32 x)
  == -x) return 0;
  n isnegative(-x); }
```

stant-time.

orse: without `-fwrapv`,
gcc can remove the
test, breaking this code.

otent gcc engineering:
f many security holes.
otent language standard.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated
only if first part is zero.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  | isnegative(-x); }
```

Constant-time logic instructions.
Safe compiler will allow this.

```
int32 is
{ return
```

ng -fwrapv:

```
(int32 x)
```

```
return 0;
```

```
ive(-x); }
```

.

out -fwrapv,

remove the

making this code.

engineering:

curity holes.

uage standard.

```
int32 isnonzero(int32 x)
```

```
{ return isnegative(x)
```

```
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated

only if first part is zero.

```
int32 isnonzero(int32 x)
```

```
{ return isnegative(x)
```

```
  | isnegative(-x); }
```

Constant-time logic instructions.

Safe compiler will allow this.

```
int32 issmaller(
```

```
{ return isnegat
```

pv:

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated
only if first part is zero.

pv,

code.

g:

es.

ard.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  | isnegative(-x); }
```

Constant-time logic instructions.

Safe compiler will allow this.

```
int32 issmaller(int32 x, i
{ return isnegative(x - y
```

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated
only if first part is zero.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  | isnegative(-x); }
```

Constant-time logic instructions.

Safe compiler will allow this.

```
int32 issmaller(int32 x,int32 y)
{ return isnegative(x - y); }
```



```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated
only if first part is zero.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  | isnegative(-x); }
```

Constant-time logic instructions.

Safe compiler will allow this.

```
int32 issmaller(int32 x, int32 y)
{ return isnegative(x - y); }
```

This code is incorrect!

Generalization of `ispositive`.

Wrong for inputs $(0, -2^{31})$.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  || isnegative(-x); }
```

Not constant-time.

Second part is evaluated
only if first part is zero.

```
int32 isnonzero(int32 x)
{ return isnegative(x)
  | isnegative(-x); }
```

Constant-time logic instructions.

Safe compiler will allow this.

```
int32 issmaller(int32 x,int32 y)
{ return isnegative(x - y); }
```

This code is incorrect!

Generalization of `ispositive`.

Wrong for inputs $(0, -2^{31})$.

Wrong for many more inputs.

Caught quickly by random tests:

```
for (j = 0; j < 10000000; ++j) {
  x += random(); y += random();
  c = issmaller(x,y);
  assert(c == -(x < y));
}
```

```
isnonzero(int32 x)
{
    if (x < 0)
        return isnegative(x)
        || isnegative(-x);
}
```

constant-time.

second part is evaluated
if first part is zero.

```
isnonzero(int32 x)
{
    if (x < 0)
        return isnegative(x)
        || isnegative(-x);
}
```

constant-time logic instructions.

compiler will allow this.

```
int32 issmaller(int32 x, int32 y)
{
    return isnegative(x - y);
}
```

This code is incorrect!

Generalization of `ispositive`.

Wrong for inputs $(0, -2^{31})$.

Wrong for many more inputs.

Caught quickly by random tests:

```
for (j = 0; j < 100000000; ++j) {
    x += random(); y += random();
    c = issmaller(x, y);
    assert(c == -(x < y));
}
```

```
int32 issmaller(int32 x, int32 y)
{
    int32 c;
    if (x < 0)
        c ^= 1;
    return (x < y) ^ c;
}
```

```
int32 x)
ive(x)
-x); }
```

evaluated
zero.

```
int32 x)
ive(x)
x); }
```

ic instructions.
allow this.

```
int32 issmaller(int32 x,int32 y)
{ return isnegative(x - y); }
```

This code is incorrect!

Generalization of `ispositive`.

Wrong for inputs $(0, -2^{31})$.

Wrong for many more inputs.

Caught quickly by random tests:

```
for (j = 0;j < 100000000;++j) {
    x += random(); y += random();
    c = issmaller(x,y);
    assert(c == -(x < y));
}
```

```
int32 issmaller(
{ int32 xy = x ^
  int32 c = x -
  c ^= xy & (c ^
return isnegat
}
```

```
int32 issmaller(int32 x,int32 y)
{ return isnegative(x - y); }
```

This code is incorrect!

Generalization of `ispositive`.

Wrong for inputs $(0, -2^{31})$.

Wrong for many more inputs.

Caught quickly by random tests:

```
for (j = 0;j < 100000000;++j) {
  x += random(); y += random();
  c = issmaller(x,y);
  assert(c == -(x < y));
}
```

```
int32 issmaller(int32 x,i
{ int32 xy = x ^ y;
  int32 c = x - y;
  c ^= xy & (c ^ x);
  return isnegative(c);
}
```

```
int32 issmaller(int32 x,int32 y)
{ return isnegative(x - y); }
```

This code is incorrect!

Generalization of `ispositive`.

Wrong for inputs $(0, -2^{31})$.

Wrong for many more inputs.

Caught quickly by random tests:

```
for (j = 0;j < 100000000;++j) {
    x += random(); y += random();
    c = issmaller(x,y);
    assert(c == -(x < y));
}
```

```
int32 issmaller(int32 x,int32 y)
{ int32 xy = x ^ y;
  int32 c = x - y;
  c ^= xy & (c ^ x);
  return isnegative(c);
}
```

```
int32 issmaller(int32 x,int32 y)
{ return isnegative(x - y); }
```

This code is incorrect!

Generalization of `ispositive`.

Wrong for inputs $(0, -2^{31})$.

Wrong for many more inputs.

Caught quickly by random tests:

```
for (j = 0;j < 10000000; ++j) {
    x += random(); y += random();
    c = issmaller(x,y);
    assert(c == -(x < y));
}
```

```
int32 issmaller(int32 x,int32 y)
{ int32 xy = x ^ y;
  int32 c = x - y;
  c ^= xy & (c ^ x);
  return isnegative(c);
}
```

Some verification strategies:

- Think this through.
- Write a proof.
- Formally verify proof.
- Automate proof construction.
- Test many random inputs.
- A bit painful: test all inputs.
- Faster: test `int16` version.

```
issmaller(int32 x,int32 y)
return isnegative(x - y); }
```

code is incorrect!

mis-implementation of `ispositive`.

fails for inputs $(0, -2^{31})$.

fails for many more inputs.

can be fixed quickly by random tests:

```
for (int i = 0; i < 100000000; ++i) {
    int x = random(); y = random();
    if (!issmaller(x,y) ||
        !test(c == -(x < y)));
}
```

```
int32 issmaller(int32 x,int32 y)
{ int32 xy = x ^ y;
  int32 c = x - y;
  c ^= xy & (c ^ x);
  return isnegative(c);
}
```

Some verification strategies:

- Think this through.
- Write a proof.
- Formally verify proof.
- Automate proof construction.
- Test many random inputs.
- A bit painful: test all inputs.
- Faster: test `int16` version.

```
void minmax(int32 x,int32 y)
{ int32 m = x;
  int32 M = x;
  if (y < m) m = y;
  if (y > M) M = y;
}
```

```
void sort(int32 a[])
{ minmax(a[0],a[1]);
  minmax(a[1],a[2]);
  minmax(a[2],a[3]);
  minmax(a[3],a[4]);
  minmax(a[4],a[5]);
  minmax(a[5],a[6]);
  minmax(a[6],a[7]);
  minmax(a[7],a[8]);
  minmax(a[8],a[9]);
}
```



```
int32 x,int32 y)
ive(x - y); }
```

rect!

ispositive.

$(0, -2^{31})$.

more inputs.

random tests:

```
0000000; ++j) {
  y += random();
  x, y);
  x < y));
```

```
int32 issmaller(int32 x,int32 y)
{ int32 xy = x ^ y;
  int32 c = x - y;
  c ^= xy & (c ^ x);
  return isnegative(c);
}
```

Some verification strategies:

- Think this through.
- Write a proof.
- Formally verify proof.
- Automate proof construction.
- Test many random inputs.
- A bit painful: test all inputs.
- Faster: test int16 version.

```
void minmax(int32
{ int32 a = *x;
  int32 b = *y;
  int32 ab = b ^
  int32 c = b -
  c ^= ab & (c ^
  c >>= 31;
  c &= ab;
  *x = a ^ c;
  *y = b ^ c;
}

void sort2(int32
{ minmax(x, x + 1
```

```
int32 y)
); }
```

ve.

s.

ests:

```
+j) {
dom();
```

```
int32 issmaller(int32 x,int32 y)
{ int32 xy = x ^ y;
  int32 c = x - y;
  c ^= xy & (c ^ x);
  return isnegative(c);
}
```

Some verification strategies:

- Think this through.
- Write a proof.
- Formally verify proof.
- Automate proof construction.
- Test many random inputs.
- A bit painful: test all inputs.
- Faster: test int16 version.

```
void minmax(int32 *x,int32 *y)
{ int32 a = *x;
  int32 b = *y;
  int32 ab = b ^ a;
  int32 c = b - a;
  c ^= ab & (c ^ b);
  c >>= 31;
  c &= ab;
  *x = a ^ c;
  *y = b ^ c;
}

void sort2(int32 *x)
{ minmax(x,x + 1); }
```

```

int32 issmaller(int32 x,int32 y)
{ int32 xy = x ^ y;
  int32 c = x - y;
  c ^= xy & (c ^ x);
  return isnegative(c);
}

```

Some verification strategies:

- Think this through.
- Write a proof.
- Formally verify proof.
- Automate proof construction.
- Test many random inputs.
- A bit painful: test all inputs.
- Faster: test int16 version.

```

void minmax(int32 *x,int32 *y)
{ int32 a = *x;
  int32 b = *y;
  int32 ab = b ^ a;
  int32 c = b - a;
  c ^= ab & (c ^ b);
  c >>= 31;
  c &= ab;
  *x = a ^ c;
  *y = b ^ c;
}

void sort2(int32 *x)
{ minmax(x,x + 1); }

```

```

isSmaller(int32 x,int32 y)
{
    xy = x ^ y;
    c = x - y;
    xy & (c ^ x);
    return isnegative(c);
}

```

Verification strategies:

test this through.

write a proof.

manually verify proof.

automate proof construction.

use many random inputs.

more painful: test all inputs.

alternatively: test int16 version.

```

void minmax(int32 *x,int32 *y)
{
    int32 a = *x;
    int32 b = *y;
    int32 ab = b ^ a;
    int32 c = b - a;
    c ^= ab & (c ^ b);
    c >>= 31;
    c &= ab;
    *x = a ^ c;
    *y = b ^ c;
}

void sort2(int32 *x)
{
    minmax(x,x + 1);
}

```

```

int32 isSmaller(int32 x,int32 y)
{
    int32 xy = x ^ y;
    int32 c = x - y;
    xy &= (c ^ x);
    return isnegative(c);
}

void sort2(int32 *x)
{
    long long n = x[0];
    for (int i = 1; i < n; i++)
        for (int j = i + 1; j < n; j++)
            minmax(x+i,x+j);
}

```

Safe compare
if array is sorted

27

```
int32 x,int32 y)
  y;
y;
x);
ive(c);
```

strategies:

gh.

proof.

construction.

om inputs.

est all inputs.

16 version.

```
void minmax(int32 *x,int32 *y)
{ int32 a = *x;
  int32 b = *y;
  int32 ab = b ^ a;
  int32 c = b - a;
  c ^= ab & (c ^ b);
  c >>= 31;
  c &= ab;
  *x = a ^ c;
  *y = b ^ c;
}
```

```
void sort2(int32 *x)
{ minmax(x,x + 1); }
```

28

```
int32 ispositive
{ int32 c = -x;
  c ^= x & c;
  return isnegat
}
```

```
void sort(int32
{ long long i,j;
  for (j = 0;j <
    for (i = j -
      minmax(x +
  }
```

Safe compiler will

if array length n is

27

```

int32 y) void minmax(int32 *x,int32 *y)
{ int32 a = *x;
  int32 b = *y;
  int32 ab = b ^ a;
  int32 c = b - a;
  c ^= ab & (c ^ b);
  c >>= 31;
  c &= ab;
  *x = a ^ c;
  *y = b ^ c;
}

void sort2(int32 *x)
{ minmax(x,x + 1); }

```

28

```

int32 ispositive(int32 x)
{ int32 c = -x;
  c ^= x & c;
  return isnegative(c);
}

void sort(int32 *x,long l
{ long long i,j;
  for (j = 0;j < n;++j)
    for (i = j - 1;i >= 0
        minmax(x + i,x + i
}

```

Safe compiler will allow this
if array length n is not secret

```

void minmax(int32 *x,int32 *y)
{ int32 a = *x;
  int32 b = *y;
  int32 ab = b ^ a;
  int32 c = b - a;
  c ^= ab & (c ^ b);
  c >>= 31;
  c &= ab;
  *x = a ^ c;
  *y = b ^ c;
}

void sort2(int32 *x)
{ minmax(x,x + 1); }

```

```

int32 ispositive(int32 x)
{ int32 c = -x;
  c ^= x & c;
  return isnegative(c);
}

void sort(int32 *x,long long n)
{ long long i,j;
  for (j = 0;j < n;++j)
    for (i = j - 1;i >= 0;--i)
      minmax(x + i,x + i + 1);
}

```

Safe compiler will allow this
if array length `n` is not secret.

```

minmax(int32 *x,int32 *y)
{
    a = *x;
    b = *y;
    ab = b ^ a;
    c = b - a;
    ab & (c ^ b);
    31;
    ab;
    a ^ c;
    b ^ c;
}

rt2(int32 *x)
{
    x(x,x + 1); }

```

```

int32 ispositive(int32 x)
{
    int32 c = -x;
    c ^= x & c;
    return isnegative(c);
}

void sort(int32 *x,long long n)
{
    long long i,j;
    for (j = 0;j < n;++j)
        for (i = j - 1;i >= 0;--i)
            minmax(x + i,x + i + 1);
}

```

Safe compiler will allow this
if array length n is not secret.

Almost a
much slo


```
int32 *x, int32 *y)
```

```
    a;  
    a;  
    b);
```

```
    *x)  
); }
```

```
int32 ispositive(int32 x)  
{ int32 c = -x;  
  c ^= x & c;  
  return isnegative(c);  
}
```

```
void sort(int32 *x, long long n)  
{ long long i, j;  
  for (j = 0; j < n; ++j)  
    for (i = j - 1; i >= 0; --i)  
      minmax(x + i, x + i + 1);  
}
```

Safe compiler will allow this
if array length `n` is not secret.

Software optimization

Almost all software
much slower than

```
int *y)
```

```
int32 ispositive(int32 x)
{ int32 c = -x;
  c ^= x & c;
  return isnegative(c);
}
```

```
void sort(int32 *x, long long n)
{ long long i, j;
  for (j = 0; j < n; ++j)
    for (i = j - 1; i >= 0; --i)
      minmax(x + i, x + i + 1);
}
```

Safe compiler will allow this
if array length `n` is not secret.

Software optimization

Almost all software is
much slower than it could be

```

int32 ispositive(int32 x)
{
    int32 c = -x;
    c ^= x & c;
    return isnegative(c);
}

void sort(int32 *x, long long n)
{
    long long i, j;
    for (j = 0; j < n; ++j)
        for (i = j - 1; i >= 0; --i)
            minmax(x + i, x + i + 1);
}

```

Safe compiler will allow this
if array length `n` is not secret.

Software optimization

Almost all software is
much slower than it could be.

```

int32 ispositive(int32 x)
{
    int32 c = -x;
    c ^= x & c;
    return isnegative(c);
}

void sort(int32 *x, long long n)
{
    long long i, j;
    for (j = 0; j < n; ++j)
        for (i = j - 1; i >= 0; --i)
            minmax(x + i, x + i + 1);
}

```

Safe compiler will allow this
if array length `n` is not secret.

Software optimization

Almost all software is
much slower than it could be.

Is software applied to much data?
Usually not. Usually the
wasted CPU time is negligible.

```

int32 ispositive(int32 x)
{
  int32 c = -x;
  c ^= x & c;
  return isnegative(c);
}

void sort(int32 *x, long long n)
{
  long long i, j;
  for (j = 0; j < n; ++j)
    for (i = j - 1; i >= 0; --i)
      minmax(x + i, x + i + 1);
}

```

Safe compiler will allow this
if array length `n` is not secret.

Software optimization

Almost all software is
much slower than it could be.

Is software applied to much data?
Usually not. Usually the
wasted CPU time is negligible.

But *crypto software* should be
applied to all communication.

Crypto that's too slow \Rightarrow
fewer users \Rightarrow fewer cryptanalysts
 \Rightarrow less attractive for everybody.

```

ispositive(int32 x)
{
    c = -x;
    x & c;
    return isnegative(c);
}

int32 *x, long long n)
long i, j;
for (j = 0; j < n; ++j)
    for (i = j - 1; i >= 0; --i)
        inmax(x + i, x + i + 1);
}

Compiler will allow this
length n is not secret.

```

Software optimization

Almost all software is
much slower than it could be.

Is software applied to much data?

Usually not. Usually the
wasted CPU time is negligible.

But *crypto software* should be
applied to all communication.

Crypto that's too slow \Rightarrow

fewer users \Rightarrow fewer cryptanalysts

\Rightarrow less attractive for everybody.

Typical s

X is a c

You hav

reference

You wan

software

as effici

You hav

(Can rep

You mea

impleme

```
(int32 x)
```

```
ive(c);
```

```
*x, long long n)
```

```
n; ++j)
```

```
1; i >= 0; --i)
```

```
i, x + i + 1);
```

allow this

is not secret.

Software optimization

Almost all software is
much slower than it could be.

Is software applied to much data?

Usually not. Usually the
wasted CPU time is negligible.

But *crypto software* should be
applied to all communication.

Crypto that's too slow \Rightarrow
fewer users \Rightarrow fewer cryptanalysts
 \Rightarrow less attractive for everybody.

Typical situation:

X is a cryptographic

You have written a
reference impleme

You want (const-t
software that com
as efficiently as po

You have chosen a
(Can repeat for ot

You measure perfor
implementation. M

Software optimization

Almost all software is much slower than it could be.

Is software applied to much data?

Usually not. Usually the wasted CPU time is negligible.

But *crypto software* should be applied to all communication.

Crypto that's too slow \Rightarrow
fewer users \Rightarrow fewer cryptanalysts
 \Rightarrow less attractive for everybody.

Typical situation:

X is a cryptographic system

You have written a (const-time) reference implementation of

You want (const-time) software that computes X as efficiently as possible.

You have chosen a target CPU
(Can repeat for other CPUs.)

You measure performance of implementation. Now what?

Software optimization

Almost all software is much slower than it could be.

Is software applied to much data?

Usually not. Usually the wasted CPU time is negligible.

But *crypto software* should be applied to all communication.

Crypto that's too slow \Rightarrow
fewer users \Rightarrow fewer cryptanalysts
 \Rightarrow less attractive for everybody.

Typical situation:

X is a cryptographic system.

You have written a (const-time) reference implementation of X .

You want (const-time) software that computes X as efficiently as possible.

You have chosen a target CPU.
(Can repeat for other CPUs.)

You measure performance of the implementation. Now what?

Performance optimization

All software is

slower than it could be.

Optimizations are applied to much data?

Not. Usually the

overhead of CPU time is negligible.

Optimizing software should be

applied to all communication.

That's too slow \Rightarrow

fewer cryptanalysts

attractive for everybody.

Typical situation:

X is a cryptographic system.

You have written a (const-time) reference implementation of X .

You want (const-time) software that computes X as efficiently as possible.

You have chosen a target CPU. (Can repeat for other CPUs.)

You measure performance of the implementation. Now what?

A simplified

Target CPU

microcontroller

one ARM

Reference

```
int sum
```

```
{
```

```
    int r
```

```
    int i
```

```
    for (i
```

```
        res
```

```
    return
```

```
}
```

tion

e is

it could be.

l to much data?

lly the

is negligible.

re should be

munication.

slow \Rightarrow

ver cryptanalysts

for everybody.

Typical situation:

X is a cryptographic system.

You have written a (const-time) reference implementation of X .

You want (const-time) software that computes X as efficiently as possible.

You have chosen a target CPU.
(Can repeat for other CPUs.)

You measure performance of the implementation. Now what?

A simplified example

Target CPU: TI LM32

microcontroller code

one ARM Cortex-M

Reference implemente

```
int sum(int *x)
{
    int result = 0;
    int i;
    for (i = 0; i <
        result += x[
    return result;
}
```

Typical situation:

X is a cryptographic system.

You have written a (const-time) reference implementation of X .

You want (const-time) software that computes X as efficiently as possible.

You have chosen a target CPU.
(Can repeat for other CPUs.)

You measure performance of the implementation. Now what?

A simplified example

Target CPU: TI LM4F120H5
microcontroller containing
one ARM Cortex-M4F core.

Reference implementation:

```
int sum(int *x)
{
    int result = 0;
    int i;
    for (i = 0; i < 1000; ++i)
        result += x[i];
    return result;
}
```

Typical situation:

X is a cryptographic system.

You have written a (const-time) reference implementation of X .

You want (const-time) software that computes X as efficiently as possible.

You have chosen a target CPU.
(Can repeat for other CPUs.)

You measure performance of the implementation. Now what?

A simplified example

Target CPU: TI LM4F120H5QR microcontroller containing one ARM Cortex-M4F core.

Reference implementation:

```
int sum(int *x)
{
    int result = 0;
    int i;
    for (i = 0; i < 1000; ++i)
        result += x[i];
    return result;
}
```

situation:

cryptographic system.

we written a (const-time)
implementation of X .

at (const-time)

that computes X
as fast as possible.

we chosen a target CPU.
(beat for other CPUs.)

measure performance of the
implementation. Now what?

A simplified example

Target CPU: TI LM4F120H5QR
microcontroller containing
one ARM Cortex-M4F core.

Reference implementation:

```
int sum(int *x)
{
    int result = 0;
    int i;
    for (i = 0; i < 1000; ++i)
        result += x[i];
    return result;
}
```

Counting

static v

*const

= (vo

...

int bef

int res

int aft

UARTpri

result

Output s

Change

A simplified example

Target CPU: TI LM4F120H5QR
 microcontroller containing
 one ARM Cortex-M4F core.

Reference implementation:

```
int sum(int *x)
{
    int result = 0;
    int i;
    for (i = 0; i < 1000; ++i)
        result += x[i];
    return result;
}
```

Counting cycles:

```
static volatile
    *const DWT_CYC
    = (void *) 0xE
    ...
```

```
int beforesum =
int result = sum
int aftersum = *
UARTprintf("sum
    result, aftersu
```

Output shows 801

Change 1000 to 50

A simplified example

Target CPU: TI LM4F120H5QR
microcontroller containing
one ARM Cortex-M4F core.

Reference implementation:

```
int sum(int *x)
{
    int result = 0;
    int i;
    for (i = 0; i < 1000; ++i)
        result += x[i];
    return result;
}
```

Counting cycles:

```
static volatile unsigned
    *const DWT_CYCCNT
    = (void *) 0xE0001004;
...
```

```
int beforesum = *DWT_CYCCNT;
int result = sum(x);
int aftersum = *DWT_CYCCNT;
UARTprintf("sum %d %d\n",
    result, aftersum - beforesum);
```

Output shows 8012 cycles.
Change 1000 to 500: 4012.

A simplified example

Target CPU: TI LM4F120H5QR
microcontroller containing
one ARM Cortex-M4F core.

Reference implementation:

```
int sum(int *x)
{
    int result = 0;
    int i;
    for (i = 0; i < 1000; ++i)
        result += x[i];
    return result;
}
```

Counting cycles:

```
static volatile unsigned int
    *const DWT_CYCCNT
    = (void *) 0xE0001004;
...

int beforesum = *DWT_CYCCNT;
int result = sum(x);
int aftersum = *DWT_CYCCNT;
UARTprintf("sum %d %d\n",
    result, aftersum - beforesum);
```

Output shows 8012 cycles.

Change 1000 to 500: 4012.

fied example

CPU: TI LM4F120H5QR

controller containing

M Cortex-M4F core.

ce implementation:

```
(int *x)
```

```
result = 0;
```

```
;
```

```
for (i = 0; i < 1000; ++i)
```

```
    result += x[i];
```

```
return result;
```

Counting cycles:

```
static volatile unsigned int
```

```
    *const DWT_CYCCNT
```

```
    = (void *) 0xE0001004;
```

```
...
```

```
int beforesum = *DWT_CYCCNT;
```

```
int result = sum(x);
```

```
int aftersum = *DWT_CYCCNT;
```

```
UARTprintf("sum %d %d\n",
```

```
    result, aftersum - beforesum);
```

Output shows 8012 cycles.

Change 1000 to 500: 4012.

“Okay, &

Um, are

really th

ole

M4F120H5QR

ntaining

M4F core.

entation:

```

;
1000; ++i)
i];

```

Counting cycles:

```

static volatile unsigned int
    *const DWT_CYCCNT
    = (void *) 0xE0001004;
...

int beforesum = *DWT_CYCCNT;
int result = sum(x);
int aftersum = *DWT_CYCCNT;
UARTprintf("sum %d %d\n",
    result, aftersum - beforesum);

```

Output shows 8012 cycles.

Change 1000 to 500: 4012.

“Okay, 8 cycles per
Um, are microcont
really this slow at

Counting cycles:

```
static volatile unsigned int
    *const DWT_CYCCNT
    = (void *) 0xE0001004;
...

int beforesum = *DWT_CYCCNT;
int result = sum(x);
int aftersum = *DWT_CYCCNT;
UARTprintf("sum %d %d\n",
    result, aftersum - beforesum);
```

Output shows 8012 cycles.

Change 1000 to 500: 4012.

“Okay, 8 cycles per addition
Um, are microcontrollers
really this slow at addition?”

Counting cycles:

```
static volatile unsigned int
    *const DWT_CYCCNT
    = (void *) 0xE0001004;
...

int beforesum = *DWT_CYCCNT;
int result = sum(x);
int aftersum = *DWT_CYCCNT;
UARTprintf("sum %d %d\n",
    result, aftersum-beforesum);
```

Output shows 8012 cycles.

Change 1000 to 500: 4012.

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Counting cycles:

```
static volatile unsigned int
    *const DWT_CYCCNT
    = (void *) 0xE0001004;
...

int beforesum = *DWT_CYCCNT;
int result = sum(x);
int aftersum = *DWT_CYCCNT;
UARTprintf("sum %d %d\n",
    result, aftersum-beforesum);
```

Output shows 8012 cycles.

Change 1000 to 500: 4012.

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Counting cycles:

```
static volatile unsigned int
    *const DWT_CYCCNT
    = (void *) 0xE0001004;
...

int beforesum = *DWT_CYCCNT;
int result = sum(x);
int aftersum = *DWT_CYCCNT;
UARTprintf("sum %d %d\n",
    result, aftersum-beforesum);
```

Output shows 8012 cycles.

Change 1000 to 500: 4012.

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Good practice:

Figure out lower bound for
cycles spent on arithmetic etc.

Understand gap between
lower bound and observed time.

g cycles:

```
volatile unsigned int
t DWT_CYCCNT
id *) 0xE0001004;
```

```
oresum = *DWT_CYCCNT;
ult = sum(x);
ersum = *DWT_CYCCNT;
ntf("sum %d %d\n",
t, aftersum-beforesum);
```

shows 8012 cycles.

1000 to 500: 4012.

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Good practice:

Figure out lower bound for
cycles spent on arithmetic etc.

Understand gap between
lower bound and observed time.

Find “A
Technical
Rely on
M4F = I


```

unsigned int
CNT
0001004;

*DWT_CYCCNT;
(x);
DWT_CYCCNT;
%d %d\n",
m-beforesum);

2 cycles.
00: 4012.

```

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Good practice:

Figure out lower bound for
cycles spent on arithmetic etc.

Understand gap between
lower bound and observed time.

Find “ARM Cortex
Technical Reference
Rely on Wikipedia
 $M4F = M4 + float$

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Good practice:

Figure out lower bound for
cycles spent on arithmetic etc.

Understand gap between
lower bound and observed time.

Find “ARM Cortex-M4 Pro
Technical Reference Manual
Rely on Wikipedia comment
M4F = M4 + floating-point

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Good practice:

Figure out lower bound for
cycles spent on arithmetic etc.

Understand gap between
lower bound and observed time.

Find “ARM Cortex-M4 Processor
Technical Reference Manual”.

Rely on Wikipedia comment that
 $M4F = M4 + \text{floating-point unit}$.

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Good practice:

Figure out lower bound for
cycles spent on arithmetic etc.
Understand gap between
lower bound and observed time.

Find “ARM Cortex-M4 Processor
Technical Reference Manual” .

Rely on Wikipedia comment that
 $M4F = M4 + \text{floating-point unit}$.

Manual says that Cortex-M4
“implements the ARMv7E-M
architecture profile” .

“Okay, 8 cycles per addition.
Um, are microcontrollers
really this slow at addition?”

Bad practice:

Apply random “optimizations”
(and tweak compiler options)
until you get bored.

Keep the fastest results.

Good practice:

Figure out lower bound for
cycles spent on arithmetic etc.
Understand gap between
lower bound and observed time.

Find “ARM Cortex-M4 Processor
Technical Reference Manual” .

Rely on Wikipedia comment that
 $M4F = M4 + \text{floating-point unit}$.

Manual says that Cortex-M4
“implements the ARMv7E-M
architecture profile” .

Points to the “ARMv7-M
Architecture Reference Manual” ,
which defines instructions:
e.g., “ADD” for 32-bit addition.

First manual says that
ADD takes just 1 cycle.

3 cycles per addition.

microcontrollers

is slow at addition?"

ctice:

andom "optimizations"

reak compiler options)

u get bored.

e fastest results.

actice:

ut lower bound for

pent on arithmetic etc.

and gap between

ound and observed time.

Find "ARM Cortex-M4 Processor
Technical Reference Manual".

Rely on Wikipedia comment that
M4F = M4 + floating-point unit.

Manual says that Cortex-M4
"implements the ARMv7E-M
architecture profile".

Points to the "ARMv7-M
Architecture Reference Manual",
which defines instructions:
e.g., "ADD" for 32-bit addition.

First manual says that
ADD takes just 1 cycle.

Inputs a

"integer

has 16 in

special-p

and "pro

er addition.
 trollers
 addition?"

optimizations"
 der options)
 d.
 results.

ound for
 arithmetic etc.
 etween
 bserved time.

Find "ARM Cortex-M4 Processor
 Technical Reference Manual".
 Rely on Wikipedia comment that
 $M4F = M4 + \text{floating-point unit}$.

Manual says that Cortex-M4
 "implements the ARMv7E-M
 architecture profile".

Points to the "ARMv7-M
 Architecture Reference Manual",
 which defines instructions:
 e.g., "ADD" for 32-bit addition.

First manual says that
 ADD takes just 1 cycle.

Inputs and output
 "integer registers"
 has 16 integer reg
 special-purpose "s
 and "program cou

Find “ARM Cortex-M4 Processor Technical Reference Manual” .

Rely on Wikipedia comment that $M4F = M4 + \text{floating-point unit}$.

Manual says that Cortex-M4 “implements the ARMv7E-M architecture profile” .

Points to the “ARMv7-M Architecture Reference Manual” , which defines instructions: e.g., “ADD” for 32-bit addition.

First manual says that ADD takes just 1 cycle.

Inputs and output of ADD a “integer registers” . ARMv7- has 16 integer registers, incl special-purpose “stack pointer” and “program counter” .

Find “ARM Cortex-M4 Processor Technical Reference Manual” .

Rely on Wikipedia comment that $M4F = M4 + \text{floating-point unit}$.

Manual says that Cortex-M4 “implements the ARMv7E-M architecture profile” .

Points to the “ARMv7-M Architecture Reference Manual” , which defines instructions: e.g., “ADD” for 32-bit addition.

First manual says that ADD takes just 1 cycle.

Inputs and output of ADD are “integer registers” . ARMv7-M has 16 integer registers, including special-purpose “stack pointer” and “program counter” .

Find “ARM Cortex-M4 Processor Technical Reference Manual” .

Rely on Wikipedia comment that $M4F = M4 + \text{floating-point unit}$.

Manual says that Cortex-M4

“implements the ARMv7E-M architecture profile” .

Points to the “ARMv7-M Architecture Reference Manual” , which defines instructions:

e.g., “ADD” for 32-bit addition.

First manual says that

ADD takes just 1 cycle.

Inputs and output of ADD are “integer registers” . ARMv7-M has 16 integer registers, including special-purpose “stack pointer” and “program counter” .

Each element of x array needs to be “loaded” into a register.

Find “ARM Cortex-M4 Processor Technical Reference Manual” .

Rely on Wikipedia comment that $M4F = M4 + \text{floating-point unit}$.

Manual says that Cortex-M4

“implements the ARMv7E-M architecture profile” .

Points to the “ARMv7-M Architecture Reference Manual” , which defines instructions:

e.g., “ADD” for 32-bit addition.

First manual says that ADD takes just 1 cycle.

Inputs and output of ADD are “integer registers” . ARMv7-M has 16 integer registers, including special-purpose “stack pointer” and “program counter” .

Each element of x array needs to be “loaded” into a register.

Basic load instruction: LDR.

Manual says 2 cycles but adds a note about “pipelining” .

Then more explanation: if next instruction is also LDR (with address not based on first LDR) then it saves 1 cycle.

ARM Cortex-M4 Processor
Reference Manual”.

Wikipedia comment that
M4 + floating-point unit.

says that Cortex-M4
implements the ARMv7E-M
architecture profile”.

to the “ARMv7-M
Architecture Reference Manual”,
defines instructions:

“ADD” for 32-bit addition.

Manual says that
takes just 1 cycle.

Inputs and output of ADD are
“integer registers”. ARMv7-M
has 16 integer registers, including
special-purpose “stack pointer”
and “program counter”.

Each element of x array needs to
be “loaded” into a register.

Basic load instruction: LDR.

Manual says 2 cycles but adds
a note about “pipelining”.

Then more explanation: if next
instruction is also LDR (with
address not based on first LDR)
then it saves 1 cycle.

n consec
takes on
(“more
pipelined

Can ach
in other
but noth

Lower bo
 $2n + 1$ c
including

Why obs
non-cons
costs of

Cortex-M4 Processor
Reference Manual”.

comment that
the unit.

Cortex-M4

ARMv7E-M

”.

ARMv7-M

Reference Manual”,

instructions:

2-bit addition.

that

cycle.

Inputs and output of ADD are
“integer registers”. ARMv7-M
has 16 integer registers, including
special-purpose “stack pointer”
and “program counter”.

Each element of x array needs to
be “loaded” into a register.

Basic load instruction: LDR.

Manual says 2 cycles but adds
a note about “pipelining”.

Then more explanation: if next
instruction is also LDR (with
address not based on first LDR)
then it saves 1 cycle.

n consecutive LDR
takes only $n + 1$ cycles
(“more multiple LDR
pipelined together”)

Can achieve this savings
in other ways (LDR
but nothing seems

Lower bound for n
 $2n + 1$ cycles,
including n cycles

Why observed time
non-consecutive LDR
costs of manipulating

Inputs and output of ADD are “integer registers”. ARMv7-M has 16 integer registers, including special-purpose “stack pointer” and “program counter”.

Each element of x array needs to be “loaded” into a register.

Basic load instruction: LDR.

Manual says 2 cycles but adds a note about “pipelining”.

Then more explanation: if next instruction is also LDR (with address not based on first LDR) then it saves 1 cycle.

n consecutive LDRs takes only $n + 1$ cycles (“more multiple LDRs can be pipelined together”).

Can achieve this speed in other ways (LDRD, LDM) but nothing seems faster.

Lower bound for n LDR + n $2n + 1$ cycles, including n cycles of arithmetic.

Why observed time is higher for non-consecutive LDRs; costs of manipulating i .

Inputs and output of ADD are “integer registers”. ARMv7-M has 16 integer registers, including special-purpose “stack pointer” and “program counter”.

Each element of x array needs to be “loaded” into a register.

Basic load instruction: LDR.

Manual says 2 cycles but adds a note about “pipelining”.

Then more explanation: if next instruction is also LDR (with address not based on first LDR) then it saves 1 cycle.

n consecutive LDRs takes only $n + 1$ cycles (“more multiple LDRs can be pipelined together”).

Can achieve this speed in other ways (LDRD, LDM) but nothing seems faster.

Lower bound for n LDR + n ADD: $2n + 1$ cycles, including n cycles of arithmetic.

Why observed time is higher: non-consecutive LDRs; costs of manipulating i .

and output of ADD are registers". ARMv7-M integer registers, including purpose "stack pointer" "program counter".

element of x array needs to be "loaded" into a register.

load instruction: LDR.

says 2 cycles but adds about "pipelining".

more explanation: if next instruction is also LDR (with pipeline not based on first LDR) saves 1 cycle.

n consecutive LDRs takes only $n + 1$ cycles ("more multiple LDRs can be pipelined together").

Can achieve this speed in other ways (LDRD, LDM) but nothing seems faster.

Lower bound for n LDR + n ADD: $2n + 1$ cycles, including n cycles of arithmetic.

Why observed time is higher: non-consecutive LDRs; costs of manipulating i .

```
int sum
{
    int r
    int *y
    int x0
        x5
    while
        x0 =
        x1 =
        x2 =
        x3 =
        x4 =
        x5 =
        x6 =
```


of ADD are
 . ARMv7-M
 isters, including
 tack pointer”
 nter”.

array needs to
 a register.

tion: LDR.

les but adds
 elining”.

ation: if next

LDR (with
 on first LDR)
 cle.

n consecutive LDRs
 takes only $n + 1$ cycles
 (“more multiple LDRs can be
 pipelined together”).

Can achieve this speed
 in other ways (LDRD, LDM)
 but nothing seems faster.

Lower bound for n LDR + n ADD:
 $2n + 1$ cycles,
 including n cycles of arithmetic.

Why observed time is higher:
 non-consecutive LDRs;
 costs of manipulating i .

```
int sum(int *x)
{
    int result = 0
    int *y = x + 1
    int x0,x1,x2,x
        x5,x6,x7,x

    while (x != y)
        x0 = 0[(vola
        x1 = 1[(vola
        x2 = 2[(vola
        x3 = 3[(vola
        x4 = 4[(vola
        x5 = 5[(vola
        x6 = 6[(vola
```

n consecutive LDRs
 takes only $n + 1$ cycles
 (“more multiple LDRs can be
 pipelined together”).

Can achieve this speed
 in other ways (LDRD, LDM)
 but nothing seems faster.

Lower bound for n LDR + n ADD:
 $2n + 1$ cycles,
 including n cycles of arithmetic.

Why observed time is higher:
 non-consecutive LDRs;
 costs of manipulating i .

```
int sum(int *x)
{
    int result = 0;
    int *y = x + 1000;
    int x0,x1,x2,x3,x4,
        x5,x6,x7,x8,x9;

    while (x != y) {
        x0 = 0[(volatile int
        x1 = 1[(volatile int
        x2 = 2[(volatile int
        x3 = 3[(volatile int
        x4 = 4[(volatile int
        x5 = 5[(volatile int
        x6 = 6[(volatile int
```

n consecutive LDRs

takes only $n + 1$ cycles

(“more multiple LDRs can be pipelined together”).

Can achieve this speed in other ways (LDRD, LDM) but nothing seems faster.

Lower bound for n LDR + n ADD:
 $2n + 1$ cycles,
including n cycles of arithmetic.

Why observed time is higher:
non-consecutive LDRs;
costs of manipulating i .

```
int sum(int *x)
{
    int result = 0;
    int *y = x + 1000;
    int x0,x1,x2,x3,x4,
        x5,x6,x7,x8,x9;

    while (x != y) {
        x0 = 0[(volatile int *)x];
        x1 = 1[(volatile int *)x];
        x2 = 2[(volatile int *)x];
        x3 = 3[(volatile int *)x];
        x4 = 4[(volatile int *)x];
        x5 = 5[(volatile int *)x];
        x6 = 6[(volatile int *)x];
    }
}
```

secutive LDRs

only $n + 1$ cycles

multiple LDRs can be
"bundled together").

to achieve this speed

different ways (LDRD, LDM)

streaming seems faster.

bound for n LDR + n ADD:

n cycles,

plus n cycles of arithmetic.

observed time is higher:

for consecutive LDRs;

and for manipulating i .

```
int sum(int *x)
```

```
{
```

```
    int result = 0;
```

```
    int *y = x + 1000;
```

```
    int x0,x1,x2,x3,x4,
```

```
        x5,x6,x7,x8,x9;
```

```
    while (x != y) {
```

```
        x0 = 0[(volatile int *)x];
```

```
        x1 = 1[(volatile int *)x];
```

```
        x2 = 2[(volatile int *)x];
```

```
        x3 = 3[(volatile int *)x];
```

```
        x4 = 4[(volatile int *)x];
```

```
        x5 = 5[(volatile int *)x];
```

```
        x6 = 6[(volatile int *)x];
```

x7 =

x8 =

x9 =

result

result

result

result

result

result

result

result

result

result

x0 =

x1 =

```

int sum(int *x)
{
    int result = 0;
    int *y = x + 1000;
    int x0,x1,x2,x3,x4,
        x5,x6,x7,x8,x9;

    while (x != y) {
        x0 = 0[(volatile int *)x];
        x1 = 1[(volatile int *)x];
        x2 = 2[(volatile int *)x];
        x3 = 3[(volatile int *)x];
        x4 = 4[(volatile int *)x];
        x5 = 5[(volatile int *)x];
        x6 = 6[(volatile int *)x];
    }
}

```

```

x7 = 7[(volatile int *)x];
x8 = 8[(volatile int *)x];
x9 = 9[(volatile int *)x];
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
result += x6;
result += x7;
result += x8;
result += x9;
x0 = 10[(volatile int *)x];
x1 = 11[(volatile int *)x];

```

```
int sum(int *x)
{
    int result = 0;
    int *y = x + 1000;
    int x0,x1,x2,x3,x4,
        x5,x6,x7,x8,x9;

    while (x != y) {
        x0 = 0[(volatile int *)x];
        x1 = 1[(volatile int *)x];
        x2 = 2[(volatile int *)x];
        x3 = 3[(volatile int *)x];
        x4 = 4[(volatile int *)x];
        x5 = 5[(volatile int *)x];
        x6 = 6[(volatile int *)x];
```

```
x7 = 7[(volatile int *)x];
x8 = 8[(volatile int *)x];
x9 = 9[(volatile int *)x];
    result += x0;
    result += x1;
    result += x2;
    result += x3;
    result += x4;
    result += x5;
    result += x6;
    result += x7;
    result += x8;
    result += x9;
    x0 = 10[(volatile int *)x];
    x1 = 11[(volatile int *)x];
```

```
int sum(int *x)
{
    int result = 0;
    int *y = x + 1000;
    int x0,x1,x2,x3,x4,
        x5,x6,x7,x8,x9;

    while (x != y) {
        x0 = 0[(volatile int *)x];
        x1 = 1[(volatile int *)x];
        x2 = 2[(volatile int *)x];
        x3 = 3[(volatile int *)x];
        x4 = 4[(volatile int *)x];
        x5 = 5[(volatile int *)x];
        x6 = 6[(volatile int *)x];
```

```
x7 = 7[(volatile int *)x];
x8 = 8[(volatile int *)x];
x9 = 9[(volatile int *)x];
    result += x0;
    result += x1;
    result += x2;
    result += x3;
    result += x4;
    result += x5;
    result += x6;
    result += x7;
    result += x8;
    result += x9;
    x0 = 10[(volatile int *)x];
    x1 = 11[(volatile int *)x];
```

```

(int *x)

result = 0;
y = x + 1000;
0, x1, x2, x3, x4,
5, x6, x7, x8, x9;

(x != y) {
= 0[(volatile int *)x];
= 1[(volatile int *)x];
= 2[(volatile int *)x];
= 3[(volatile int *)x];
= 4[(volatile int *)x];
= 5[(volatile int *)x];
= 6[(volatile int *)x];

```

```

x7 = 7[(volatile int *)x];
x8 = 8[(volatile int *)x];
x9 = 9[(volatile int *)x];
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
result += x6;
result += x7;
result += x8;
result += x9;
x0 = 10[(volatile int *)x];
x1 = 11[(volatile int *)x];

```

```

x2 =
x3 =
x4 =
x5 =
x6 =
x7 =
x8 =
x9 =
x +=
result
result
result
result
result
result

```



```

;
000;
3,x4,
8,x9;

{
tile int *)x];
tile int *)x];
tile int *)x];
tile int *)x];
tile int *)x];
tile int *)x];
tile int *)x];

```

```

x7 = 7[(volatile int *)x];
x8 = 8[(volatile int *)x];
x9 = 9[(volatile int *)x];
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
result += x6;
result += x7;
result += x8;
result += x9;
x0 = 10[(volatile int *)x];
x1 = 11[(volatile int *)x];

```

```

x2 = 12[(vol
x3 = 13[(vol
x4 = 14[(vol
x5 = 15[(vol
x6 = 16[(vol
x7 = 17[(vol
x8 = 18[(vol
x9 = 19[(vol
x += 20;
result += x0
result += x1
result += x2
result += x3
result += x4
result += x5

```

```
x7 = 7[(volatile int *)x];
x8 = 8[(volatile int *)x];
x9 = 9[(volatile int *)x];
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
result += x6;
result += x7;
result += x8;
result += x9;
x0 = 10[(volatile int *)x];
x1 = 11[(volatile int *)x];
```

```
x2 = 12[(volatile int *)x];
x3 = 13[(volatile int *)x];
x4 = 14[(volatile int *)x];
x5 = 15[(volatile int *)x];
x6 = 16[(volatile int *)x];
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];
x9 = 19[(volatile int *)x];
x += 20;
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
```

```
x7 = 7[(volatile int *)x];
x8 = 8[(volatile int *)x];
x9 = 9[(volatile int *)x];
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
result += x6;
result += x7;
result += x8;
result += x9;
x0 = 10[(volatile int *)x];
x1 = 11[(volatile int *)x];
```

```
x2 = 12[(volatile int *)x];
x3 = 13[(volatile int *)x];
x4 = 14[(volatile int *)x];
x5 = 15[(volatile int *)x];
x6 = 16[(volatile int *)x];
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];
x9 = 19[(volatile int *)x];
x += 20;
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
```

```
= 7[(volatile int *)x];
= 8[(volatile int *)x];
= 9[(volatile int *)x];

ult += x0;
ult += x1;
ult += x2;
ult += x3;
ult += x4;
ult += x5;
ult += x6;
ult += x7;
ult += x8;
ult += x9;

= 10[(volatile int *)x];
= 11[(volatile int *)x];
```

```
x2 = 12[(volatile int *)x];
x3 = 13[(volatile int *)x];
x4 = 14[(volatile int *)x];
x5 = 15[(volatile int *)x];
x6 = 16[(volatile int *)x];
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];
x9 = 19[(volatile int *)x];

x += 20;

result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;

result
result
result
result
}
return
}
```

39

```
tile int *)x];  
tile int *)x];  
tile int *)x];  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
atile int *)x];  
atile int *)x];
```

40

```
x2 = 12[(volatile int *)x];  
x3 = 13[(volatile int *)x];  
x4 = 14[(volatile int *)x];  
x5 = 15[(volatile int *)x];  
x6 = 16[(volatile int *)x];  
x7 = 17[(volatile int *)x];  
x8 = 18[(volatile int *)x];  
x9 = 19[(volatile int *)x];  
x += 20;  
result += x0;  
result += x1;  
result += x2;  
result += x3;  
result += x4;  
result += x5;  
result += x6;  
result += x7;  
result += x8;  
result += x9  
}  
return result;  
}
```

39

`*)x];``*)x];``*)x];``x2 = 12[(volatile int *)x];``x3 = 13[(volatile int *)x];``x4 = 14[(volatile int *)x];``x5 = 15[(volatile int *)x];``x6 = 16[(volatile int *)x];``x7 = 17[(volatile int *)x];``x8 = 18[(volatile int *)x];``x9 = 19[(volatile int *)x];``x += 20;``result += x0;``result += x1;``result += x2;``result += x3;``result += x4;``result += x5;``*)x];``*)x];`

40

`result += x6;``result += x7;``result += x8;``result += x9;``}``return result;``}`

```
x2 = 12[(volatile int *)x];
x3 = 13[(volatile int *)x];
x4 = 14[(volatile int *)x];
x5 = 15[(volatile int *)x];
x6 = 16[(volatile int *)x];
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];
x9 = 19[(volatile int *)x];
x += 20;
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
```

```
    result += x6;
    result += x7;
    result += x8;
    result += x9;
}

return result;
}
```

```
x2 = 12[(volatile int *)x];
x3 = 13[(volatile int *)x];
x4 = 14[(volatile int *)x];
x5 = 15[(volatile int *)x];
x6 = 16[(volatile int *)x];
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];
x9 = 19[(volatile int *)x];
x += 20;
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
```

```
    result += x6;
    result += x7;
    result += x8;
    result += x9;
}

return result;
}
```

2526 cycles. Even better in asm.


```
x2 = 12[(volatile int *)x];
x3 = 13[(volatile int *)x];
x4 = 14[(volatile int *)x];
x5 = 15[(volatile int *)x];
x6 = 16[(volatile int *)x];
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];
x9 = 19[(volatile int *)x];
x += 20;
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
```

```
    result += x6;
    result += x7;
    result += x8;
    result += x9;
}

return result;
}
```

2526 cycles. Even better in asm.

Wikipedia: “By the late 1990s for even performance sensitive code, optimizing compilers exceeded the performance of human experts.”

```

x2 = 12[(volatile int *)x];
x3 = 13[(volatile int *)x];
x4 = 14[(volatile int *)x];
x5 = 15[(volatile int *)x];
x6 = 16[(volatile int *)x];
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];
x9 = 19[(volatile int *)x];
x += 20;
result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;

```

```

    result += x6;
    result += x7;
    result += x8;
    result += x9;
}

return result;
}

```

2526 cycles. Even better in asm.

Wikipedia: “By the late 1990s for even performance sensitive code, optimizing compilers exceeded the performance of human experts.”

— [citation needed]