

Engineering
cryptographic software

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

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.

ring
raphic software

. Bernstein

ty of Illinois at Chicago &
che Universiteit Eindhoven

easy, right?

general principles

ftware engineering.

y principles to crypto.

y 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 softwa
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.

ago &
hoven

D.

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
- SAAAA
- SEAAA

Attacker
deduces
A few h
reveal se

2

the criteria
composing
ules” :
ead that
list of
isions or
hich are
Each module
to hide such
e 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. SECF
- SAAAAA vs. SECF
- SEAAAA vs. SECF

Attacker sees com
deduces position o
A few hundred trie
reveal secret passv

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. SECRET: stop at 6th character
- SAAAAA vs. SECRET: stop at 1st character
- SEAAAA vs. SECRET: stop at 2nd character

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

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. SECRET: stop at 1.
- SAAAAA vs. SECRET: stop at 2.
- SEAAAA vs. SECRET: 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
the 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. SECRET: stop at 1.
- SAAAAA vs. SECRET: stop at 2.
- SEAAAA vs. SECRET: 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
retur

Fix, elim
from sec
diff
for
di
retur

Notice t
makes th
and the

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. SECRET: stop at 1.
- SAAAAA vs. SECRET: stop at 2.
- SEAAAA vs. SECRET: stop at 3.

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

How typical software
 16-byte authentication

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

Fix, eliminating in
 from secrets to tim

```
diff = 0;
for (i = 0; i < 16; i++)
    diff |= x[i] != secret[i];
return 1 & (diff > 0);
```

Notice that the lar
 makes the wrong t
 and the right thing

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. SECRET: stop at 1.
- SAAAAA vs. SECRET: stop at 2.
- SEAAAA vs. SECRET: 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]) re
return 1;
```

Fix, eliminating information from secrets to timings:

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

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

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. SECRET: stop at 1.
- SAAAAA vs. SECRET: stop at 2.
- SEAAAA vs. SECRET: 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. SECRET: stop at 1.

A vs. SECRET: stop at 2.

A vs. SECRET: 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:

RET: stop at 1.

RET: stop at 2.

RET: 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;
return 1;
```

minating information flow

crets to timings:

```
= 0;
(i = 0; i < 16; ++i)
diff |= x[i] ^ y[i];
return 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];
diff-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.

attacks continue to happen.

In many examples,
the reference software for
a candidate CLOC:

```
are the tag */
```

```
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 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

Defender

Some of

[1996 Ko](#)

attacks

Briefly

Kocher

Schneier

secret

affect

[2002 Pa](#)

Suzaki-S

timing

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;  
  
CCESS;
```

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–Sasaki–Suzuki–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.
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.
Same issues described in 2004.

2016 Yarom–Genkin–Heninger

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

“steed” countermeasure:
 write table into cache.
 /2005.04 Bernstein:
 attacks on AES.
 countermeasure isn't safe;
 secret array indices can affect
 via cache-bank collisions.
 safe: kill all data flow
 secrets to array indices.
 Omer–Osvik–Shamir:
 steal Linux AES key
 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.
 Same issues described in 2004.

2016 Yarom–Genkin–Heninger

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

2008 RFC
 Layer Se
 Version
 small tim
 performa
 extent o
 fragmen
 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.

Shamir:
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.
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.
Same issues described in 2004.

2016 Yarom–Genkin–Heninger

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

affect

ons.

ow

.

t

y

n.

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.”

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.

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 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.
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 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.
 issues 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.
 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 document
 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 processed, 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
of instruction selection.

r documentation

ng variability: e.g.,

n operations terminate

e divide operation

es, with the number of

required 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.

Software

Almost a

much slo

Constant-time code

code in asm

ion selection.

tation

ity: e.g.,

ns terminate

operation

ne number of

pendent on the

t operands.”

ther than

umentation.

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.

Software optimization

Almost all software

much slower than

code

m

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.

Software optimization

Almost all software is much slower than it could b

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.

Software optimization

Almost all software is much slower than it could be.

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.

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.

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.

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.

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.

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

You want

software

as efficien

Starting

You hav

reference

You hav

(Can rep

You mea

impleme

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:

You want (constant) software that comes as efficiently as possible.

Starting point:

You have written a reference implementation.

You have chosen a (Can repeat for other)

You measure performance of 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:

You want (constant-time) software that computes ciphers as efficiently as possible.

Starting point:

You have written a reference implementation of

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:

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

Starting point:

You have written a reference implementation of X .

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

reference CPU time is negligible.

Opto software should be

central to all communication.

That's too slow \Rightarrow

fewer cryptanalysts

attractive for everybody.

Typical situation:

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

Starting point:

You have written a
reference implementation of X .

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

You measure performance of the
implementation. Now what?

A simplified

Target C

microcon

one ARM

Referenc

```
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:

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

Starting point:

You have written a
reference implementation of X .

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 LM3S9464
microcontroller code
one ARM Cortex-M3

Reference implementation

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

Typical situation:

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

Starting point:

You have written a reference implementation of X .

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:

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

Starting point:

You have written a reference implementation of X .

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:

at (constant-time)

that computes cipher X
as fast as possible.

point:

we written a

reference implementation of X .

(We have chosen a target CPU.
(We will try to beat for other CPUs.)

To measure performance of the
reference 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

*const

= (vo

...

int bef

int res

int aft

UARTpri

result

Output s

Change

15

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;
}
```

16

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

16

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.

17

“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.

17

“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.

18

Find “A
Technical
Rely on
M4F =
Manual
“implem
architect
Points to
Architec
which de
e.g., “A
First ma
ADD tal

```

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
 $M4F = M4 + \text{float}$
Manual says that
“implements the A
architecture profile
Points to the “AR
Architecture Refer
which defines instr
e.g., “ADD” for 3
First manual says
ADD takes just 1

“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 Proc
Technical Reference Manual
Rely on Wikipedia comment
M4F = M4 + floating-point

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

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

First manual says that
ADD takes just 1 cycle.

“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.

18

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.

19

Inputs a
"integer

has 16 i
special-p
and "pro

Each ele
be "load

Basic loa
Manual

a note a
Then mo

instructi
address
then it s

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 registers and special-purpose “s” and “program counter” .

Each element of x can be “loaded” into a register.

Basic load instruction takes 2 cycles. Manual says 2 cycles, with a note about “pipeline” . Then more explanation: “The load instruction is also address not based, so it saves 1 cycle” .

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- 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 add 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.

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 =
```

Rs
cycles
DRs can be
").

peed
RD, LDM)
s faster.

n LDR + n ADD:

of arithmetic.

e is higher:

DRs;

ing 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 = 7[(vola
x8 = 8[(vola
x9 = 9[(vola
result += x0
result += x1
result += x2
result += x3
result += x4
result += x5
result += x6
result += x7
result += x8
result += x9
x0 = 10[(vol
x1 = 11[(vol
```

```
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
x8 = 8[(volatile int
x9 = 9[(volatile int
    result += x0;
    result += x1;
    result += x2;
    result += x3;
    result += x4;
    result += x5;
    result += x6;
    result += x7;
    result += x8;
    result += x9;
    x0 = 10[(volatile int
    x1 = 11[(volatile int
```

```
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

```

```

;
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];      resu
x3 = 13[(volatile int *)x];      resu
x4 = 14[(volatile int *)x];      resu
x5 = 15[(volatile int *)x];      resu
x6 = 16[(volatile int *)x];      }
x7 = 17[(volatile int *)x];
x8 = 18[(volatile int *)x];      return
x9 = 19[(volatile int *)x];      }

x += 20;

result += x0;
result += x1;
result += x2;
result += x3;
result += x4;
result += x5;
```

23

```
tile int *)x];  
tile int *)x];  
tile int *)x];  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
;  
atile int *)x];  
atile int *)x];
```

24

```
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;  
}
```

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

```
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]

```

= 12[(volatile int *)x];
= 13[(volatile int *)x];
= 14[(volatile int *)x];
= 15[(volatile int *)x];
= 16[(volatile int *)x];
= 17[(volatile int *)x];
= 18[(volatile int *)x];
= 19[(volatile int *)x];
= 20;

ult += x0;
ult += x1;
ult += x2;
ult += x3;
ult += x4;
ult += 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]

A real ex

Salsa20
30.25 cy

Lower bo
64 bytes
21 · 16 1
20 · 16 1
so at lea

ARMv7-
includes
as part o
(Compile

```

atile int *)x];
atile int *)x];
atile int *)x];
atile int *)x];
atile int *)x];
atile int *)x];
atile int *)x];
atile int *)x];
atile int *)x];

```

```

    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]

A real example

Salsa20 reference
30.25 cycles/byte

Lower bound for a
64 bytes require
21 · 16 1-cycle AD
20 · 16 1-cycle XO
so at least 10.25 c

ARMv7-M instruct
includes free rotat
as part of XOR ins
(Compiler knows t

```

*)x];
    result += x6;
*)x];
    result += x7;
*)x];
    result += x8;
*)x];
    result += x9;
*)x];
}
*)x];
    return result;
*)x];
}

```

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]

A real example

Salsa20 reference software:
30.25 cycles/byte on this CPU

Lower bound for arithmetic:
64 bytes require
21 · 16 1-cycle ADDs,
20 · 16 1-cycle XORs,
so at least 10.25 cycles/byte

ARMv7-M instruction set
includes free rotation
as part of XOR instruction.
(Compiler knows this.)

```

    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]

A real example

Salsa20 reference software:
30.25 cycles/byte on this CPU.

Lower bound for arithmetic:

64 bytes require

21 · 16 1-cycle ADDs,

20 · 16 1-cycle XORs,

so at least 10.25 cycles/byte.

ARMv7-M instruction set

includes free rotation

as part of XOR instruction.

(Compiler knows this.)

```

ult += x6;
ult += x7;
ult += x8;
ult += x9;

```

```

n result;

```

cles. Even better in asm.

ia: “By the late 1990s for
performance sensitive code,
ng compilers exceeded the
ance of human experts.”
tion needed]

A real example

Salsa20 reference software:
30.25 cycles/byte on this CPU.

Lower bound for arithmetic:

64 bytes require

21 · 16 1-cycle ADDs,

20 · 16 1-cycle XORs,

so at least 10.25 cycles/byte.

ARMv7-M instruction set

includes free rotation

as part of XOR instruction.

(Compiler knows this.)

Detailed

several c

load_li

store_l

Can repl

(Compile

Then ob

18 cycles

plus 5 cy

Still far

A real example

Salsa20 reference software:
30.25 cycles/byte on this CPU.

Lower bound for arithmetic:

64 bytes require

21 · 16 1-cycle ADDs,

20 · 16 1-cycle XORs,

so at least 10.25 cycles/byte.

ARMv7-M instruction set

includes free rotation

as part of XOR instruction.

(Compiler knows this.)

Detailed benchmark

several cycles/byte

load_littleendian

store_littleendian

Can replace with L

(Compiler doesn't

Then observe 23 c

18 cycles/byte for

plus 5 cycles/byte

Still far above 10.2

better in asm.

the late 1990s for

sensitive code,

ers exceeded the

man experts.”

d]

A real example

Salsa20 reference software:
30.25 cycles/byte on this CPU.

Lower bound for arithmetic:
64 bytes require
21 · 16 1-cycle ADDs,
20 · 16 1-cycle XORs,
so at least 10.25 cycles/byte.

ARMv7-M instruction set
includes free rotation
as part of XOR instruction.
(Compiler knows this.)

Detailed benchmarks show
several cycles/byte spent on
load_littleendian and
store_littleendian.

Can replace with LDR and STR
(Compiler doesn't see this.)

Then observe 23 cycles/byte
18 cycles/byte for rounds,
plus 5 cycles/byte overhead.
Still far above 10.25 cycles/

asm.

00s for
code,
ed the
rts.”

A real example

Salsa20 reference software:
30.25 cycles/byte on this CPU.

Lower bound for arithmetic:

64 bytes require

21 · 16 1-cycle ADDs,

20 · 16 1-cycle XORs,

so at least 10.25 cycles/byte.

ARMv7-M instruction set

includes free rotation

as part of XOR instruction.

(Compiler knows this.)

Detailed benchmarks show
several cycles/byte spent on
load_littleendian and
store_littleendian.

Can replace with LDR and STR.
(Compiler doesn't see this.)

Then observe 23 cycles/byte:
18 cycles/byte for rounds,
plus 5 cycles/byte overhead.
Still far above 10.25 cycles/byte.

A real example

Salsa20 reference software:
30.25 cycles/byte on this CPU.

Lower bound for arithmetic:

64 bytes require

21 · 16 1-cycle ADDs,

20 · 16 1-cycle XORs,

so at least 10.25 cycles/byte.

ARMv7-M instruction set

includes free rotation

as part of XOR instruction.

(Compiler knows this.)

Detailed benchmarks show
several cycles/byte spent on
load_littleendian and
store_littleendian.

Can replace with LDR and STR.
(Compiler doesn't see this.)

Then observe 23 cycles/byte:
18 cycles/byte for rounds,
plus 5 cycles/byte overhead.
Still far above 10.25 cycles/byte.

Gap is mostly loads, stores.
Minimize load/store cost by
choosing “spills” carefully.

example

reference software:

cycles/byte on this CPU.

bound for arithmetic:

require

1-cycle ADDs,

1-cycle XORs,

at least 10.25 cycles/byte.

M instruction set

free rotation

of XOR instruction.

(Compiler knows this.)

Detailed benchmarks show
several cycles/byte spent on
`load_littleendian` and
`store_littleendian`.

Can replace with `LDR` and `STR`.
(Compiler doesn't see this.)

Then observe 23 cycles/byte:
18 cycles/byte for rounds,
plus 5 cycles/byte overhead.
Still far above 10.25 cycles/byte.

Gap is mostly loads, stores.
Minimize load/store cost by
choosing "spills" carefully.

Which o
should b
Don't tr
optimize

Make lo
Don't tr
optimize

Spill to
Don't tr
optimize

On bigg
selecting
is critica

software:
on this CPU.

arithmetic:

Ds,

Rs,

cycles/byte.

tion set

ion

struction.

this.)

Detailed benchmarks show
several cycles/byte spent on
`load_littleendian` and
`store_littleendian`.

Can replace with LDR and STR.
(Compiler doesn't see this.)

Then observe 23 cycles/byte:
18 cycles/byte for rounds,
plus 5 cycles/byte overhead.
Still far above 10.25 cycles/byte.

Gap is mostly loads, stores.
Minimize load/store cost by
choosing "spills" carefully.

Which of the 16 S
should be in regist
Don't trust compi
optimize register a

Make loads consec
Don't trust compi
optimize instructio

Spill to FPU instea
Don't trust compi
optimize instructio

On bigger CPUs,
selecting vector ins
is critical for perfo

Detailed benchmarks show several cycles/byte spent on `load_littleendian` and `store_littleendian`.

Can replace with `LDR` and `STR`.
(Compiler doesn't see this.)

Then observe 23 cycles/byte:
18 cycles/byte for rounds,
plus 5 cycles/byte overhead.
Still far above 10.25 cycles/byte.

Gap is mostly loads, stores.
Minimize load/store cost by
choosing "spills" carefully.

Which of the 16 Salsa20 words should be in registers?
Don't trust compiler to optimize register allocation.

Make loads consecutive?
Don't trust compiler to optimize instruction schedule.

Spill to FPU instead of stack?
Don't trust compiler to optimize instruction selection.

On bigger CPUs,
selecting vector instructions is critical for performance.

Detailed benchmarks show several cycles/byte spent on `load_littleendian` and `store_littleendian`.

Can replace with `LDR` and `STR`.
(Compiler doesn't see this.)

Then observe 23 cycles/byte:
18 cycles/byte for rounds,
plus 5 cycles/byte overhead.
Still far above 10.25 cycles/byte.

Gap is mostly loads, stores.
Minimize load/store cost by
choosing "spills" carefully.

Which of the 16 Salsa20 words should be in registers?

Don't trust compiler to optimize register allocation.

Make loads consecutive?

Don't trust compiler to optimize instruction scheduling.

Spill to FPU instead of stack?

Don't trust compiler to optimize instruction selection.

On bigger CPUs,
selecting vector instructions
is critical for performance.

benchmarks show
cycles/byte spent on
littleendian and
littleendian.

ace with LDR and STR.
(compiler doesn't see this.)

serve 23 cycles/byte:

s/byte for rounds,
cycles/byte overhead.

above 10.25 cycles/byte.

mostly loads, stores.

the load/store cost by

g "spills" carefully.

Which of the 16 Salsa20 words
should be in registers?

Don't trust compiler to
optimize register allocation.

Make loads consecutive?

Don't trust compiler to
optimize instruction scheduling.

Spill to FPU instead of stack?

Don't trust compiler to
optimize instruction selection.

On bigger CPUs,
selecting vector instructions
is critical for performance.

The big

CPUs are
farther a
from na

arks show
e spent on
n and
an.

LDR and STR.
(see this.)

cycles/byte:

rounds,
overhead.

25 cycles/byte.

ls, stores.
re cost by
carefully.

Which of the 16 Salsa20 words
should be in registers?

Don't trust compiler to
optimize register allocation.

Make loads consecutive?

Don't trust compiler to
optimize instruction scheduling.

Spill to FPU instead of stack?

Don't trust compiler to
optimize instruction selection.

On bigger CPUs,
selecting vector instructions
is critical for performance.

The big picture

CPUs are evolving
farther and farther
from naive models

Which of the 16 Salsa20 words
should be in registers?

Don't trust compiler to
optimize register allocation.

Make loads consecutive?

Don't trust compiler to
optimize instruction scheduling.

Spill to FPU instead of stack?

Don't trust compiler to
optimize instruction selection.

On bigger CPUs,
selecting vector instructions
is critical for performance.

The big picture

CPUs are evolving
farther and farther away
from naive models of CPUs.

Which of the 16 Salsa20 words should be in registers?

Don't trust compiler to optimize register allocation.

Make loads consecutive?

Don't trust compiler to optimize instruction scheduling.

Spill to FPU instead of stack?

Don't trust compiler to optimize instruction selection.

On bigger CPUs, selecting vector instructions is critical for performance.

The big picture

CPUs are evolving farther and farther away from naive models of CPUs.

Which of the 16 Salsa20 words should be in registers?

Don't trust compiler to optimize register allocation.

Make loads consecutive?

Don't trust compiler to optimize instruction scheduling.

Spill to FPU instead of stack?

Don't trust compiler to optimize instruction selection.

On bigger CPUs, selecting vector instructions is critical for performance.

The big picture

CPUs are evolving farther and farther away from naive models of CPUs.

Minor optimization challenges:

- Pipelining.
- Superscalar processing.

Major optimization challenges:

- Vectorization.
- Many threads; many cores.
- The memory hierarchy; the ring; the mesh.
- Larger-scale parallelism.
- Larger-scale networking.

of the 16 Salsa20 words
 be in registers?
 must compiler to
 e register allocation.
 ads consecutive?
 must compiler to
 e instruction scheduling.
 FPU instead of stack?
 must compiler to
 e instruction selection.
 er CPUs,
 g vector instructions
 l for performance.

The big picture

CPUs are evolving
 farther and farther away
 from naive models of CPUs.

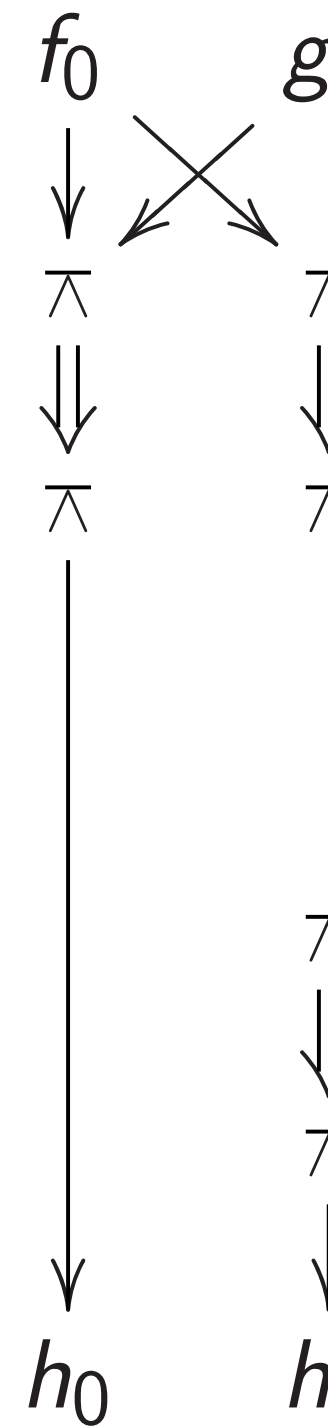
Minor optimization challenges:

- Pipelining.
- Superscalar processing.

Major optimization challenges:

- Vectorization.
- Many threads; many cores.
- The memory hierarchy;
the ring; the mesh.
- Larger-scale parallelism.
- Larger-scale networking.

CPU des



Gates π
 product
 of integers

also 20 words
 registers?
 order to
 allocation.
 cutive?
 order to
 on scheduling.
 ad of stack?
 order to
 on selection.
 instructions
 performance.

The big picture

CPUs are evolving
 farther and farther away
 from naive models of CPUs.

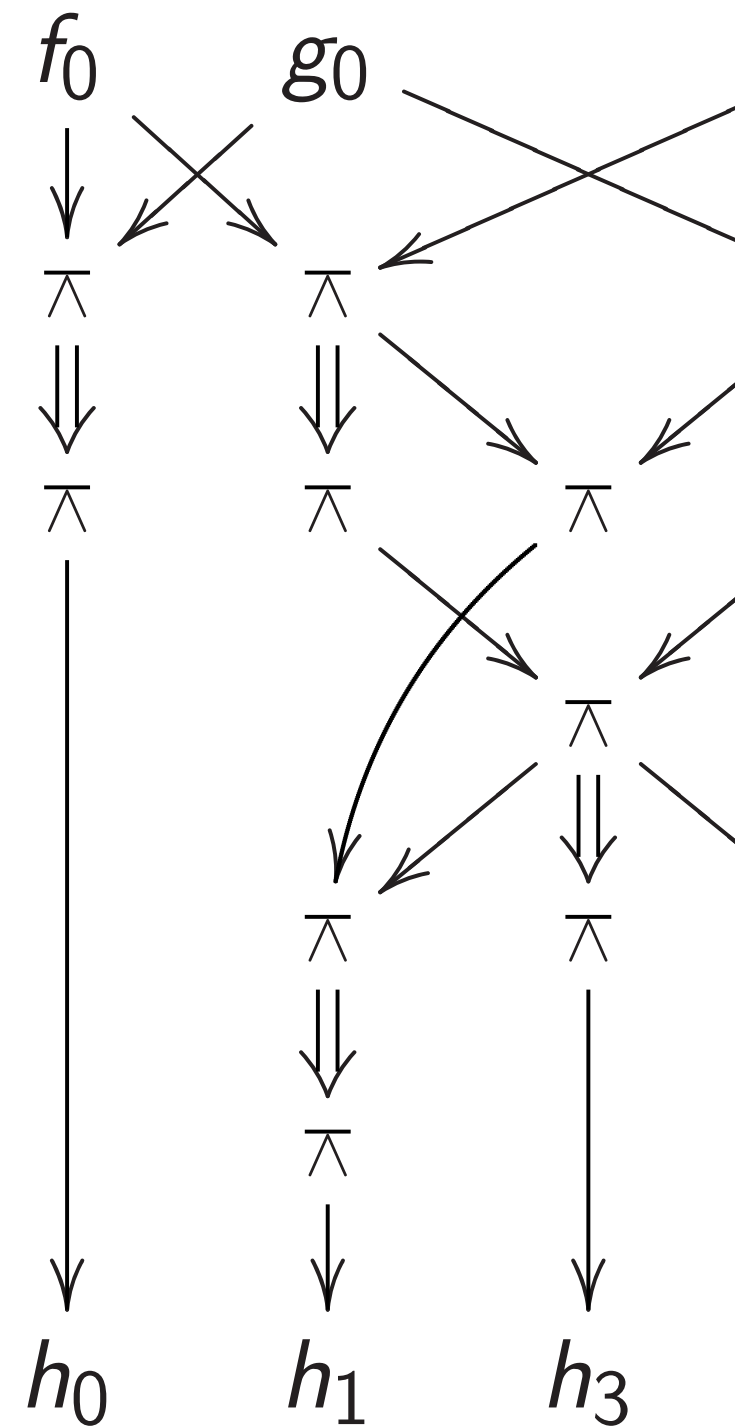
Minor optimization challenges:

- Pipelining.
- Superscalar processing.

Major optimization challenges:

- Vectorization.
- Many threads; many cores.
- The memory hierarchy;
the ring; the mesh.
- Larger-scale parallelism.
- Larger-scale networking.

CPU design in a n



Gates $\pi: a, b \mapsto 1$
 product $h_0 + 2h_1 + h_3$
 of integers $f_0 + 2f_1 + f_3$

The big picture

CPUs are evolving farther and farther away from naive models of CPUs.

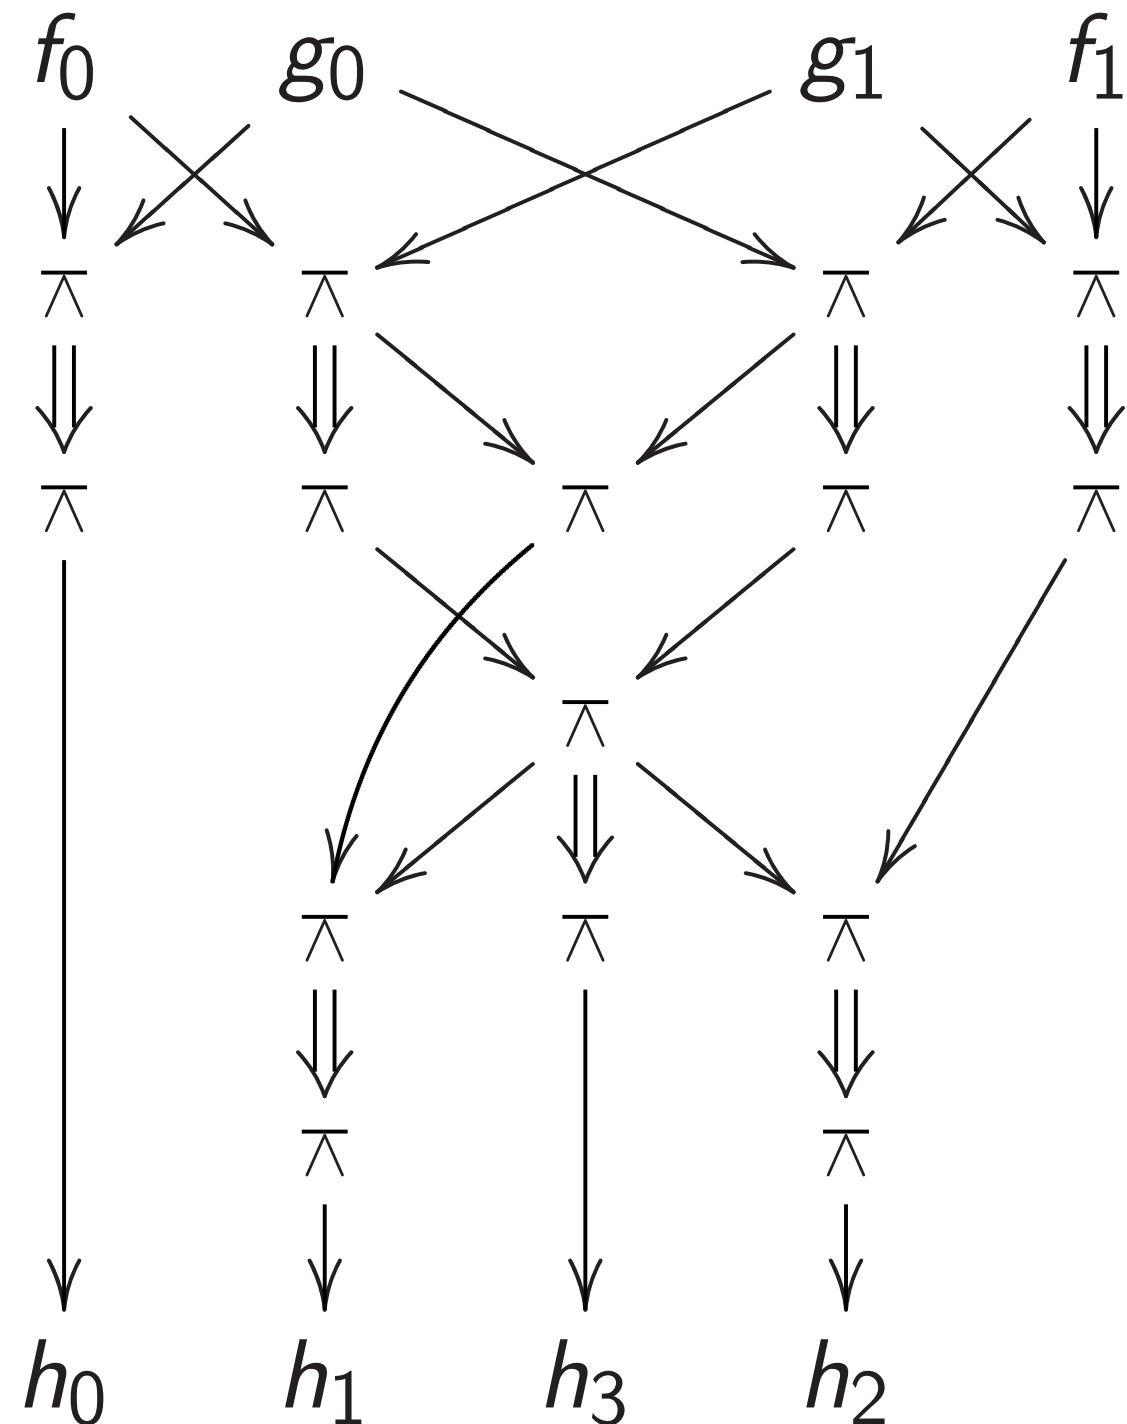
Minor optimization challenges:

- Pipelining.
- Superscalar processing.

Major optimization challenges:

- Vectorization.
- Many threads; many cores.
- The memory hierarchy; the ring; the mesh.
- Larger-scale parallelism.
- Larger-scale networking.

CPU design in a nutshell



Gates $\pi : a, b \mapsto 1 - ab$ compute the product $h_0 + 2h_1 + 4h_2 + 8h_3$ of integers $f_0 + 2f_1, g_0 + 2g_1$.

The big picture

CPUs are evolving farther and farther away from naive models of CPUs.

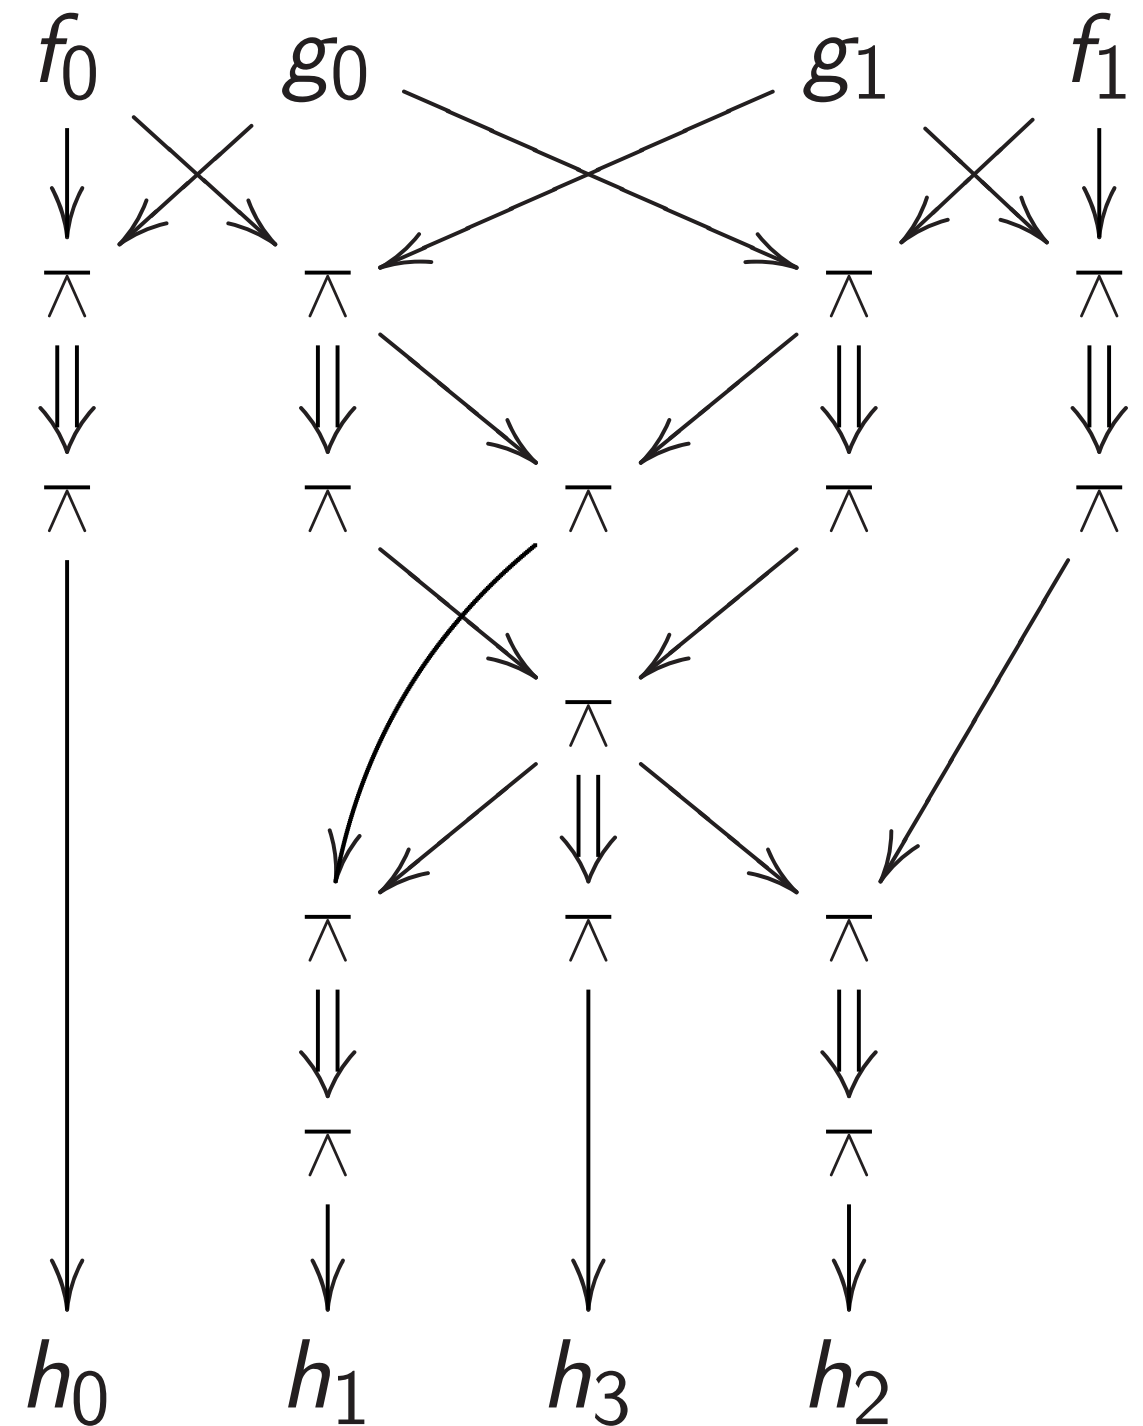
Minor optimization challenges:

- Pipelining.
- Superscalar processing.

Major optimization challenges:

- Vectorization.
- Many threads; many cores.
- The memory hierarchy; the ring; the mesh.
- Larger-scale parallelism.
- Larger-scale networking.

CPU design in a nutshell



Gates $\pi : a, b \mapsto 1 - ab$ computing product $h_0 + 2h_1 + 4h_2 + 8h_3$ of integers $f_0 + 2f_1, g_0 + 2g_1$.

picture

evolving
and farther away
ve models of CPUs.

optimization challenges:
ning.

scalar processing.

optimization challenges:
rization.

threads; many cores.

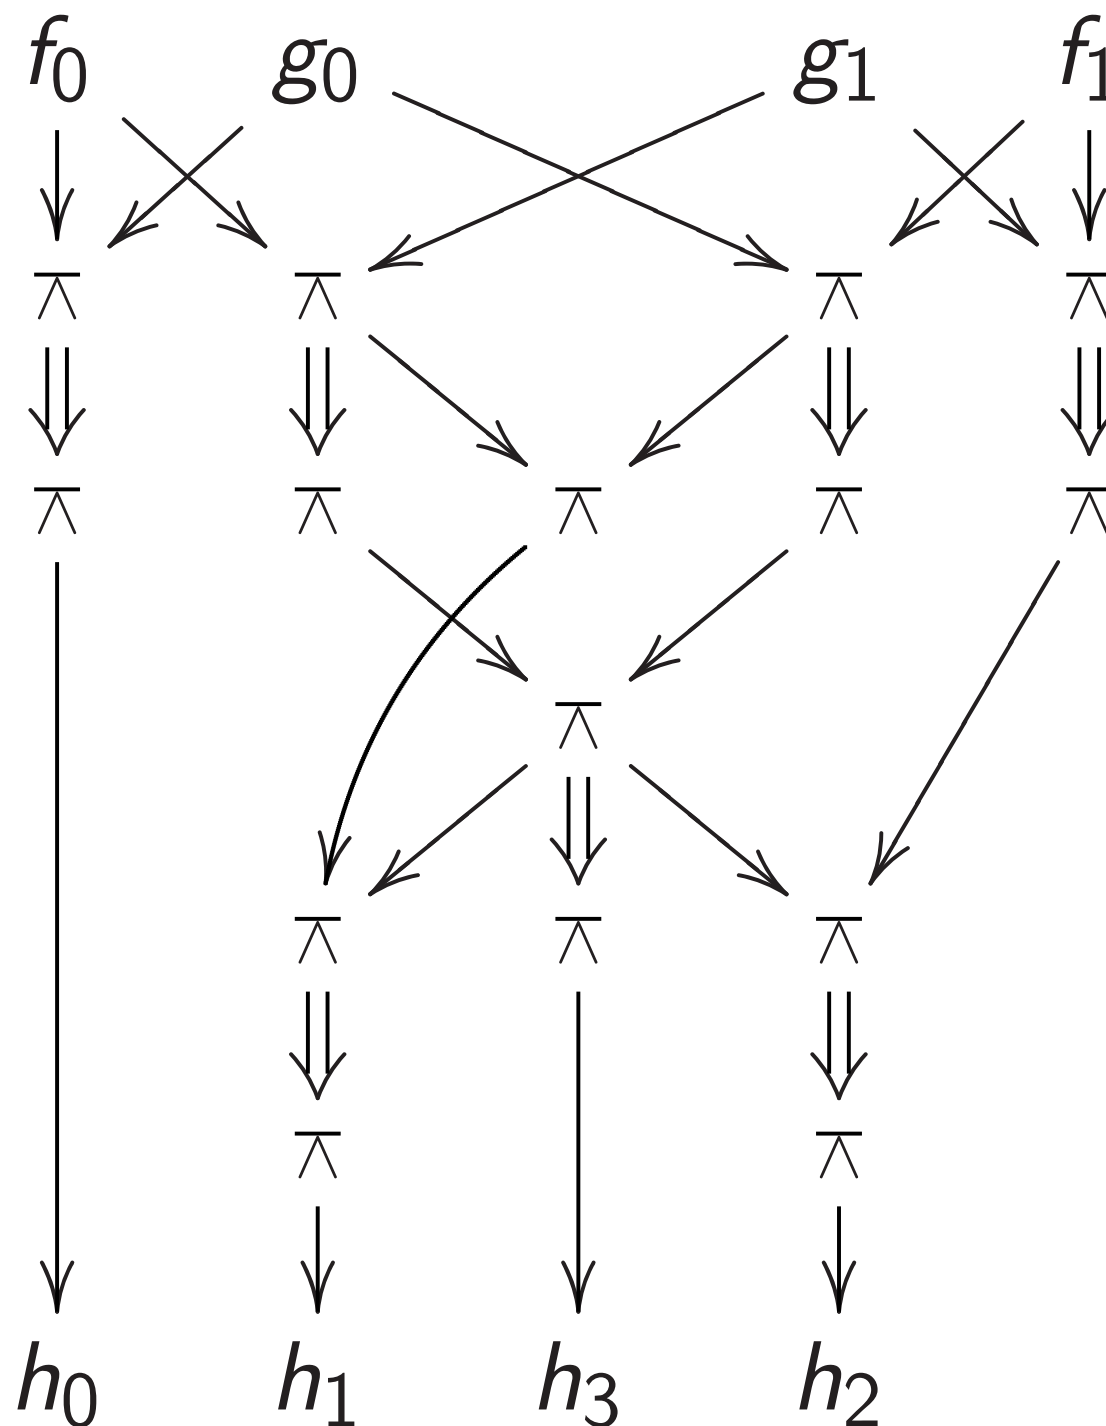
memory hierarchy;

g; the mesh.

-scale parallelism.

-scale networking.

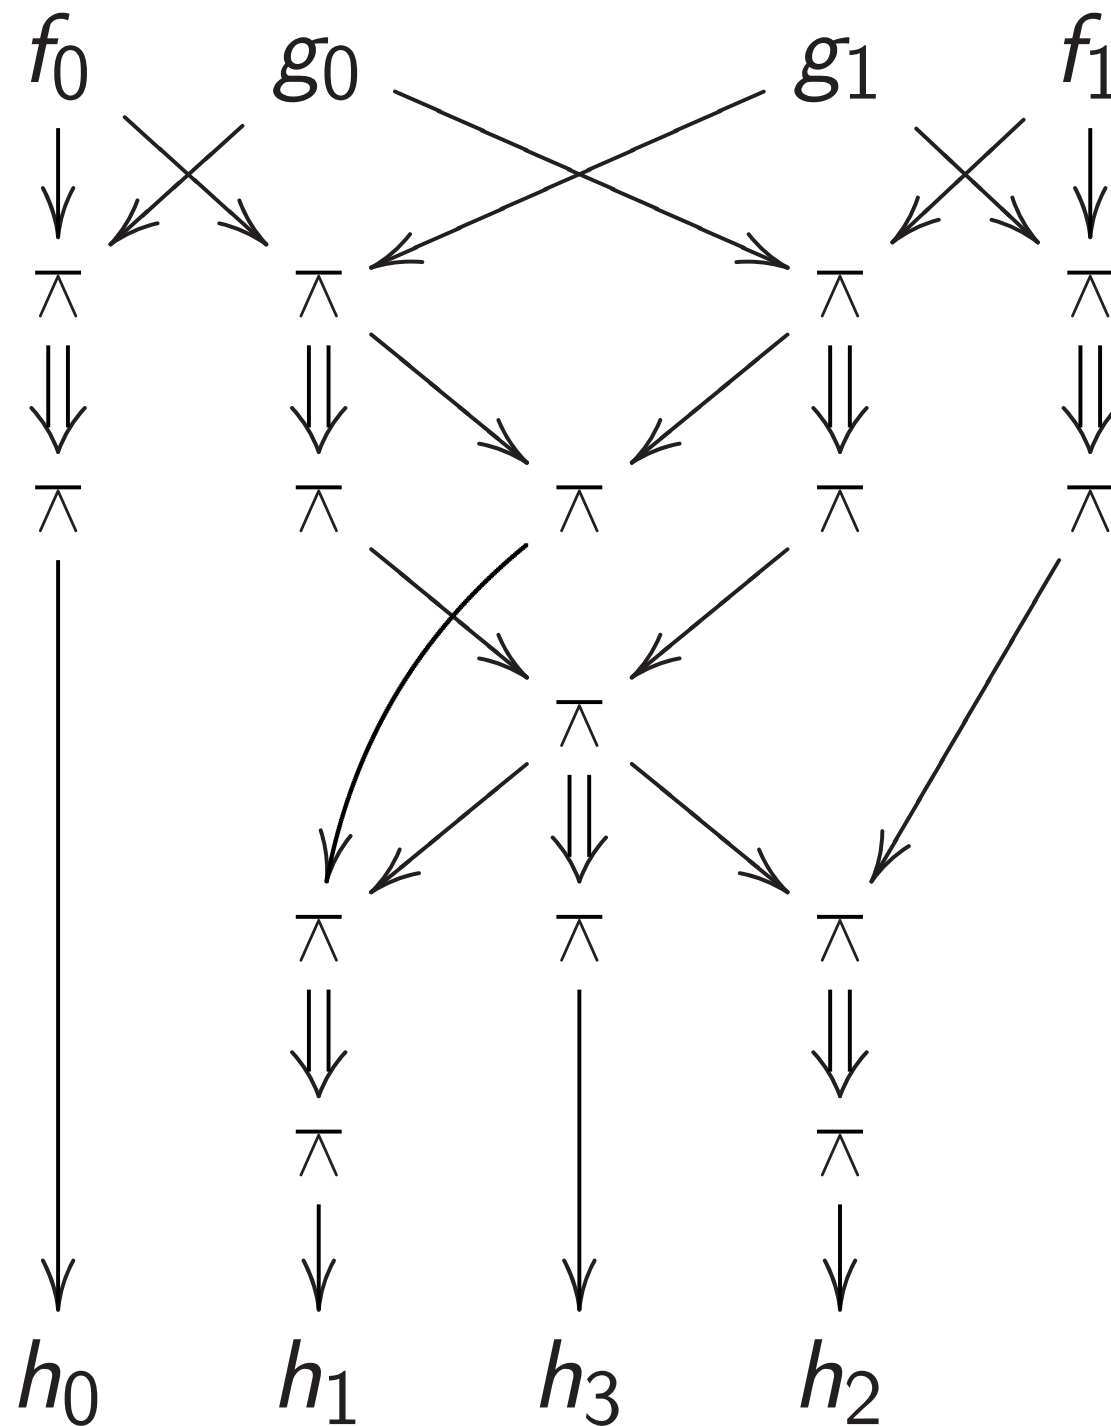
CPU design in a nutshell



Gates $\pi : a, b \mapsto 1 - ab$ computing
product $h_0 + 2h_1 + 4h_2 + 8h_3$
of integers $f_0 + 2f_1, g_0 + 2g_1$.

Electricity
percolates
If f_0, f_1, \dots
then h_0, \dots
a few m

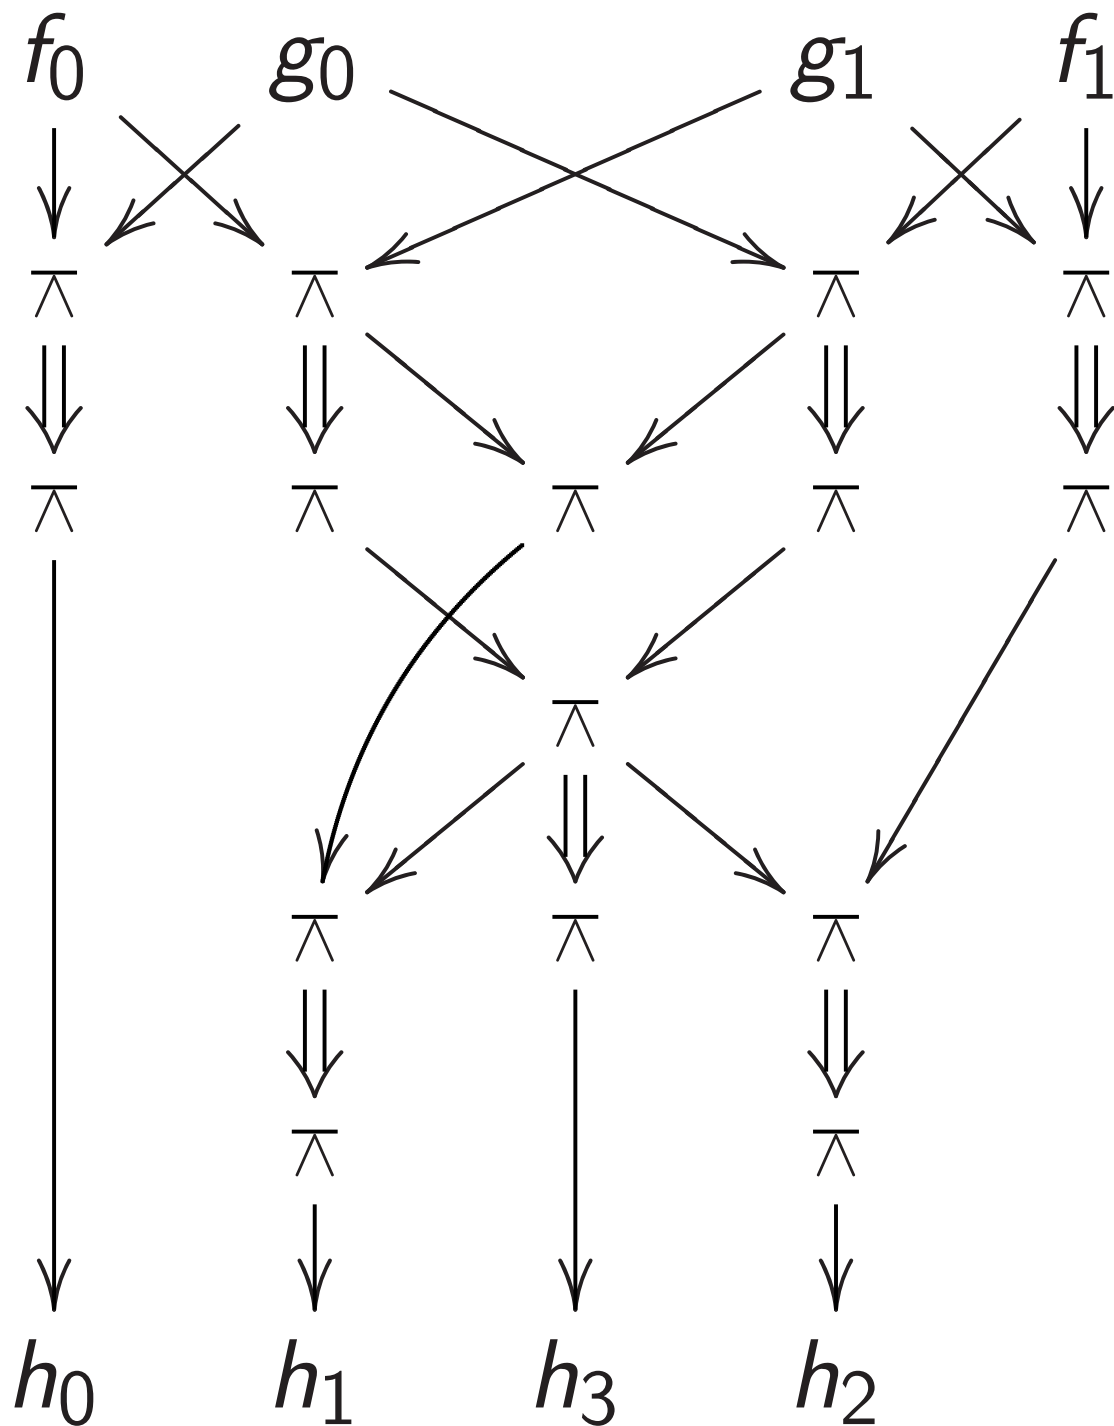
CPU design in a nutshell



Gates $\neg : a, b \mapsto 1 - ab$ computing
 product $h_0 + 2h_1 + 4h_2 + 8h_3$
 of integers $f_0 + 2f_1, g_0 + 2g_1$.

Electricity takes time
 to percolate through
 the circuit.
 If f_0, f_1, g_0, g_1 are
 integers, then h_0, h_1, h_2, h_3
 are integers.
 a few moments later

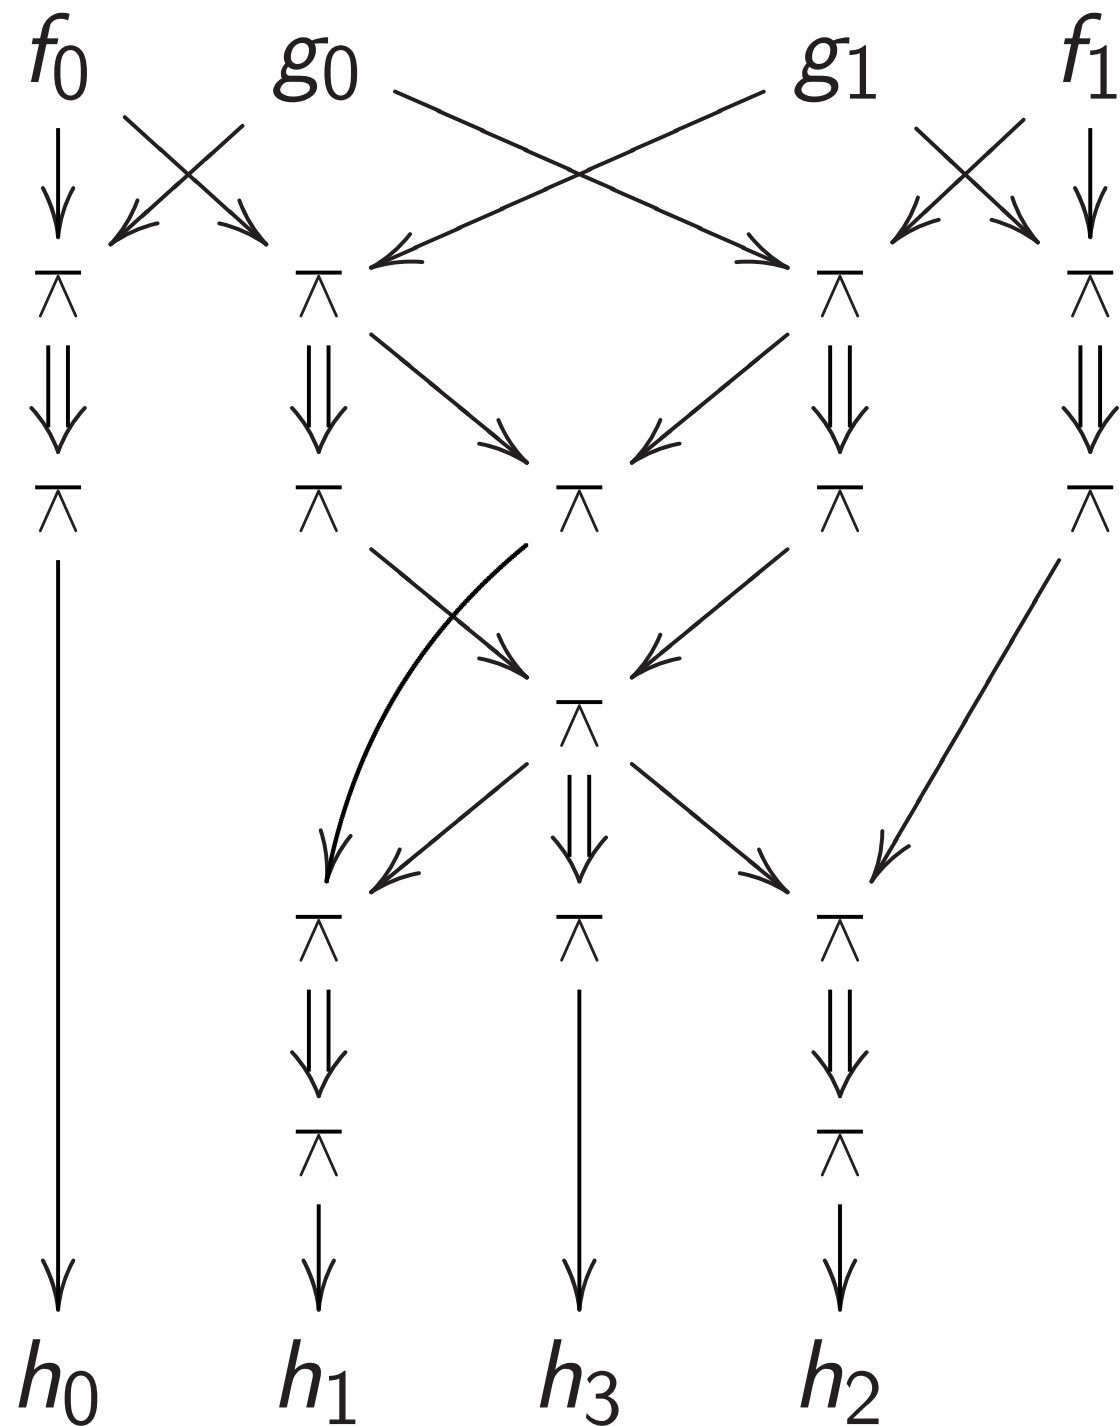
CPU design in a nutshell



Gates \neg : $a, b \mapsto 1 - ab$ computing
 product $h_0 + 2h_1 + 4h_2 + 8h_3$
 of integers $f_0 + 2f_1, g_0 + 2g_1$.

Electricity takes time to
 percolate through wires and
 If f_0, f_1, g_0, g_1 are stable
 then h_0, h_1, h_2, h_3 are stable
 a few moments later.

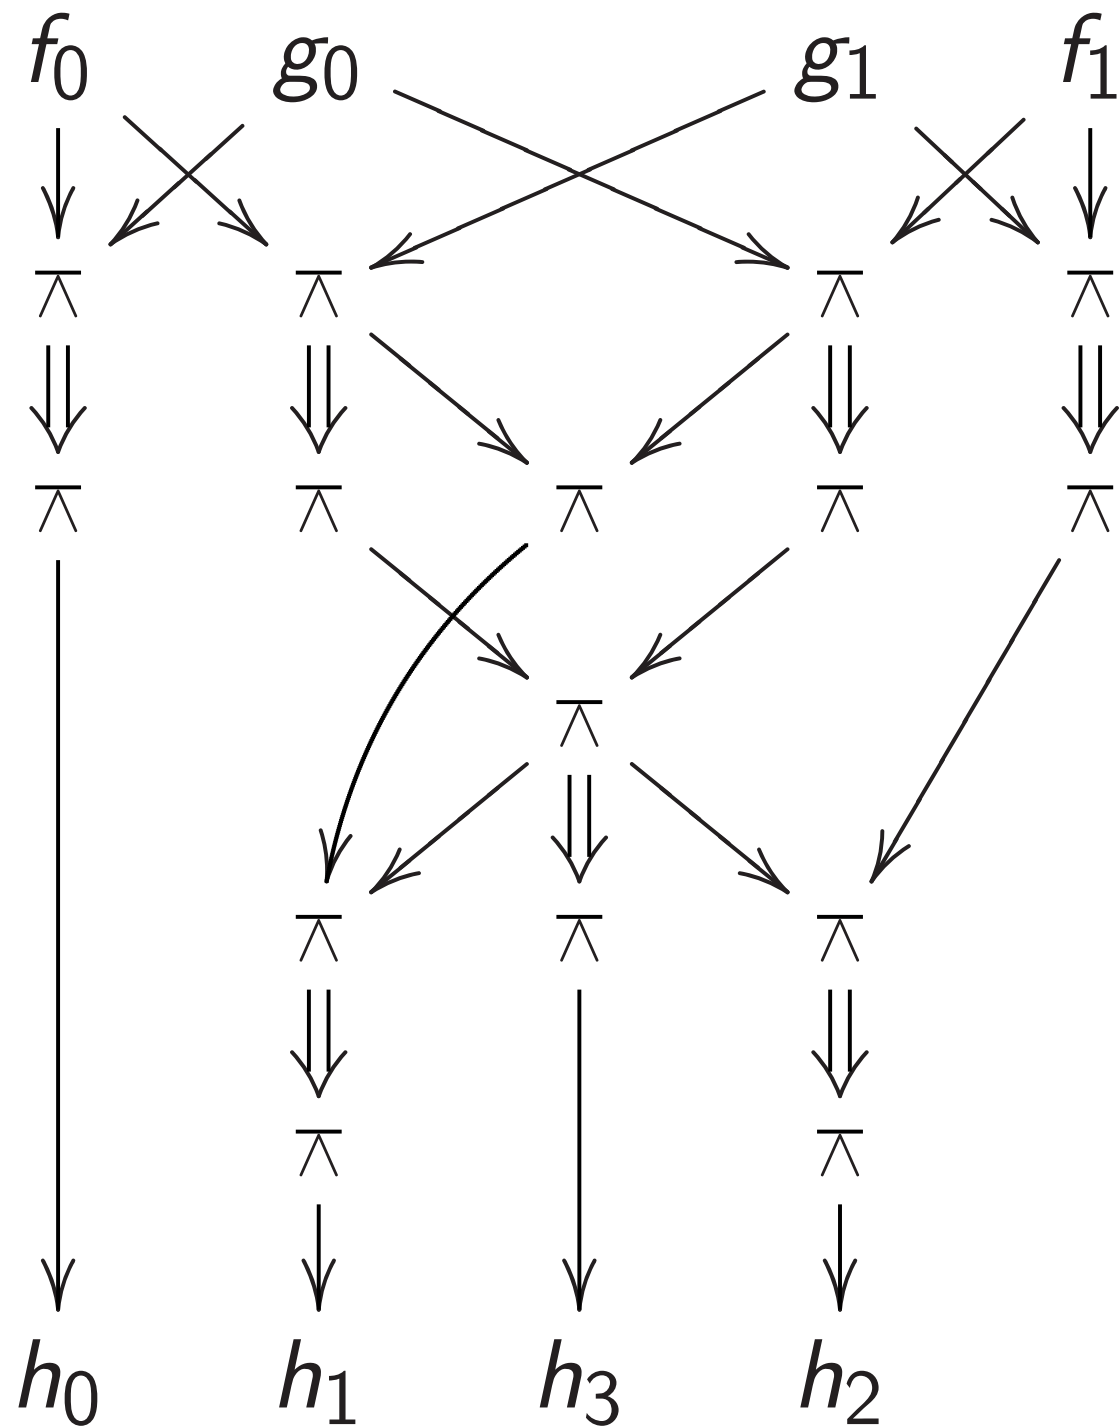
CPU design in a nutshell



Gates $\neg : a, b \mapsto 1 - ab$ computing product $h_0 + 2h_1 + 4h_2 + 8h_3$ of integers $f_0 + 2f_1, g_0 + 2g_1$.

Electricity takes time to percolate through wires and gates. If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

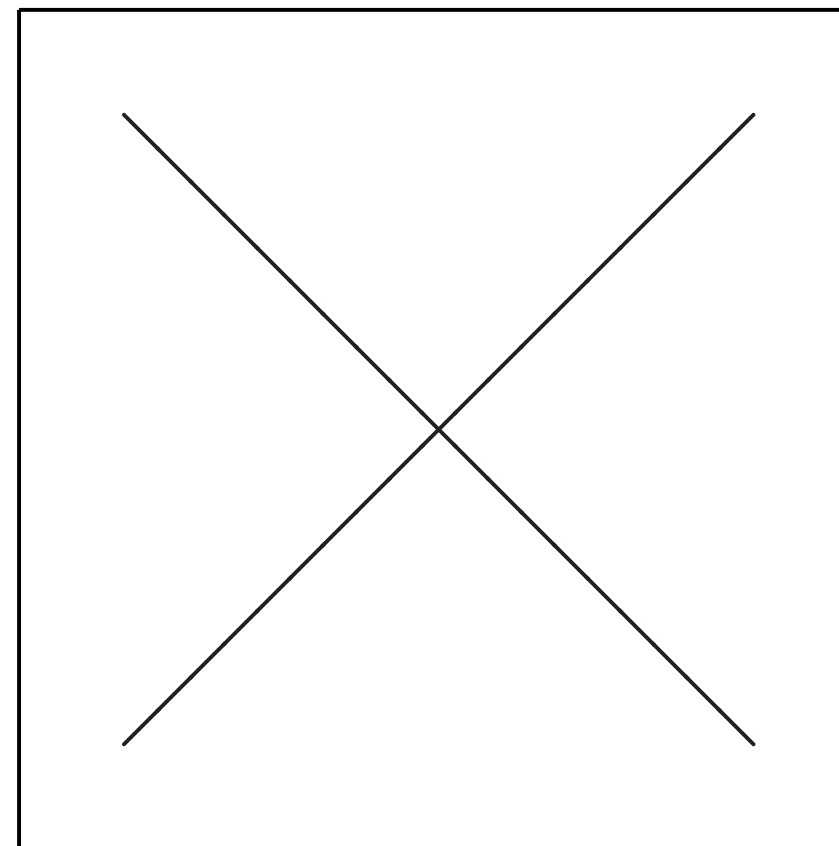
CPU design in a nutshell



Gates $\wedge : a, b \mapsto 1 - ab$ computing product $h_0 + 2h_1 + 4h_2 + 8h_3$ of integers $f_0 + 2f_1, g_0 + 2g_1$.

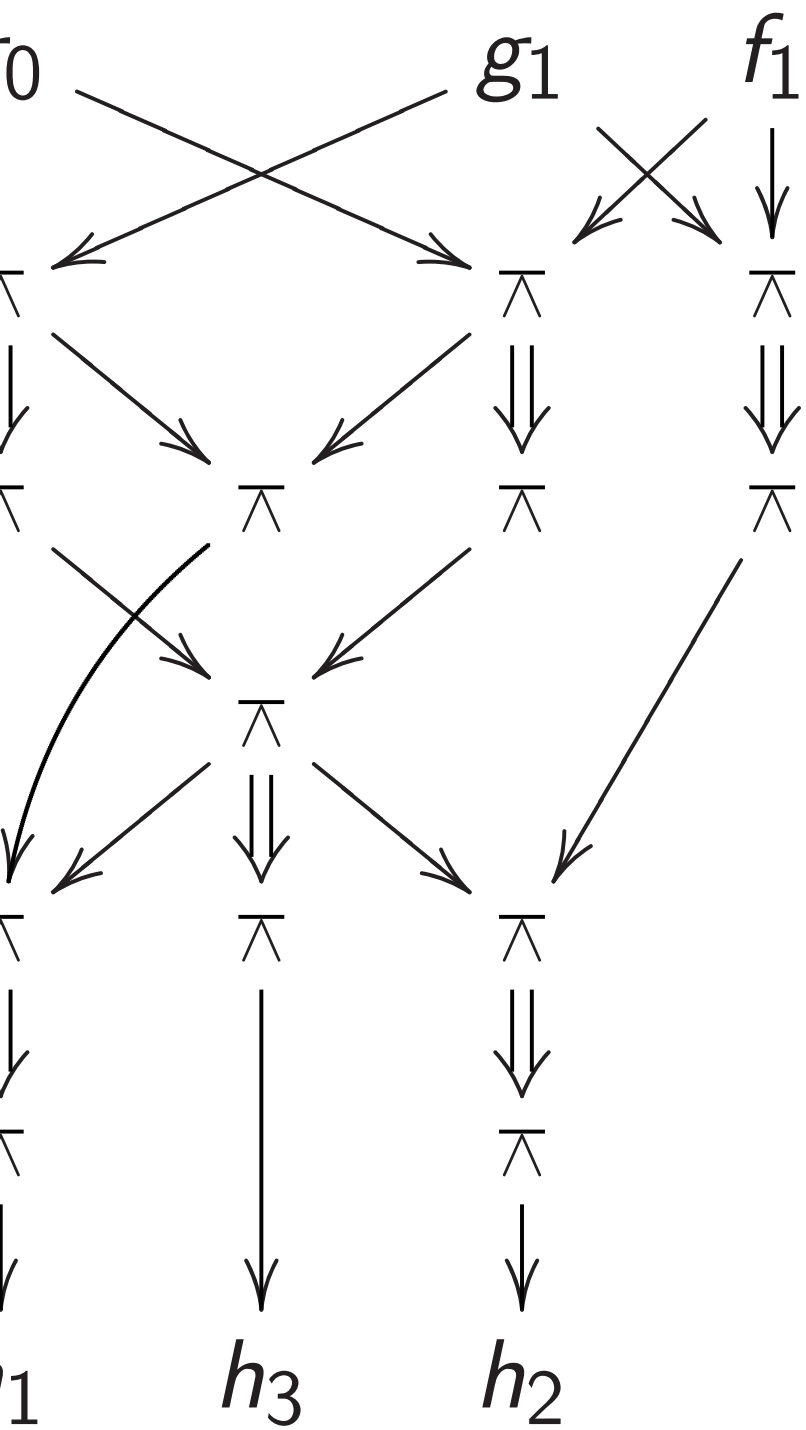
Electricity takes time to percolate through wires and gates. If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

Build circuit with more gates to multiply (e.g.) 32-bit integers:



(Details omitted.)

Design in a nutshell



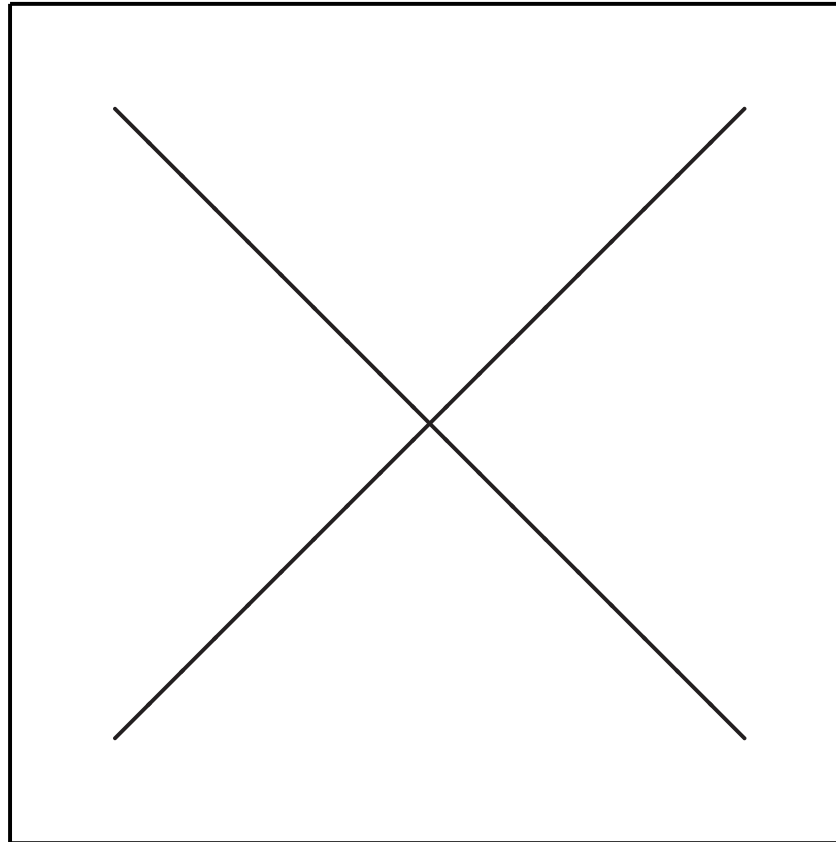
$: a, b \mapsto 1 - ab$ computing

$h_0 + 2h_1 + 4h_2 + 8h_3$

ers $f_0 + 2f_1, g_0 + 2g_1$.

Electricity takes time to percolate through wires and gates. If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

Build circuit with more gates to multiply (e.g.) 32-bit integers:

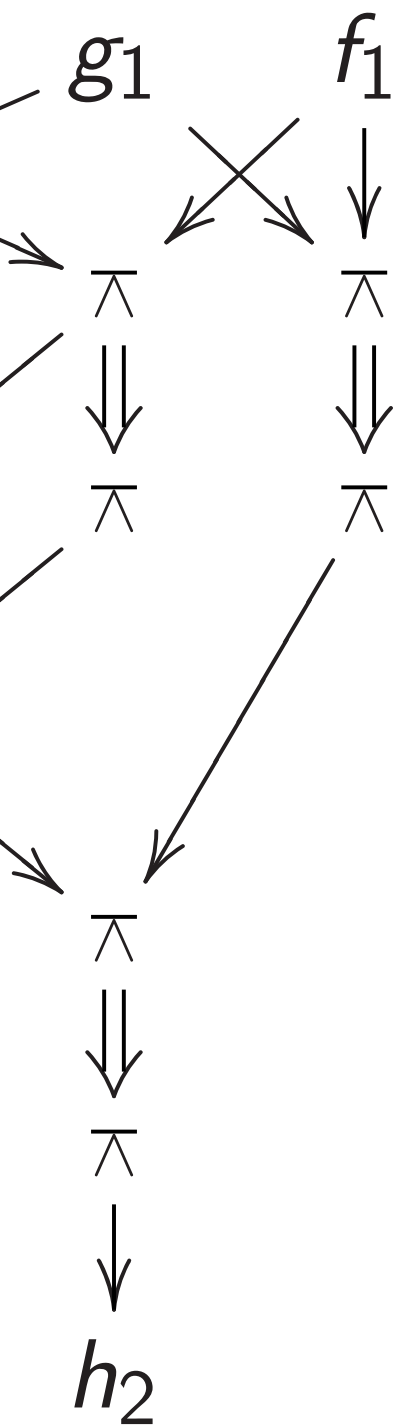


(Details omitted.)

Build circuit
32-bit integers
given 4-bit
and 32-bit

reg
re

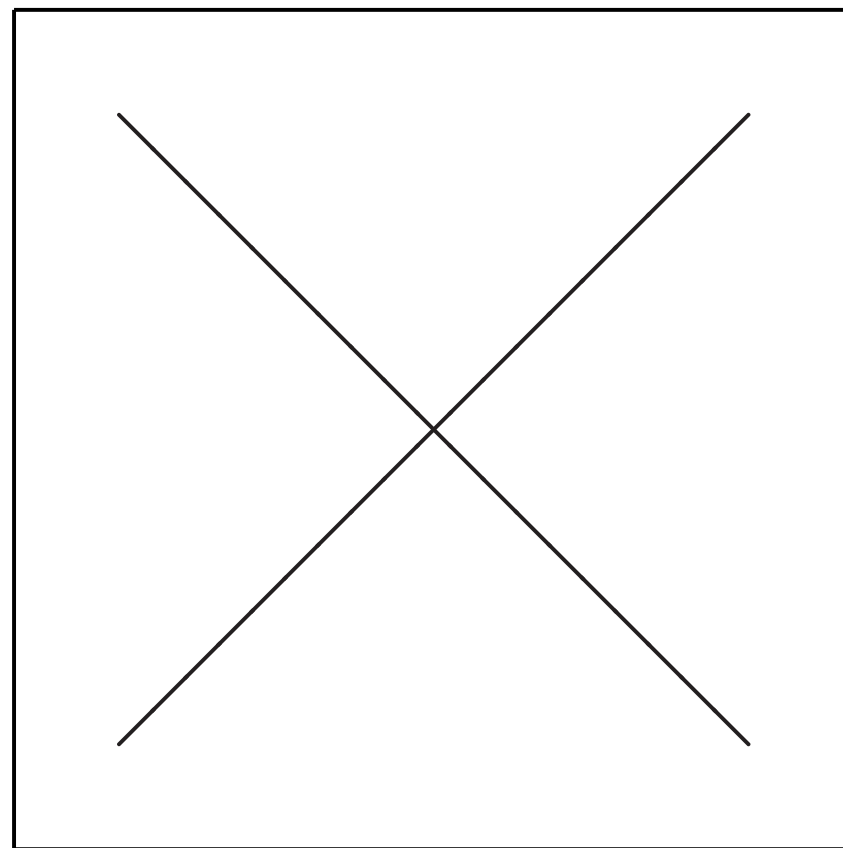
utshell



– ab computing
 $+ 4h_2 + 8h_3$
 $+ f_1, g_0 + 2g_1.$

Electricity takes time to percolate through wires and gates. If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

Build circuit with more gates to multiply (e.g.) 32-bit integers:



(Details omitted.)

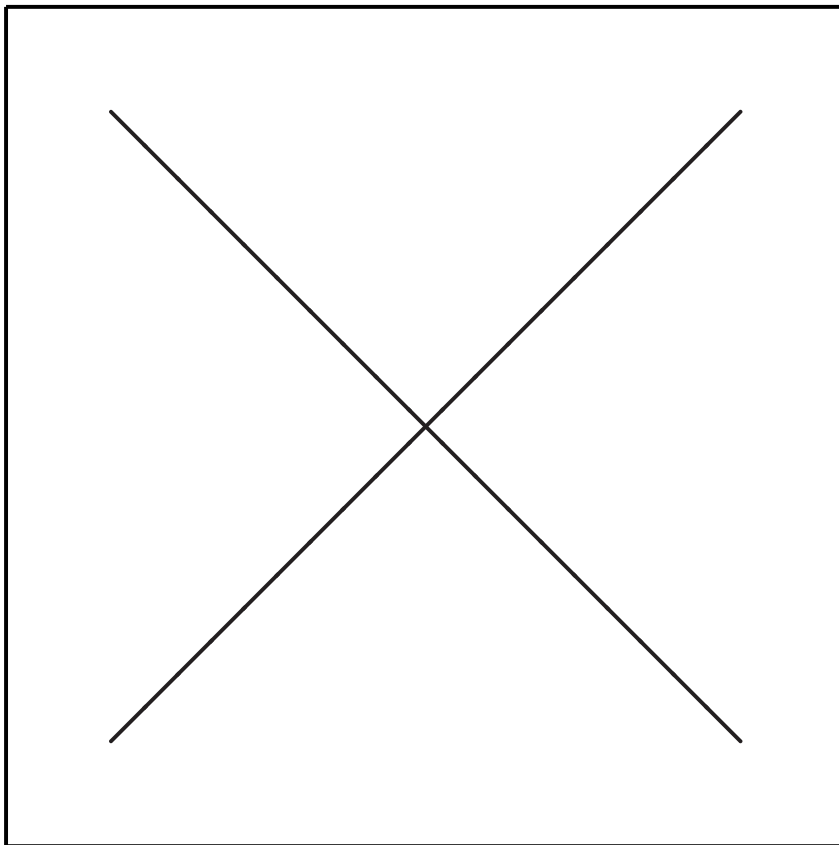
Build circuit to compute 32-bit integer r_i given 4-bit integers a_i and 32-bit integers r_i .

register
read

Electricity takes time to percolate through wires and gates.

If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

Build circuit with more gates to multiply (e.g.) 32-bit integers:



(Details omitted.)

Build circuit to compute 32-bit integer r_i

given 4-bit integer i

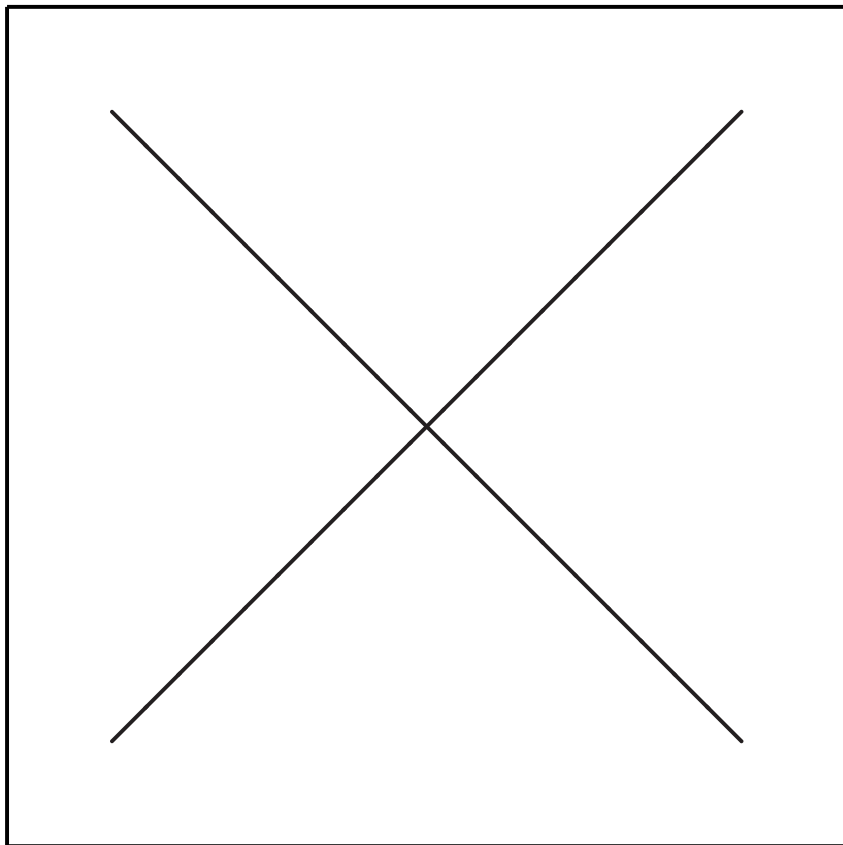
and 32-bit integers r_0, r_1, \dots

register
read

Electricity takes time to percolate through wires and gates.

If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

Build circuit with more gates to multiply (e.g.) 32-bit integers:



(Details omitted.)

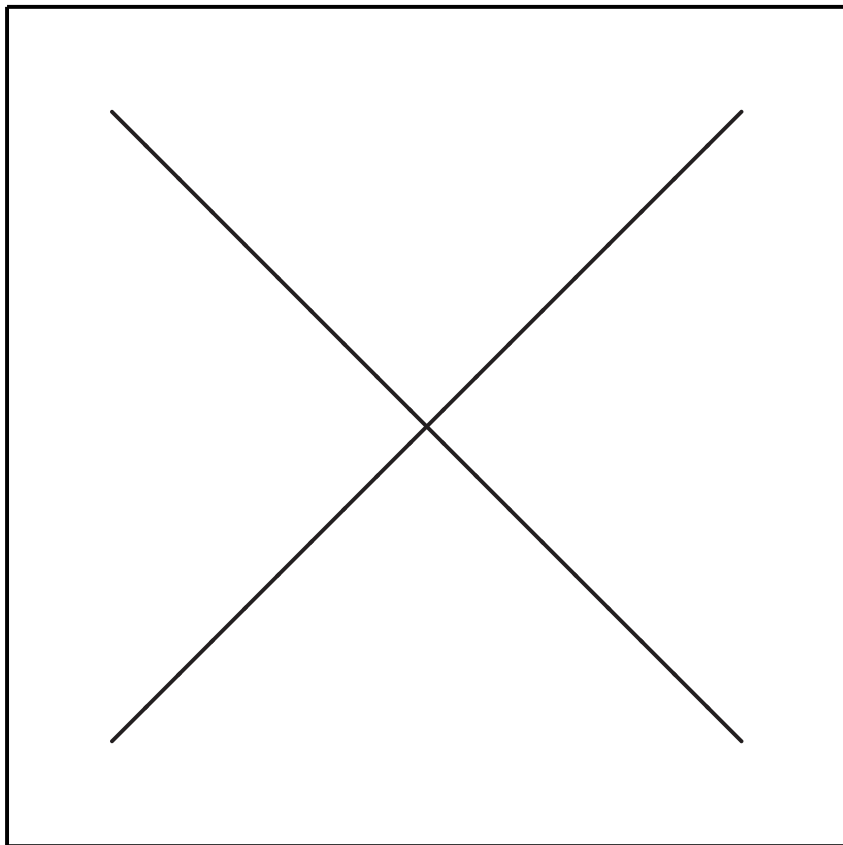
Build circuit to compute 32-bit integer r_i given 4-bit integer i and 32-bit integers r_0, r_1, \dots, r_{15} :

register
read

Electricity takes time to percolate through wires and gates.

If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

Build circuit with more gates to multiply (e.g.) 32-bit integers:



(Details omitted.)

Build circuit to compute 32-bit integer r_i given 4-bit integer i and 32-bit integers r_0, r_1, \dots, r_{15} :

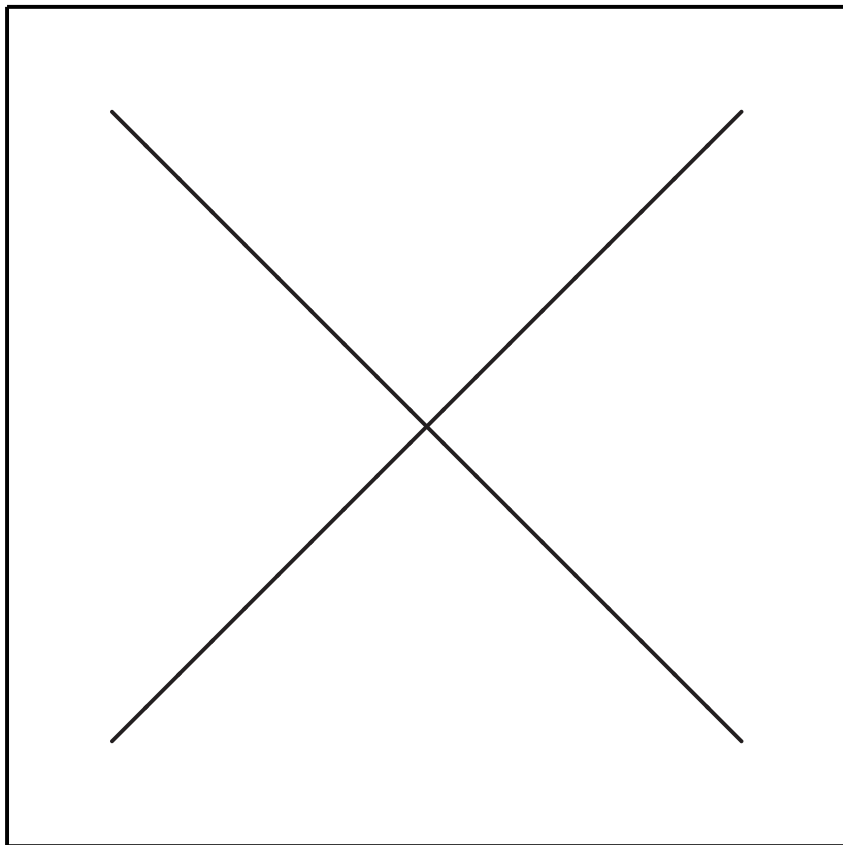
register
read

Build circuit for “register write”:
 $r_0, \dots, r_{15}, s, i \mapsto r'_0, \dots, r'_{15}$
 where $r'_j = r_j$ except $r'_i = s$.

Electricity takes time to percolate through wires and gates.

If f_0, f_1, g_0, g_1 are stable then h_0, h_1, h_2, h_3 are stable a few moments later.

Build circuit with more gates to multiply (e.g.) 32-bit integers:



(Details omitted.)

Build circuit to compute 32-bit integer r_i given 4-bit integer i and 32-bit integers r_0, r_1, \dots, r_{15} :

register
read

Build circuit for “register write”:

$r_0, \dots, r_{15}, s, i \mapsto r'_0, \dots, r'_{15}$

where $r'_j = r_j$ except $r'_i = s$.

Build circuit for addition. Etc.

ity takes time to
e through wires and gates.

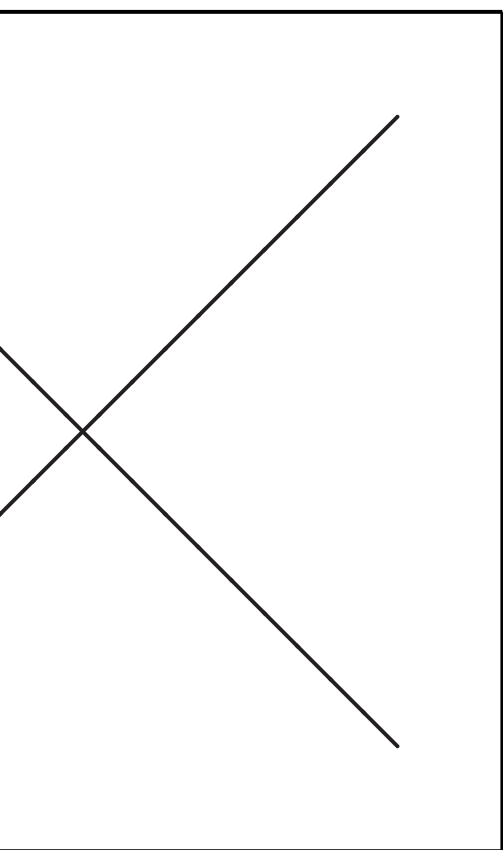
g_0, g_1 are stable

h_1, h_2, h_3 are stable

oments later.

circuit with more gates

ply (e.g.) 32-bit integers:



omitted.)

Build circuit to compute
32-bit integer r_i
given 4-bit integer i
and 32-bit integers r_0, r_1, \dots, r_{15} :

register
read

Build circuit for “register write”:

$r_0, \dots, r_{15}, s, i \mapsto r'_0, \dots, r'_{15}$

where $r'_j = r_j$ except $r'_i = s$.

Build circuit for addition. Etc.

r_0, \dots, r_{15}
where r'_ℓ

regist
rea

r

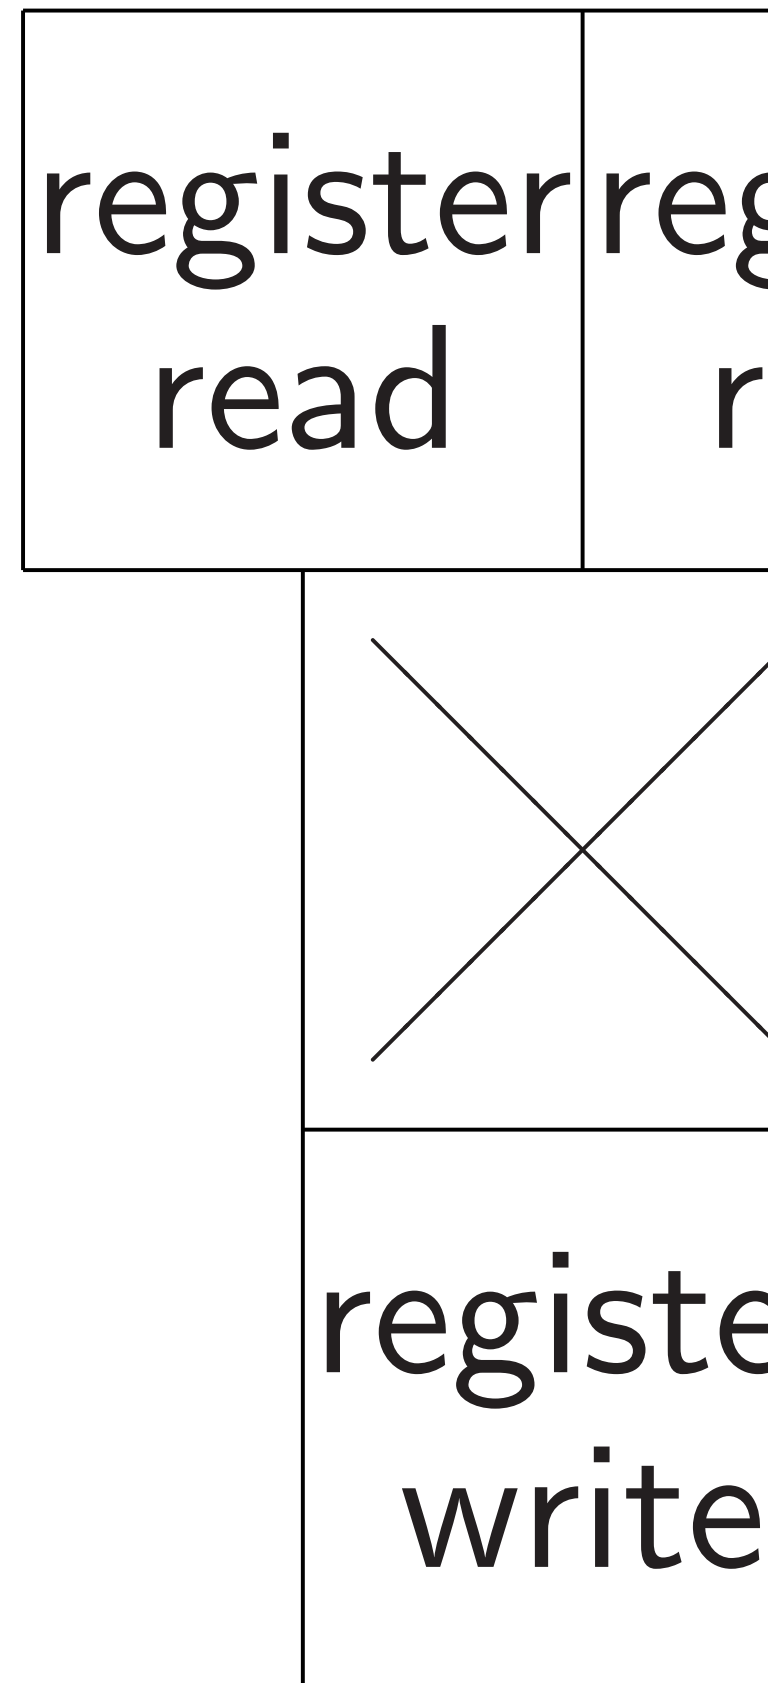
me to
wires and gates.
stable
are stable
ter.
more gates
32-bit integers:

Build circuit to compute
32-bit integer r_i
given 4-bit integer i
and 32-bit integers r_0, r_1, \dots, r_{15} :

register
read

Build circuit for “register write”:
 $r_0, \dots, r_{15}, s, i \mapsto r'_0, \dots, r'_{15}$
where $r'_j = r_j$ except $r'_i = s$.
Build circuit for addition. Etc.

$r_0, \dots, r_{15}, i, j, k \mapsto$
where $r'_\ell = r_\ell$ except



gates.

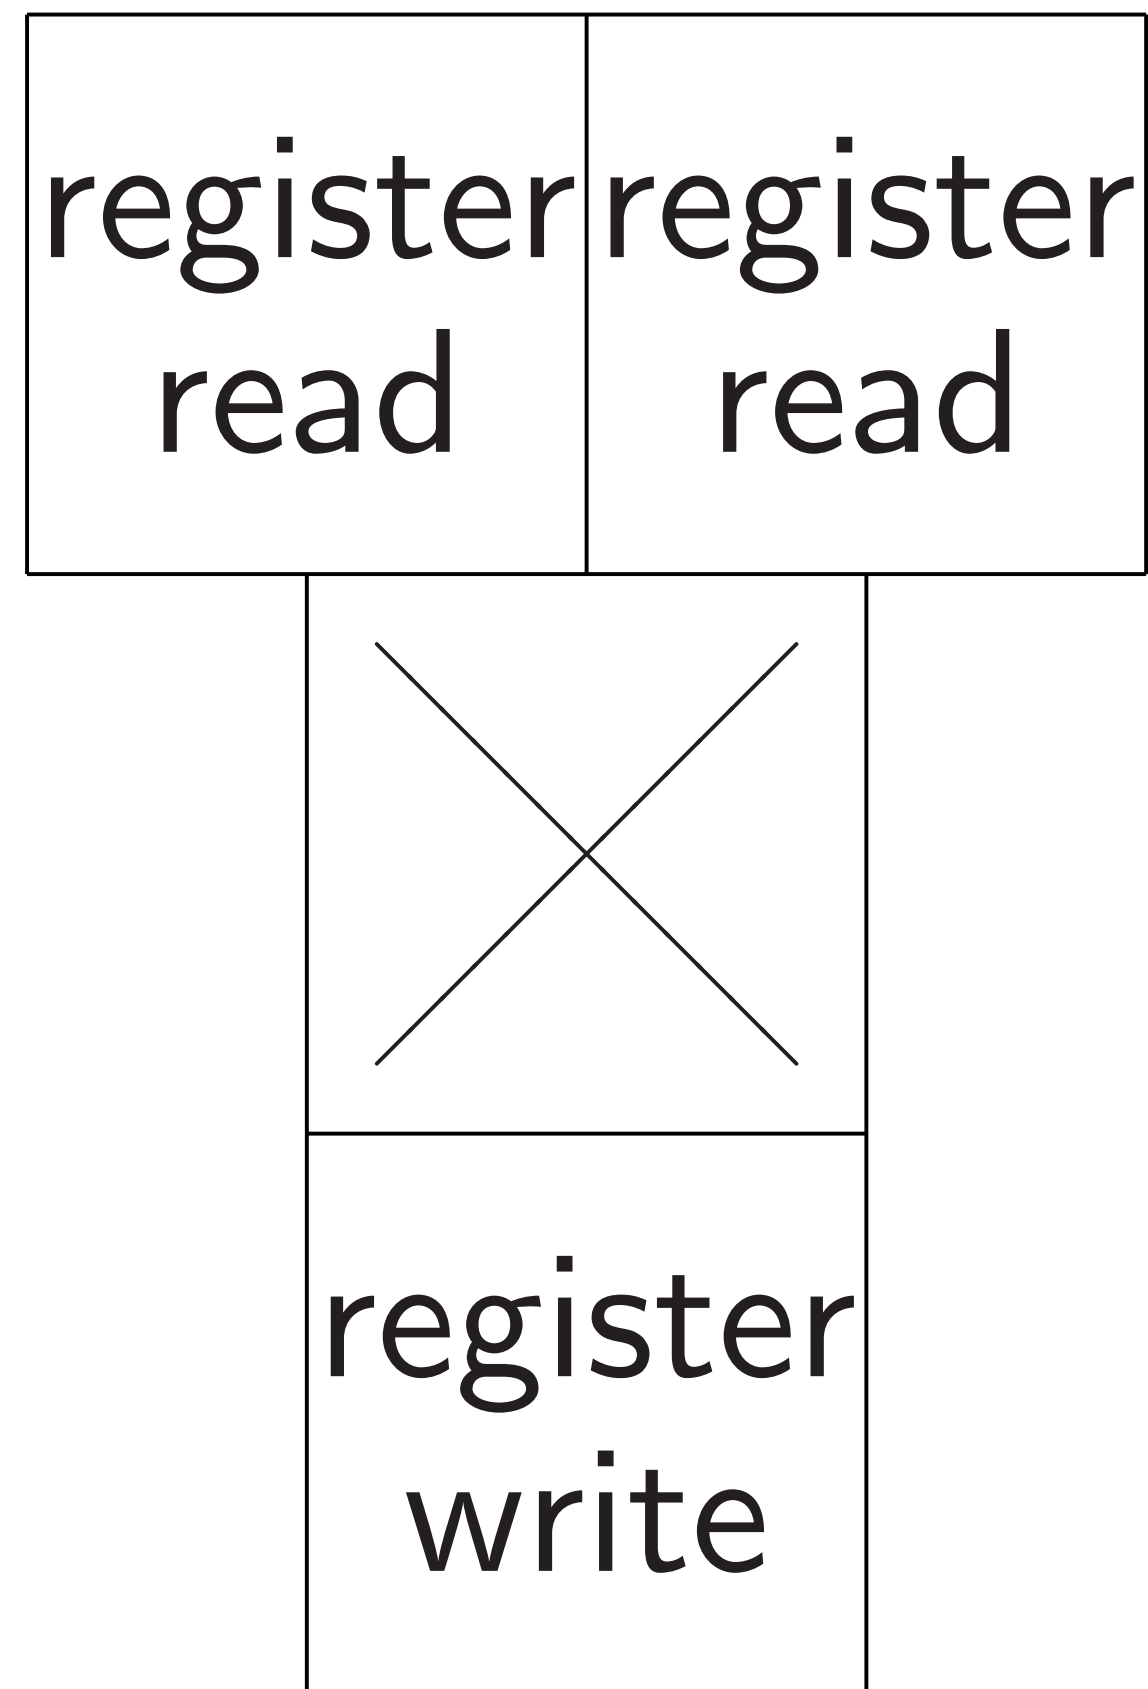
Build circuit to compute
32-bit integer r_i
given 4-bit integer i
and 32-bit integers r_0, r_1, \dots, r_{15} :



s
egers:

Build circuit for “register write”:
 $r_0, \dots, r_{15}, s, i \mapsto r'_0, \dots, r'_{15}$
where $r'_j = r_j$ except $r'_i = s$.
Build circuit for addition. Etc.

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
where $r'_\ell = r_\ell$ except $r'_i = r_j$

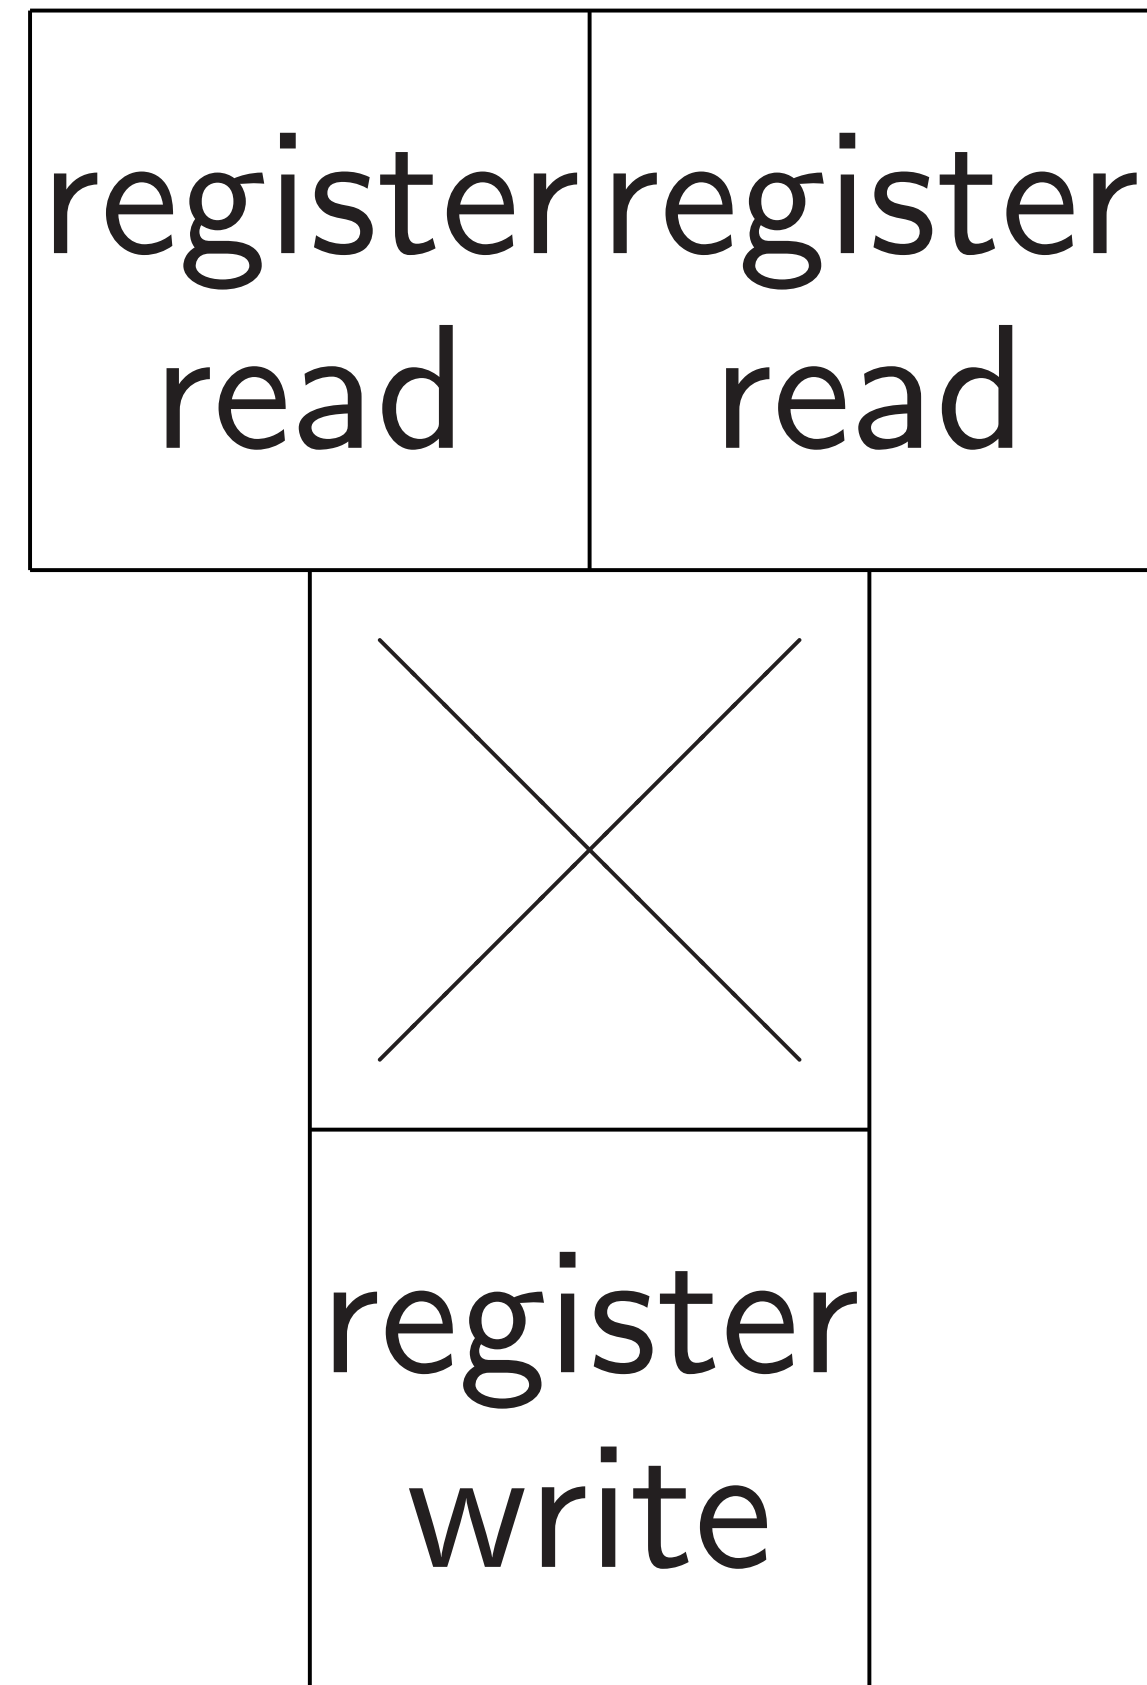


Build circuit to compute
32-bit integer r_i
given 4-bit integer i
and 32-bit integers r_0, r_1, \dots, r_{15} :

register
read

Build circuit for “register write”:
 $r_0, \dots, r_{15}, s, i \mapsto r'_0, \dots, r'_{15}$
where $r'_j = r_j$ except $r'_i = s$.
Build circuit for addition. Etc.

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



circuit to compute

integer r_i

bit integer i

bit integers r_0, r_1, \dots, r_{15} :

register
read

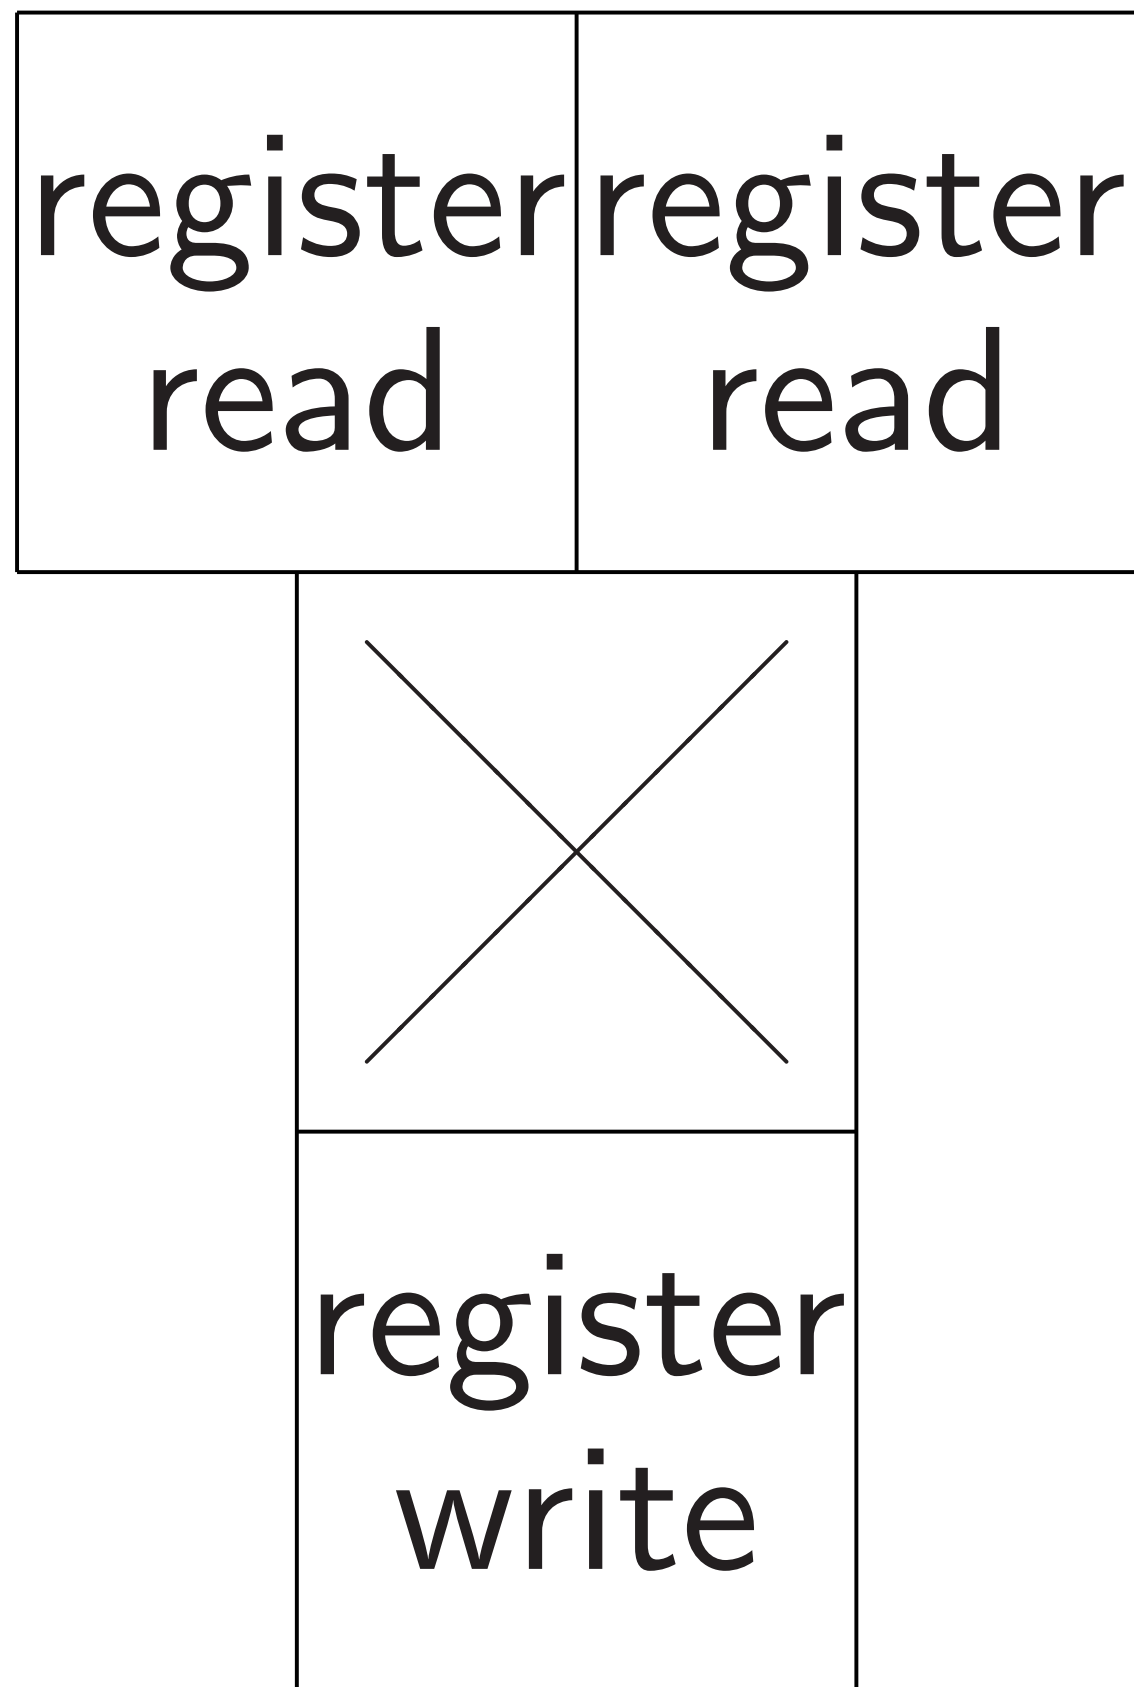
circuit for “register write”:

$r_0, \dots, r_{15}, s, i \mapsto r'_0, \dots, r'_{15}$

$r'_\ell = r_\ell$ except $r'_i = s$.

circuit for addition. Etc.

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



Add more

More arith

replace (

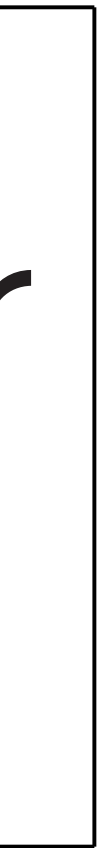
(“×”, i, j)

(“+”, i, j)

compute

i

registers r_0, r_1, \dots, r_{15} :



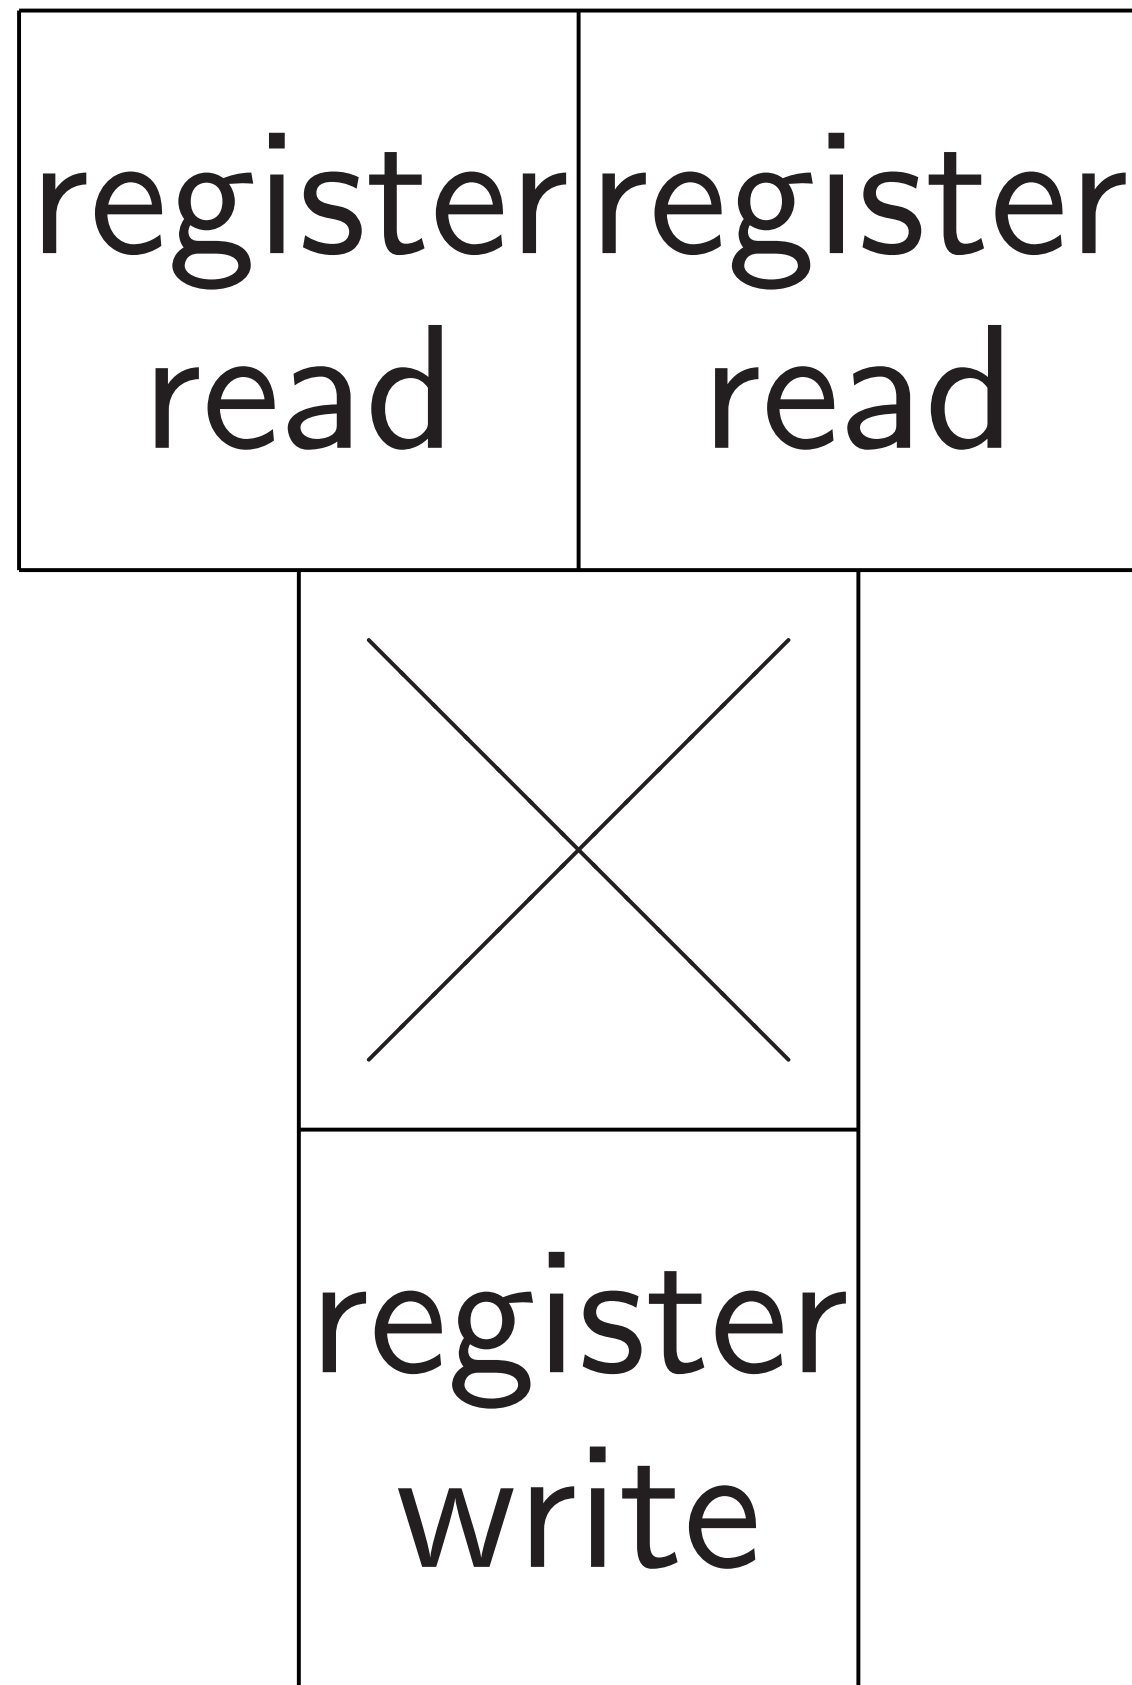
“register write”:

r'_0, \dots, r'_{15}

except $r'_i = s$.

addition. Etc.

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
 where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



Add more flexibility

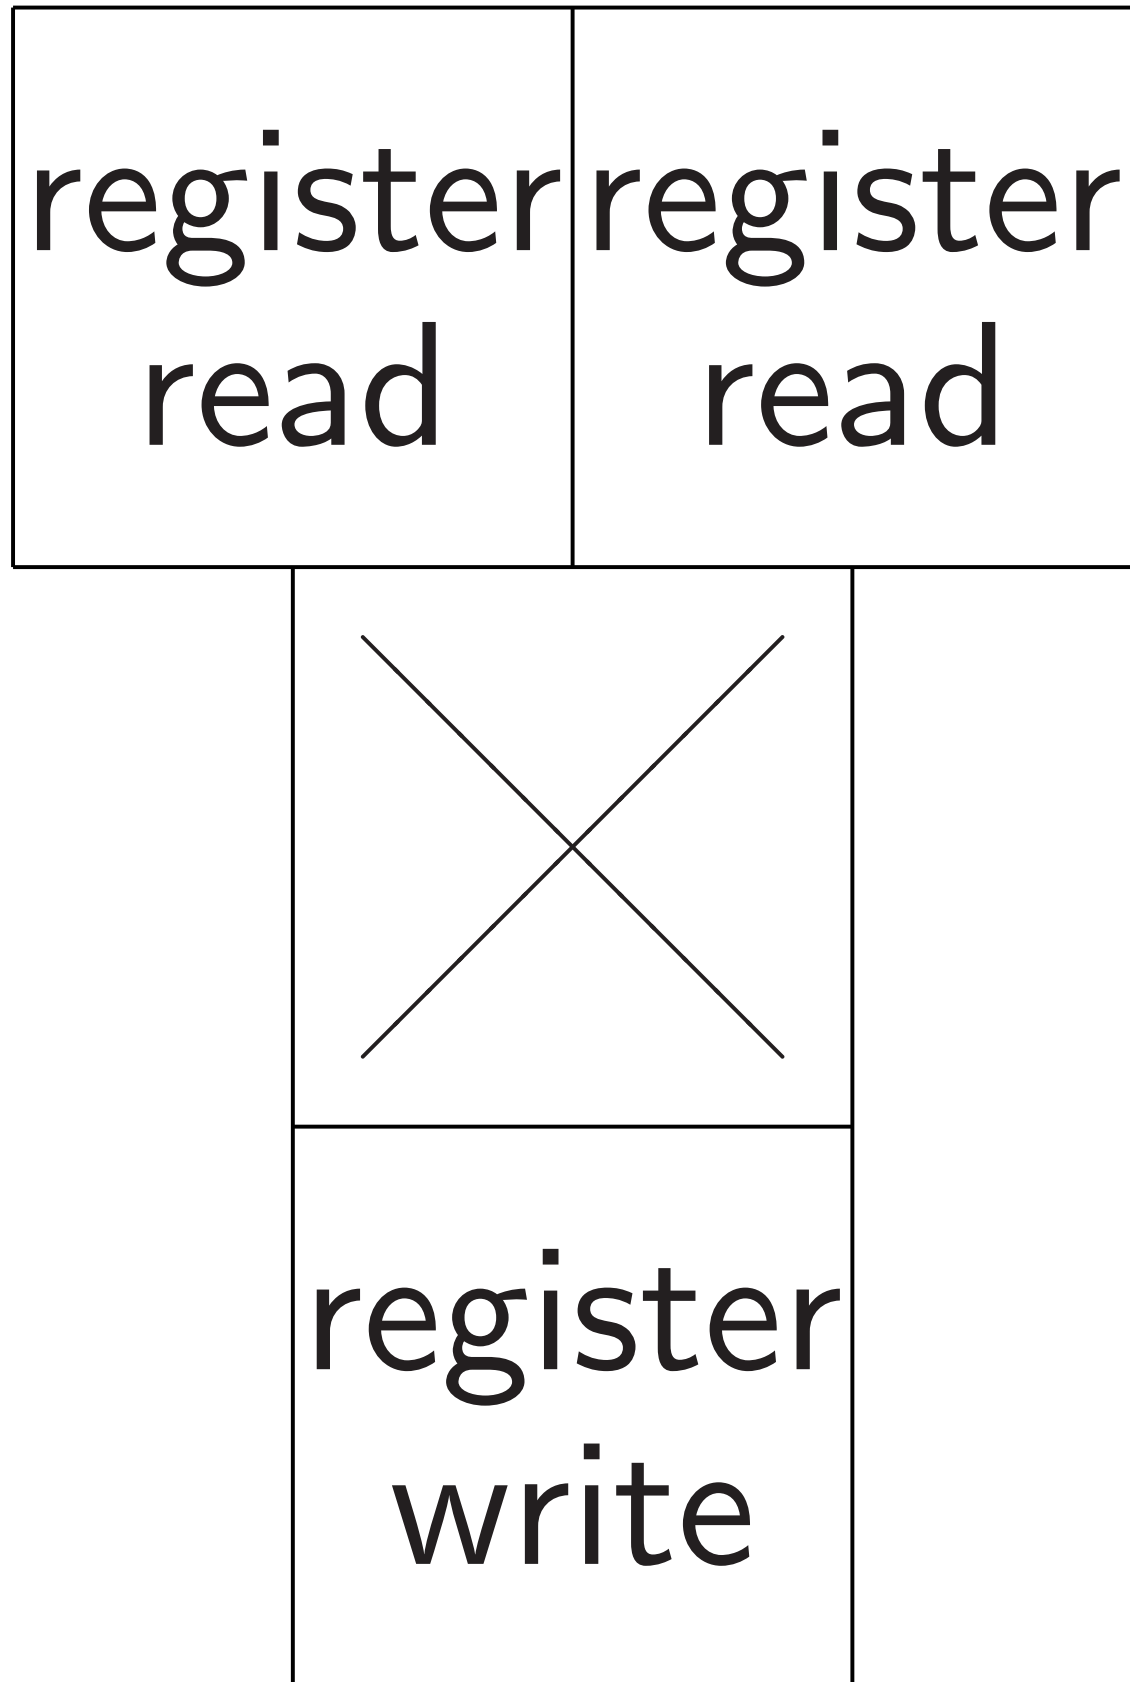
More arithmetic:

replace (i, j, k) with

$(“\times”, i, j, k)$ and

$(“+”, i, j, k)$ and

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
 where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



Add more flexibility.

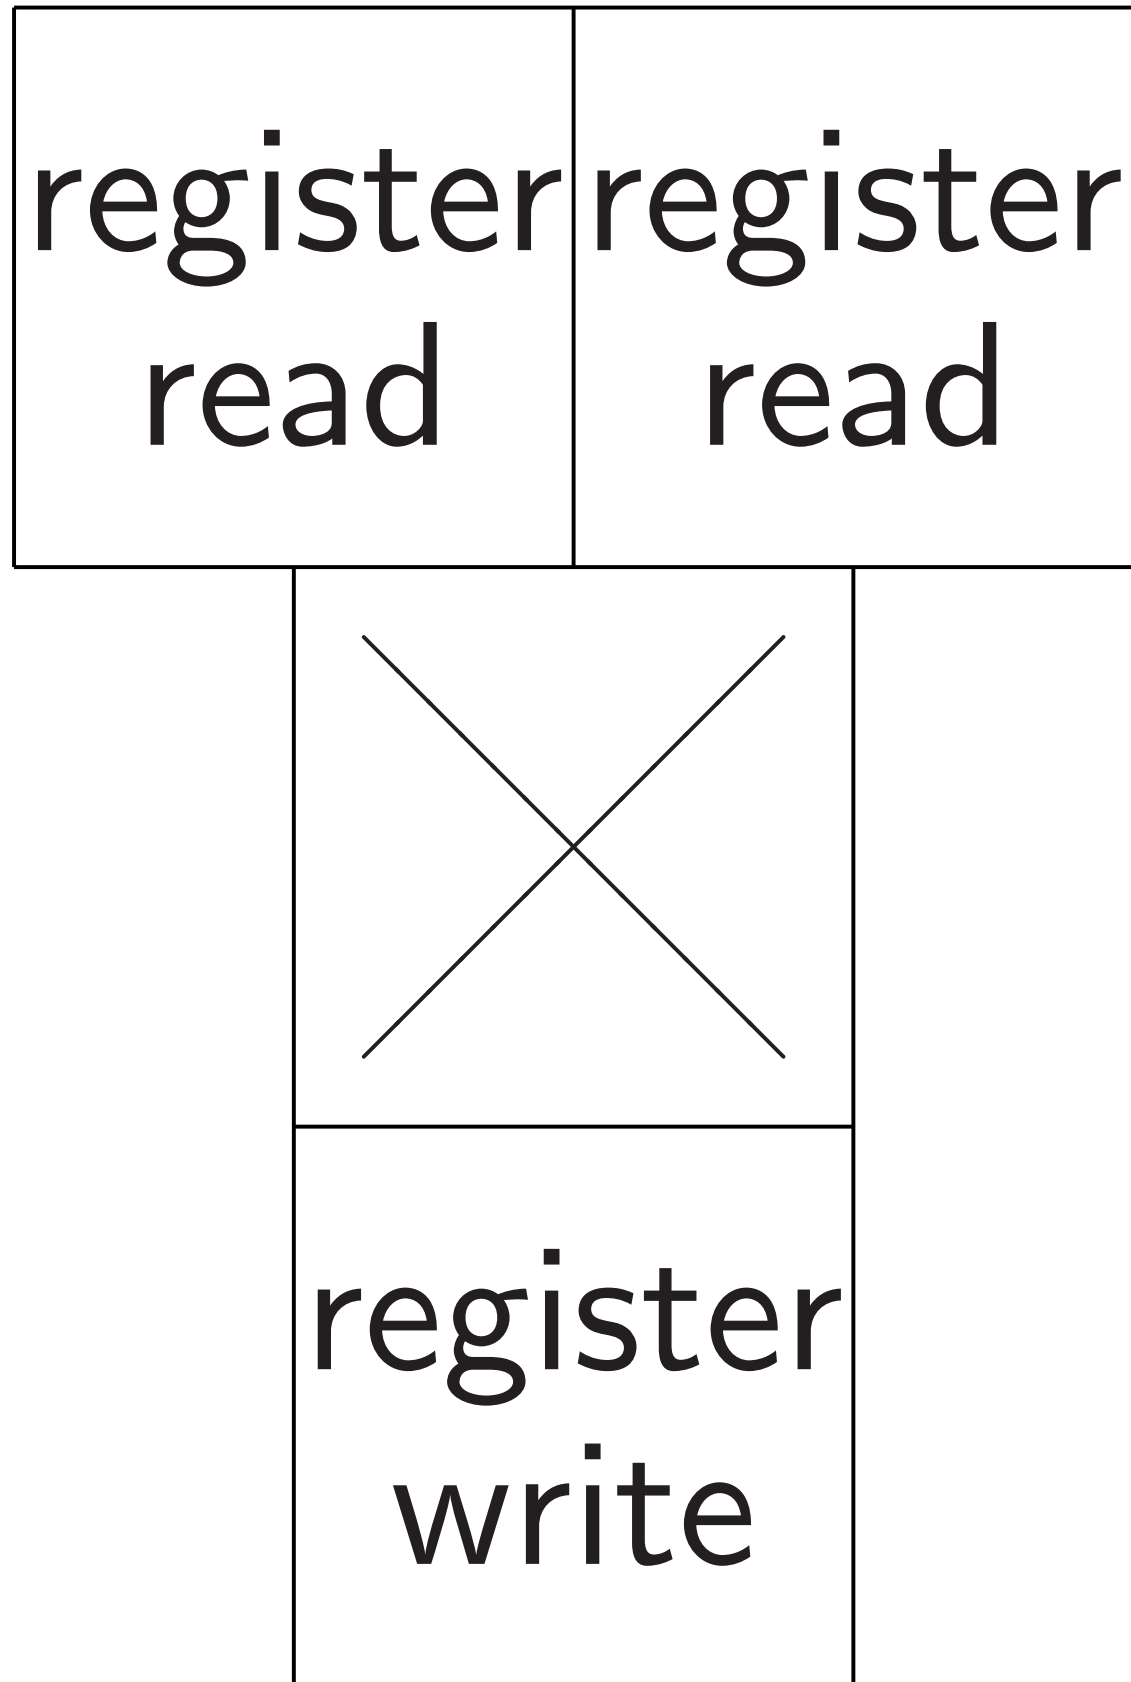
More arithmetic:

replace (i, j, k) with

$(\text{"\times"}, i, j, k)$ and

$(\text{"+"}, i, j, k)$ and more options

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
 where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



Add more flexibility.

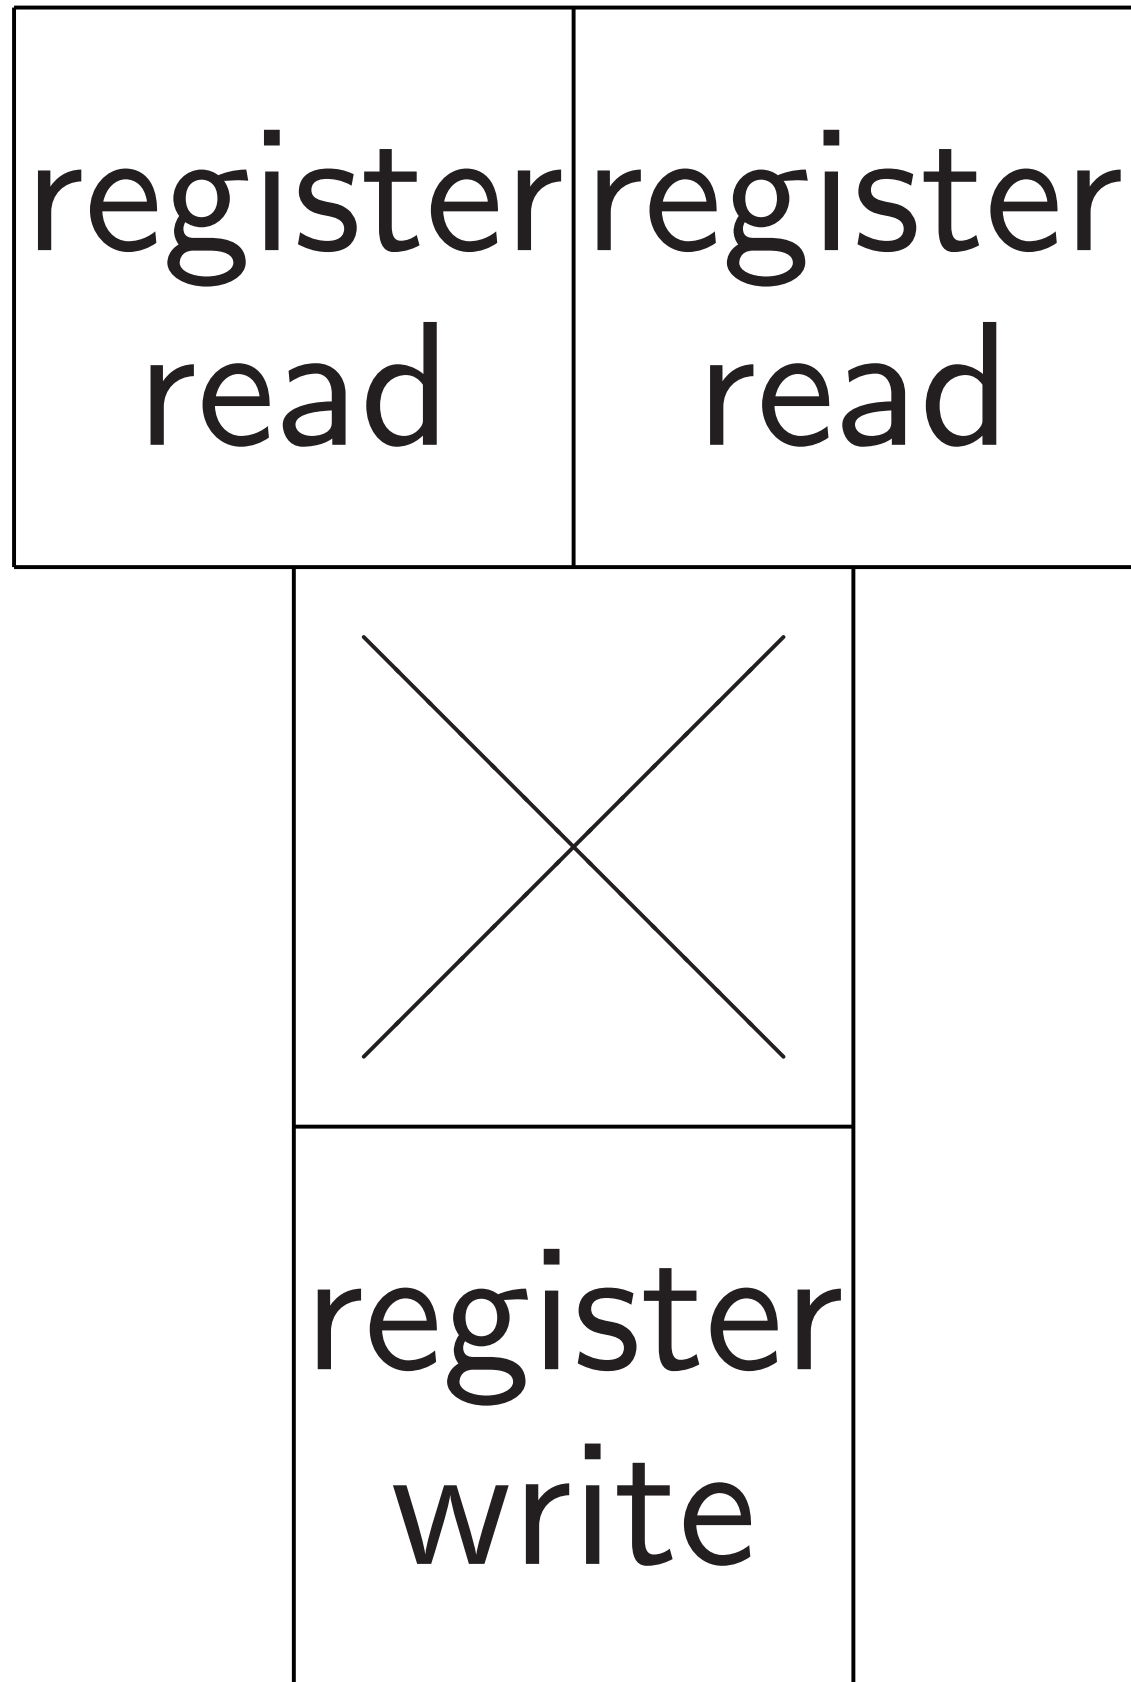
More arithmetic:

replace (i, j, k) with

$(“\times”, i, j, k)$ and

$(“+”, i, j, k)$ and more options.

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
 where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



Add more flexibility.

More arithmetic:

replace (i, j, k) with

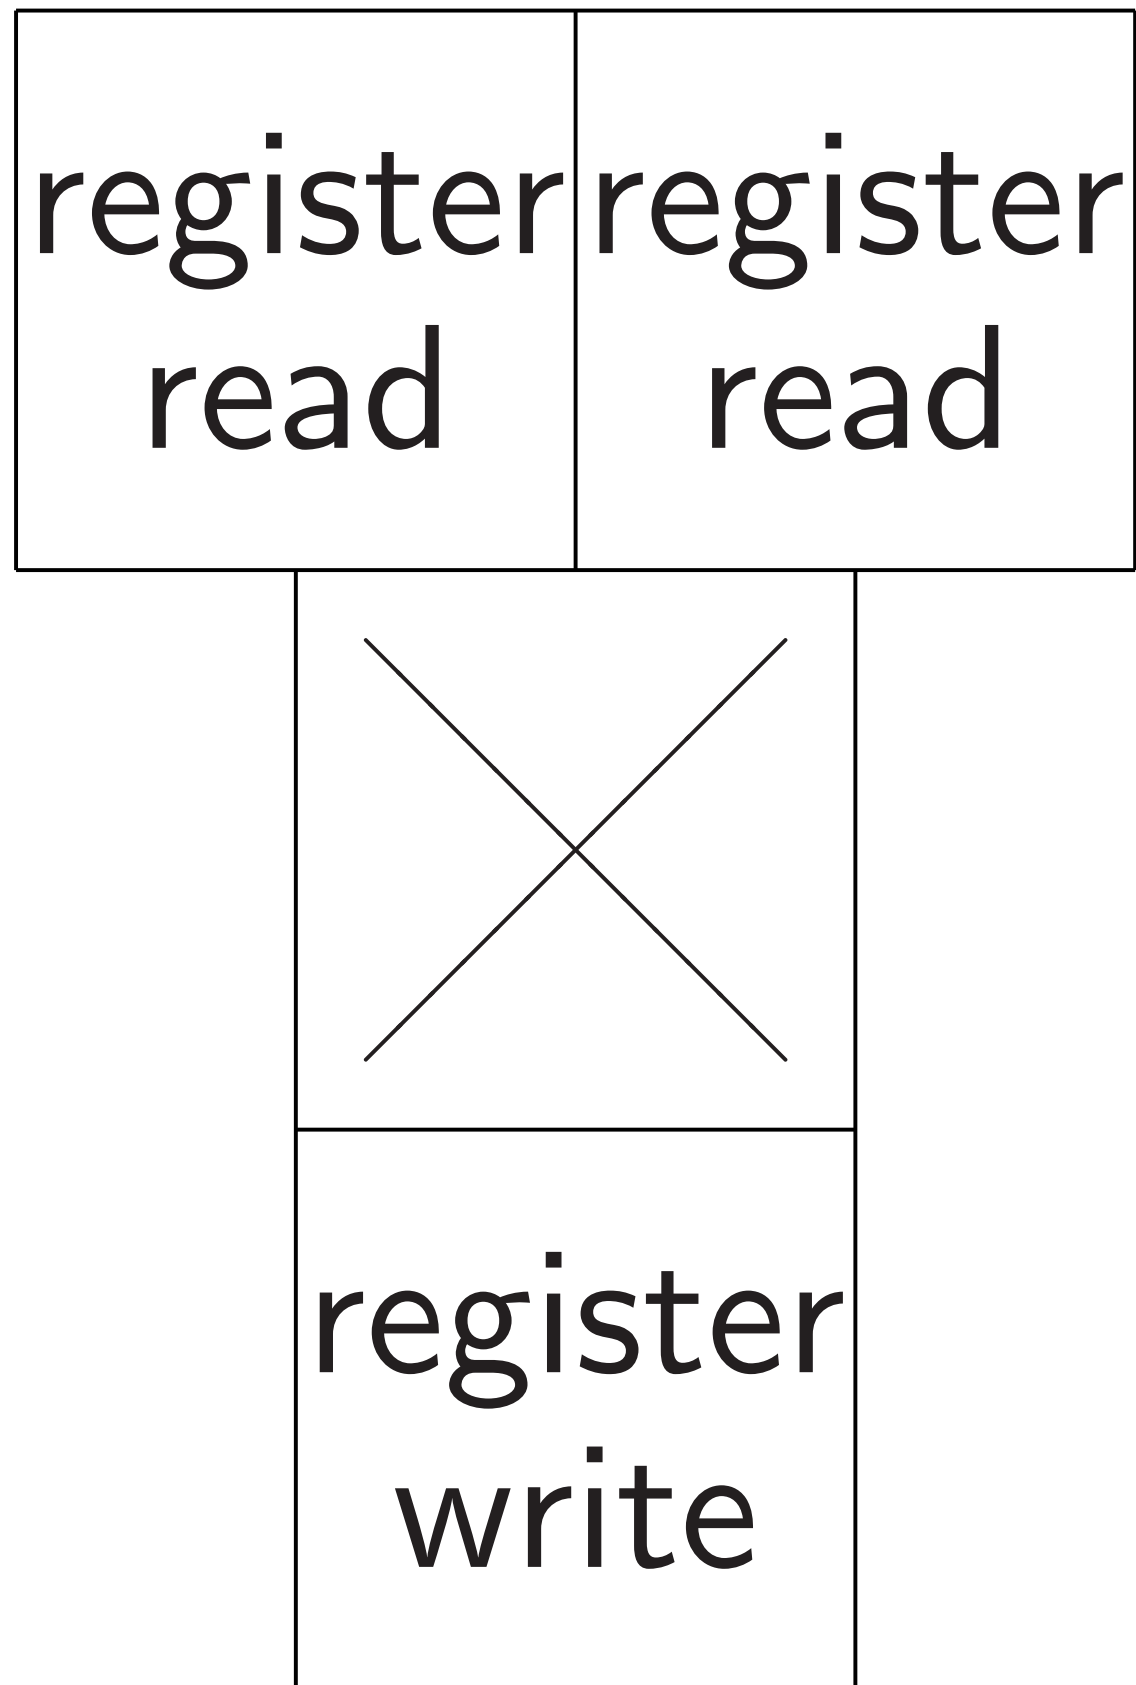
$(“\times”, i, j, k)$ and

$(“+”, i, j, k)$ and more options.

“Instruction fetch”:

$p \mapsto o_p, i_p, j_p, k_p, p'$.

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
 where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



Add more flexibility.

More arithmetic:

replace (i, j, k) with

$(“\times”, i, j, k)$ and

$(“+”, i, j, k)$ and more options.

“Instruction fetch”:

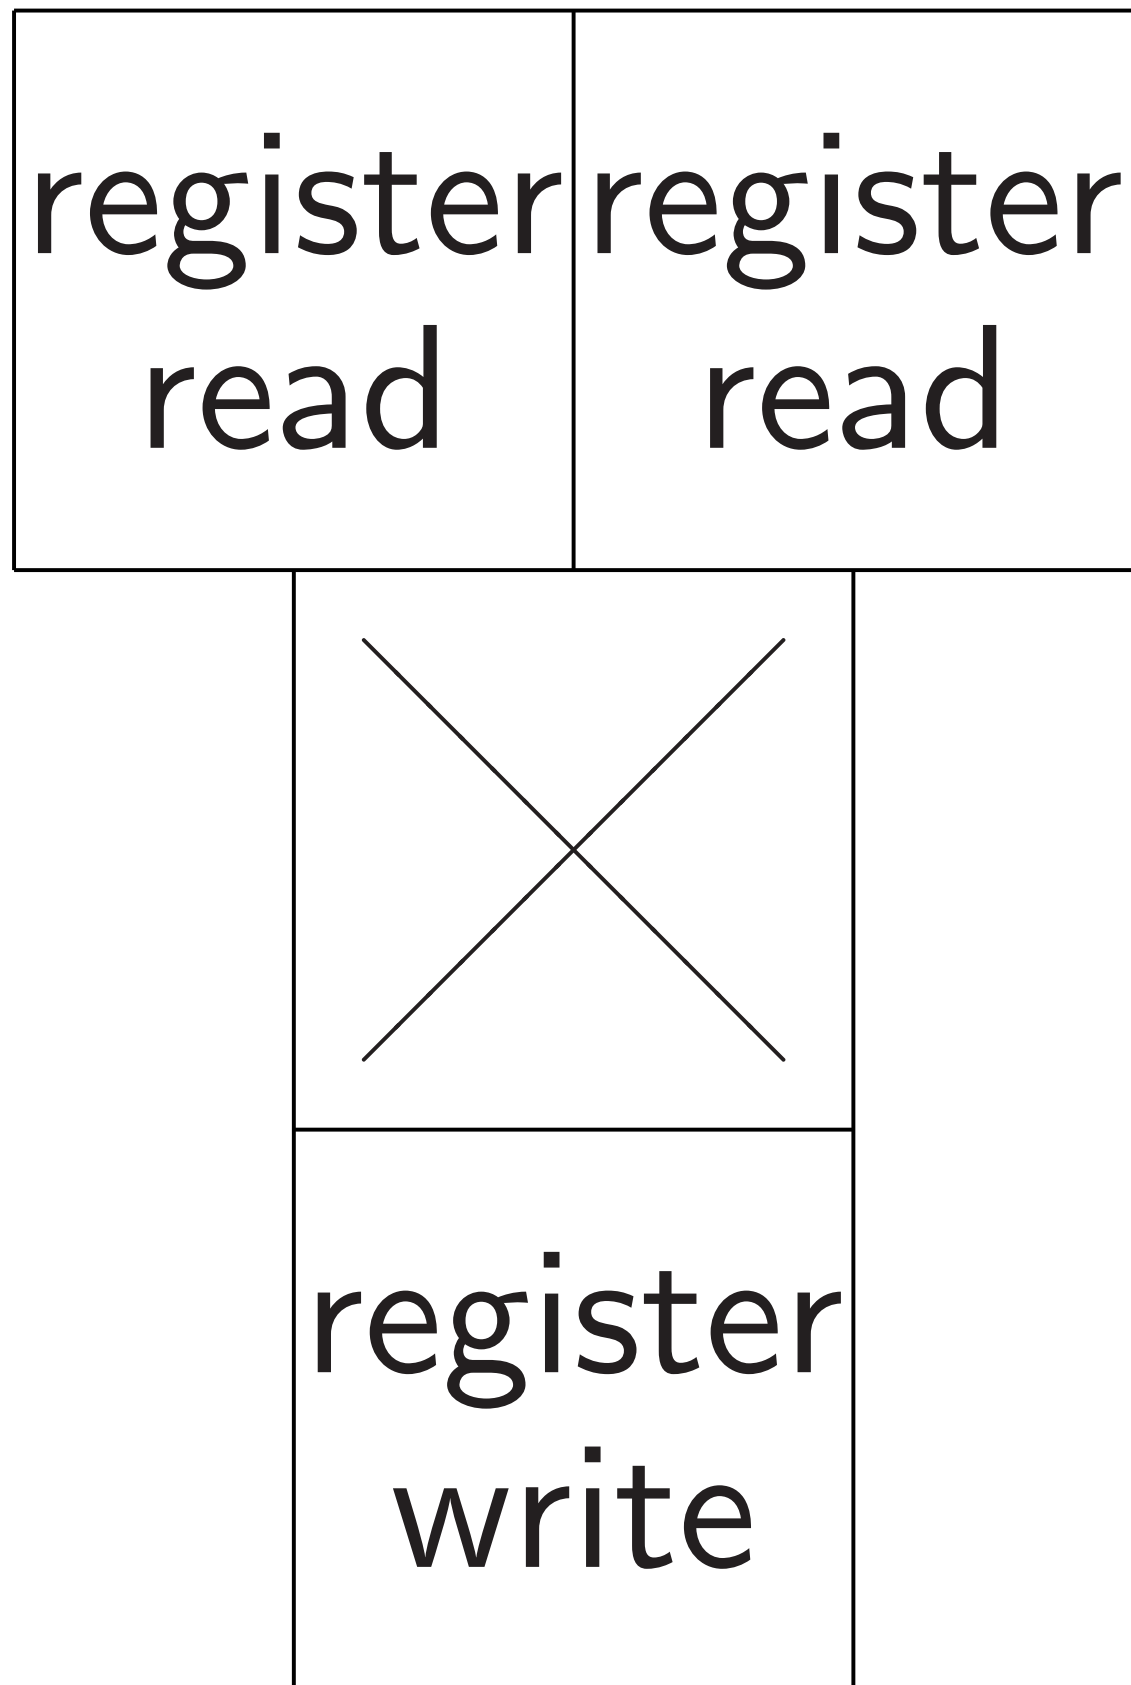
$p \mapsto o_p, i_p, j_p, k_p, p'$.

“Instruction decode”:

decompression of compressed

format for o_p, i_p, j_p, k_p, p' .

$r_0, \dots, r_{15}, i, j, k \mapsto r'_0, \dots, r'_{15}$
 where $r'_\ell = r_\ell$ except $r'_i = r_j r_k$:



Add more flexibility.

More arithmetic:

replace (i, j, k) with

$(“\times”, i, j, k)$ and

$(“+”, i, j, k)$ and more options.

“Instruction fetch”:

$p \mapsto o_p, i_p, j_p, k_p, p'$.

“Instruction decode”:

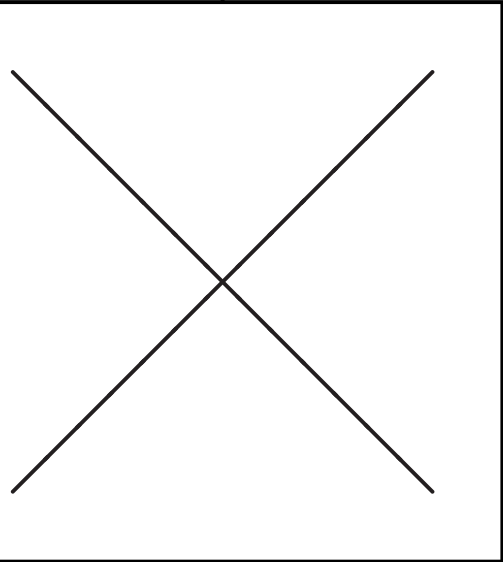
decompression of compressed
 format for o_p, i_p, j_p, k_p, p' .

More (but slower) storage:

“load” from and “store” to
 larger “RAM” arrays.

$15, i, j, k \mapsto r'_0, \dots, r'_{15}$
 $= r_\ell$ except $r'_i = r_j r_k$:

register read	register read
---------------	---------------



register write

Add more flexibility.

More arithmetic:

replace (i, j, k) with

$(\text{"}\times\text{"}, i, j, k)$ and

$(\text{"}\text{+}\text{"}, i, j, k)$ and more options.

“Instruction fetch”:

$p \mapsto o_p, i_p, j_p, k_p, p'$.

“Instruction decode”:

decompression of compressed format for o_p, i_p, j_p, k_p, p' .

More (but slower) storage:

“load” from and “store” to larger “RAM” arrays.

Build “fl
 storing (

Hook (p
 flip-flops

Hook ou
 into the

At each
 flip-flops

with the

Clock ne
 for elect

all the w
 from flip

$\rightarrow r'_0, \dots, r'_{15}$
except $r'_i = r_j r_k$:

register
head

er

33

Add more flexibility.

More arithmetic:

replace (i, j, k) with

$(\text{"}\times\text{"}, i, j, k)$ and

$(\text{"}\text{+}\text{"}, i, j, k)$ and more options.

"Instruction fetch":

$p \mapsto o_p, i_p, j_p, k_p, p'$.

"Instruction decode":

decompression of compressed
format for o_p, i_p, j_p, k_p, p' .

More (but slower) storage:

"load" from and "store" to
larger "RAM" arrays.

34

Build "flip-flops"
storing $(p, r_0, \dots,$

Hook (p, r_0, \dots, r_{15})
flip-flops into circuit.

Hook outputs $(p',$
into the same flip-

At each "clock tick"
flip-flops are overwrit-
ten with the outputs.

Clock needs to be
for electricity to pass
all the way through
from flip-flops to f

Add more flexibility.

More arithmetic:

replace (i, j, k) with

$(“\times”, i, j, k)$ and

$(“+”, i, j, k)$ and more options.

“Instruction fetch”:

$p \mapsto o_p, i_p, j_p, k_p, p'$.

“Instruction decode”:

decompression of compressed

format for o_p, i_p, j_p, k_p, p' .

More (but slower) storage:

“load” from and “store” to

larger “RAM” arrays.

Build “flip-flops”

storing (p, r_0, \dots, r_{15}) .

Hook (p, r_0, \dots, r_{15})

flip-flops into circuit inputs.

Hook outputs $(p', r'_0, \dots, r'_{15})$

into the same flip-flops.

At each “clock tick”,

flip-flops are overwritten

with the outputs.

Clock needs to be slow enough

for electricity to percolate

all the way through the circuit

from flip-flops to flip-flops.

Add more flexibility.

More arithmetic:

replace (i, j, k) with

$(“\times”, i, j, k)$ and

$(“+”, i, j, k)$ and more options.

“Instruction fetch”:

$p \mapsto o_p, i_p, j_p, k_p, p'$.

“Instruction decode”:

decompression of compressed

format for o_p, i_p, j_p, k_p, p' .

More (but slower) storage:

“load” from and “store” to
larger “RAM” arrays.

Build “flip-flops”

storing (p, r_0, \dots, r_{15}) .

Hook (p, r_0, \dots, r_{15})

flip-flops into circuit inputs.

Hook outputs $(p', r'_0, \dots, r'_{15})$

into the same flip-flops.

At each “clock tick”,

flip-flops are overwritten

with the outputs.

Clock needs to be slow enough

for electricity to percolate

all the way through the circuit,

from flip-flops to flip-flops.

re flexibility.

ithmetic:

(i, j, k) with

(j, k) and

(j, k) and more options.

tion fetch”:

i_p, j_p, k_p, p' .

tion decode”:

ression of compressed

or o_p, i_p, j_p, k_p, p' .

ut slower) storage:

rom and “store” to

RAM” arrays.

Build “flip-flops”

storing (p, r_0, \dots, r_{15}) .

Hook (p, r_0, \dots, r_{15})

flip-flops into circuit inputs.

Hook outputs $(p', r'_0, \dots, r'_{15})$

into the same flip-flops.

At each “clock tick”,

flip-flops are overwritten

with the outputs.

Clock needs to be slow enough

for electricity to percolate

all the way through the circuit,

from flip-flops to flip-flops.

Now hav



Further

e.g., rota

Build “flip-flops”

storing (p, r_0, \dots, r_{15}) .

Hook (p, r_0, \dots, r_{15})

flip-flops into circuit inputs.

Hook outputs $(p', r'_0, \dots, r'_{15})$

into the same flip-flops.

At each “clock tick”,

flip-flops are overwritten

with the outputs.

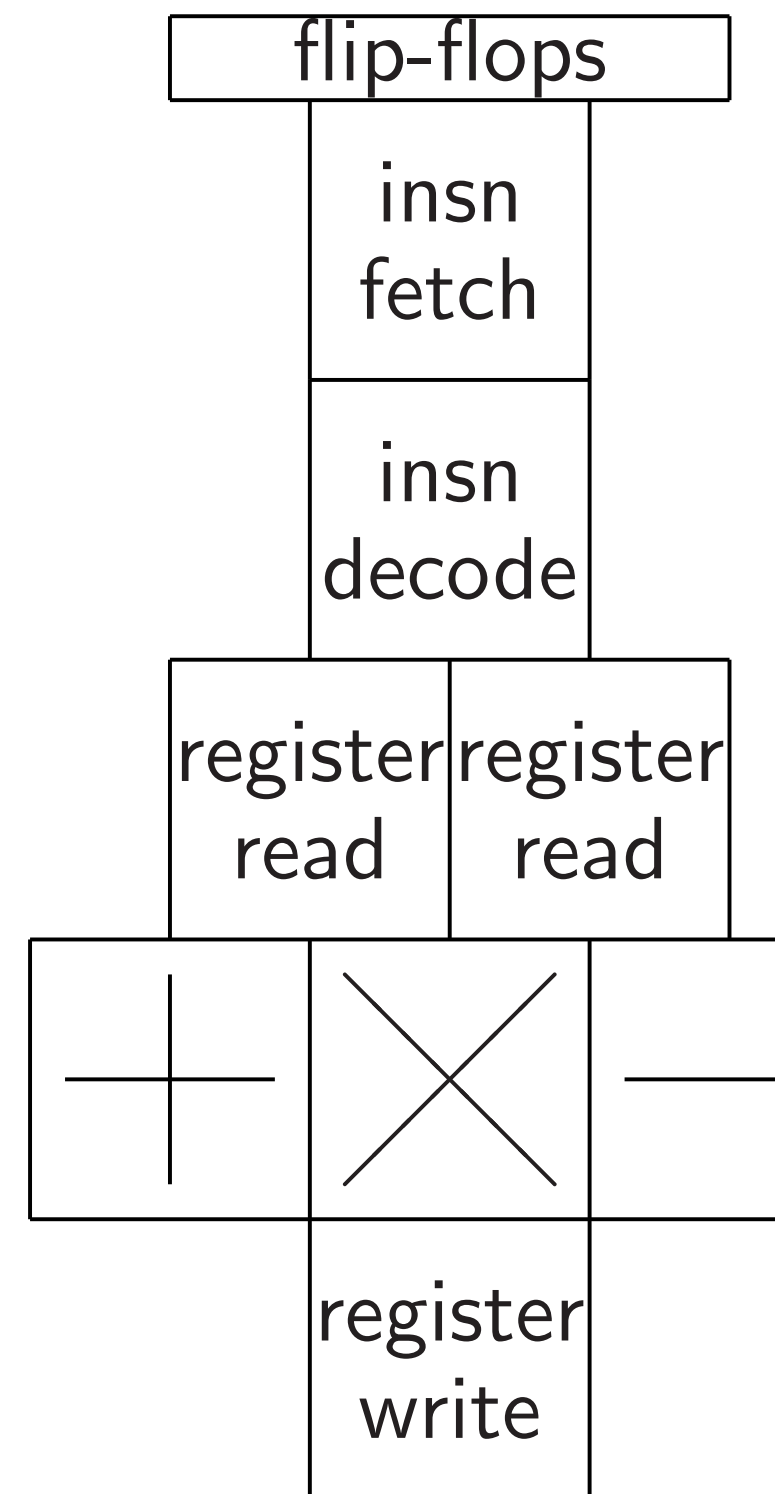
Clock needs to be slow enough

for electricity to percolate

all the way through the circuit,

from flip-flops to flip-flops.

Now have semi-flip-flops



Further flexibility in

e.g., rotation instr

Build “flip-flops”
storing (p, r_0, \dots, r_{15}) .

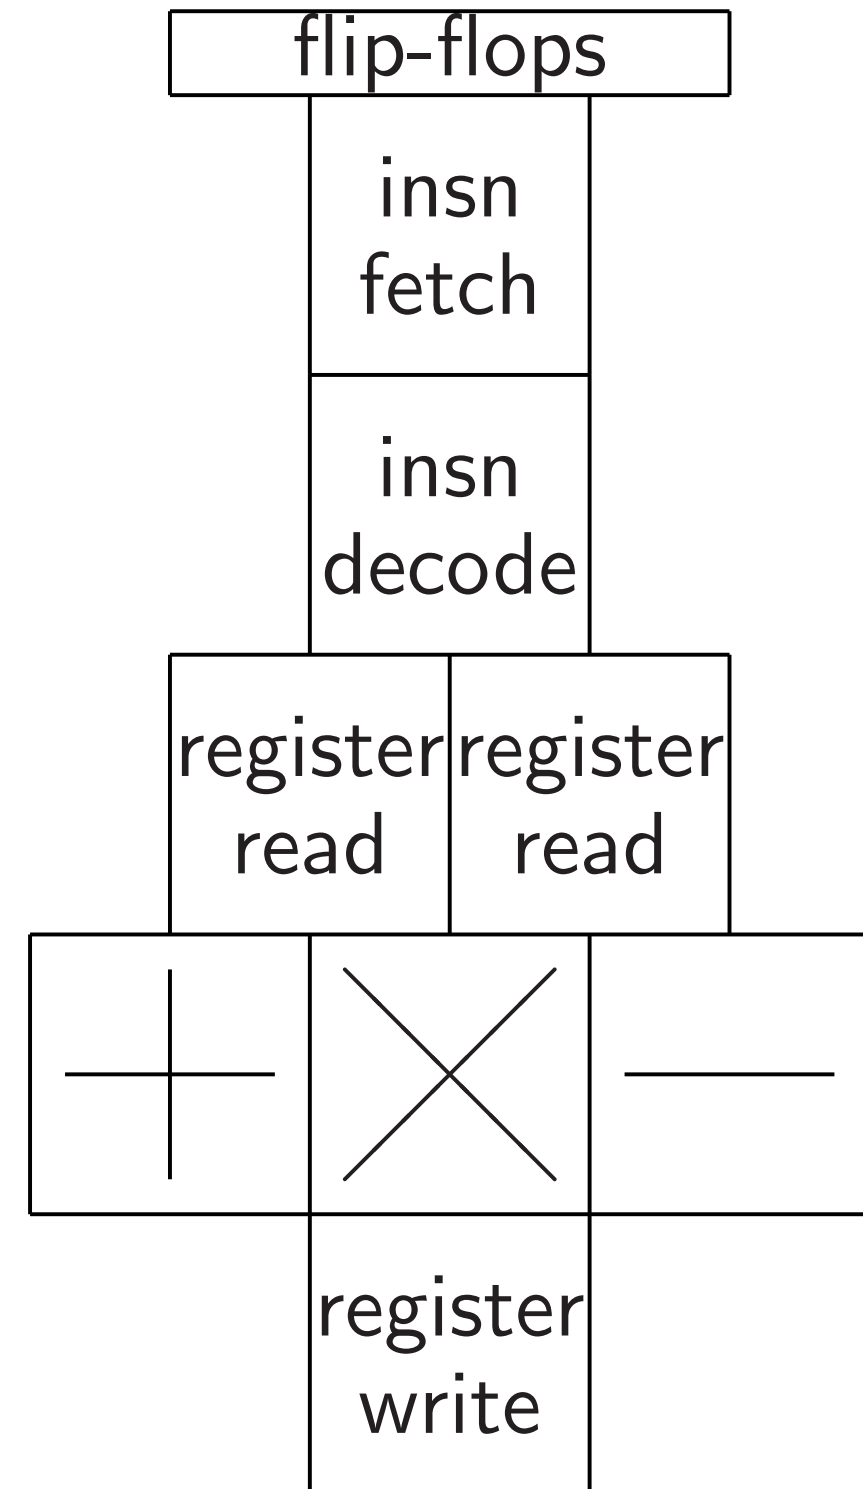
Hook (p, r_0, \dots, r_{15})
flip-flops into circuit inputs.

Hook outputs $(p', r'_0, \dots, r'_{15})$
into the same flip-flops.

At each “clock tick”,
flip-flops are overwritten
with the outputs.

Clock needs to be slow enough
for electricity to percolate
all the way through the circuit,
from flip-flops to flip-flops.

Now have semi-flexible CPU



Further flexibility is useful:
e.g., rotation instructions.

Build “flip-flops”

storing (p, r_0, \dots, r_{15}) .

Hook (p, r_0, \dots, r_{15})

flip-flops into circuit inputs.

Hook outputs $(p', r'_0, \dots, r'_{15})$

into the same flip-flops.

At each “clock tick”,

flip-flops are overwritten

with the outputs.

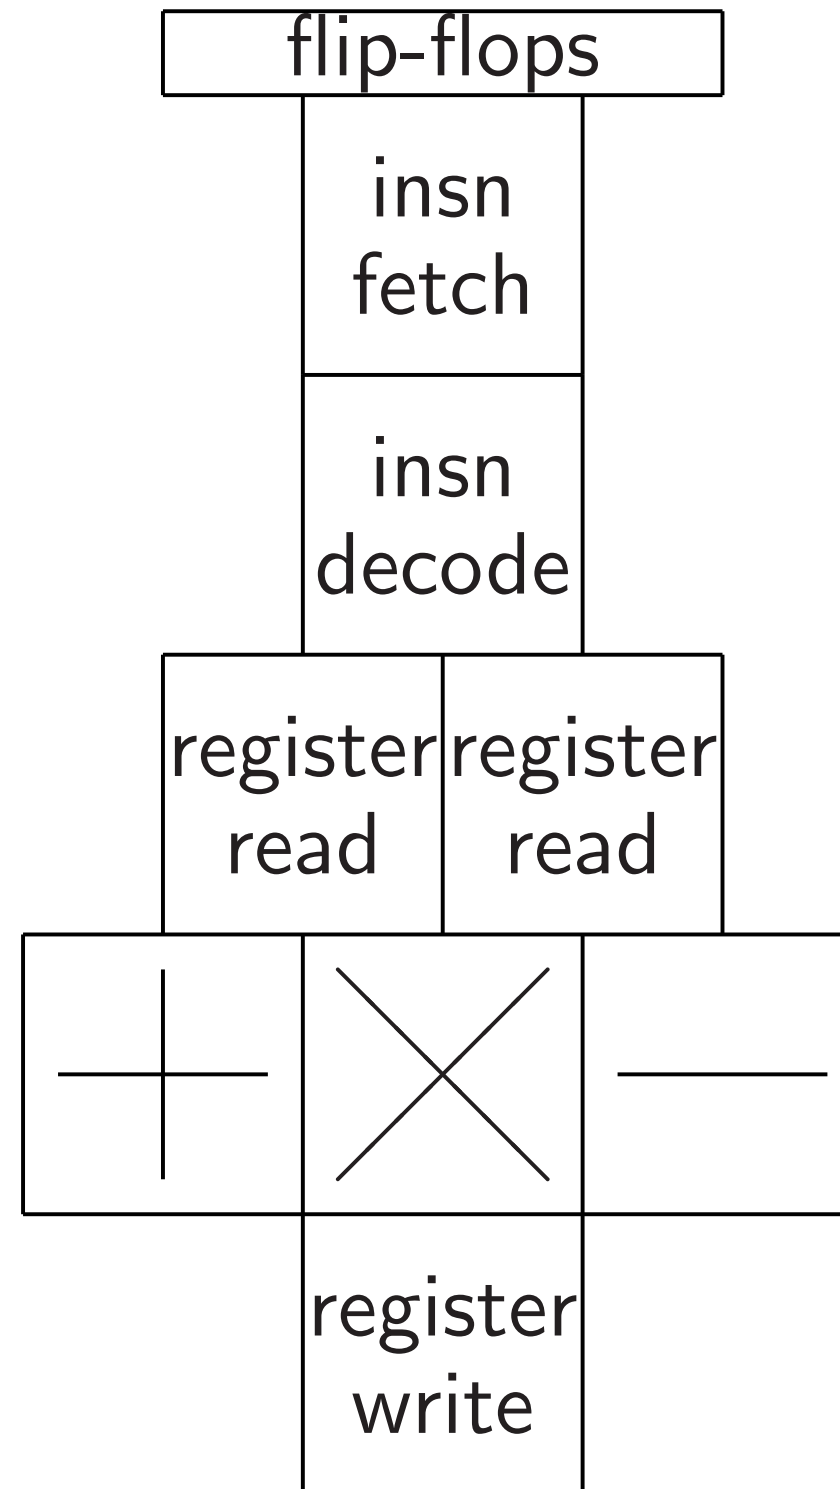
Clock needs to be slow enough

for electricity to percolate

all the way through the circuit,

from flip-flops to flip-flops.

Now have semi-flexible CPU:



Further flexibility is useful:

e.g., rotation instructions.

flip-flops”

(p, r_0, \dots, r_{15}) .

(p, r_0, \dots, r_{15})

into circuit inputs.

outputs $(p', r'_0, \dots, r'_{15})$

same flip-flops.

“clock tick”,

are overwritten

outputs.

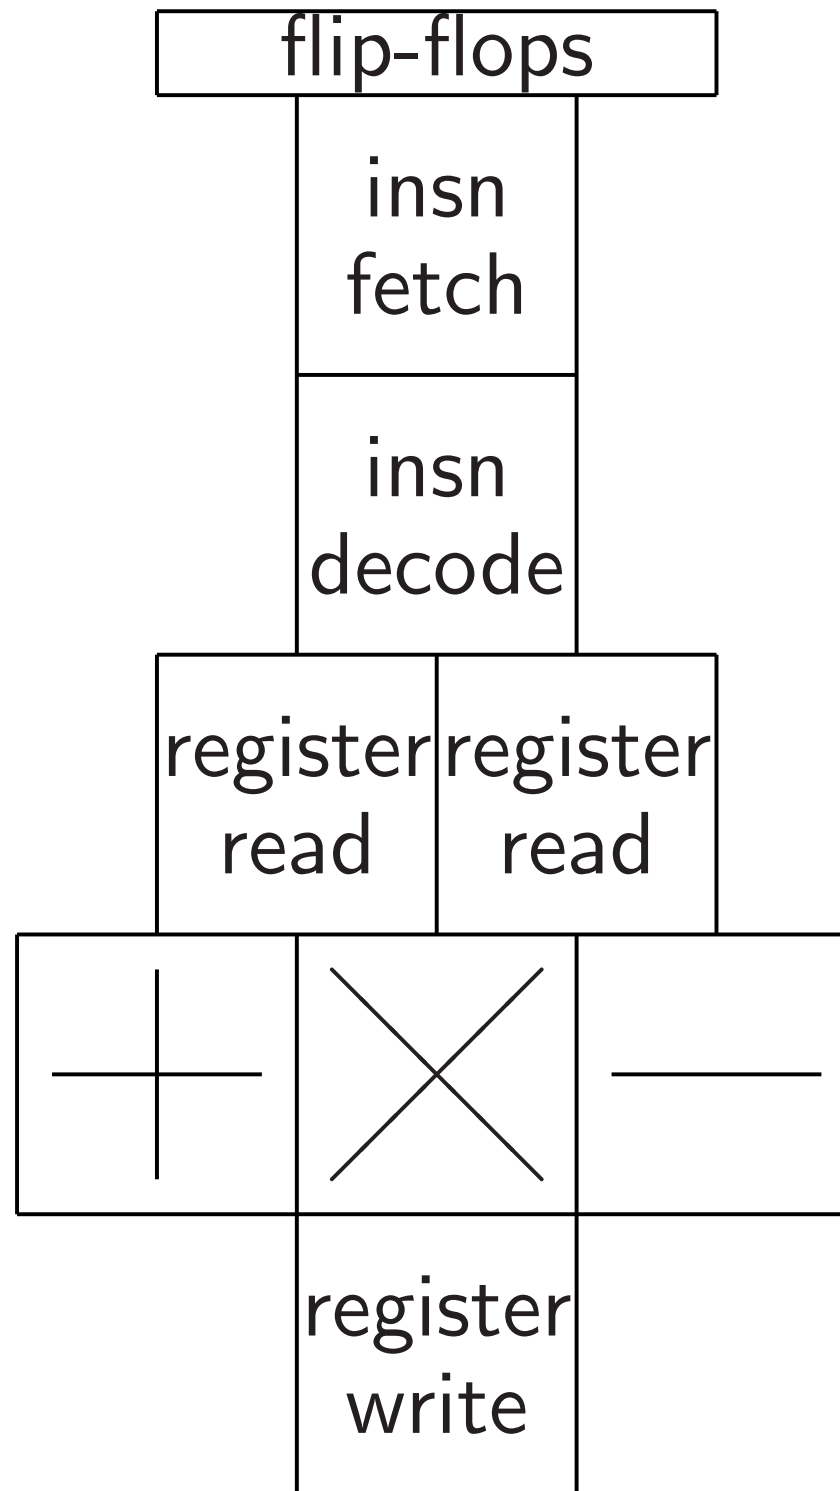
needs to be slow enough

ricity to percolate

way through the circuit,

o-flops to flip-flops.

Now have semi-flexible CPU:

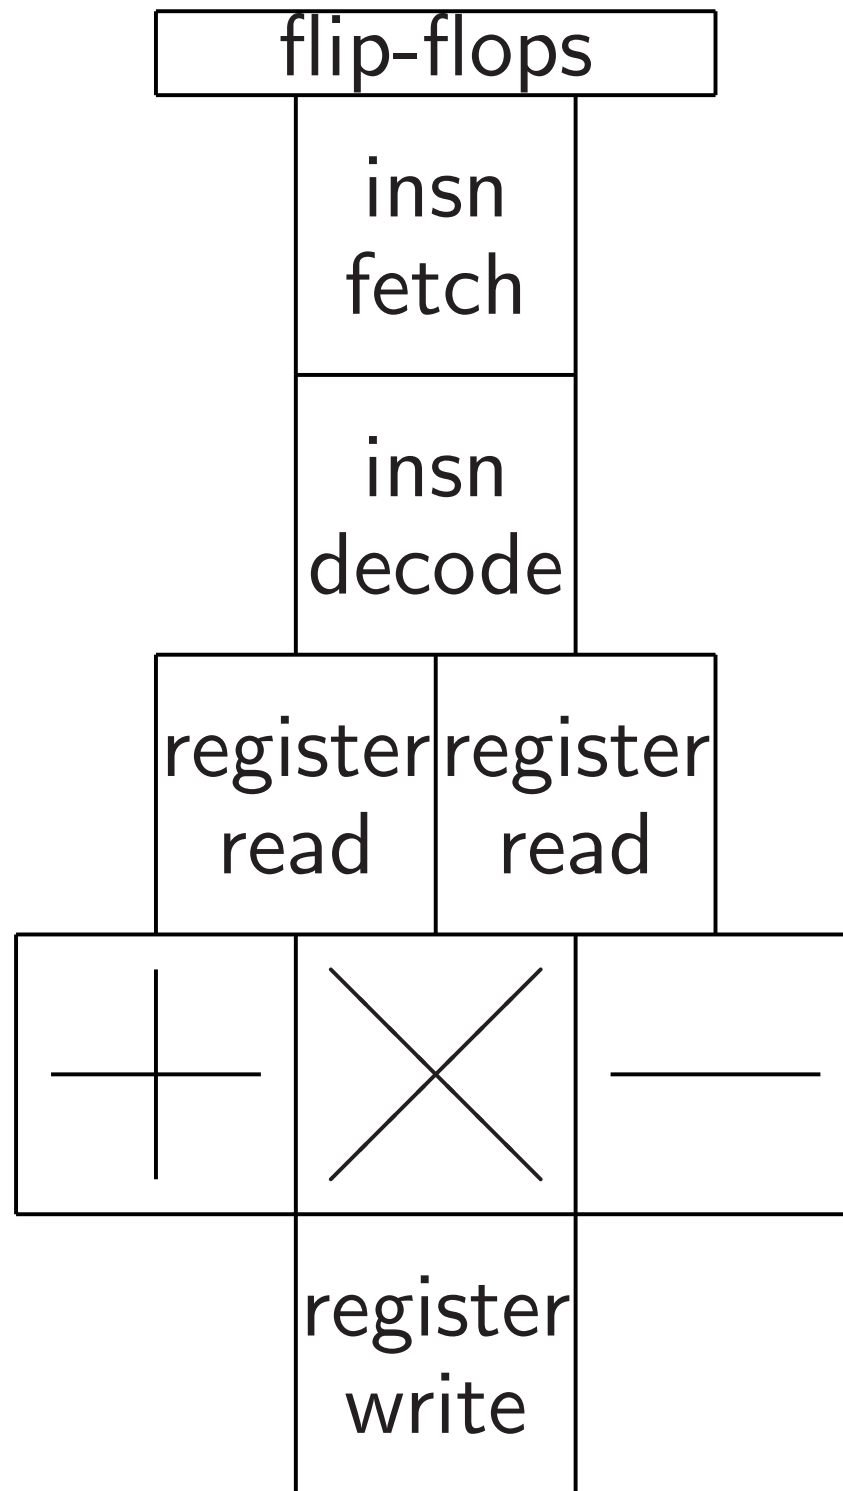


Further flexibility is useful:
e.g., rotation instructions.

“Pipelined

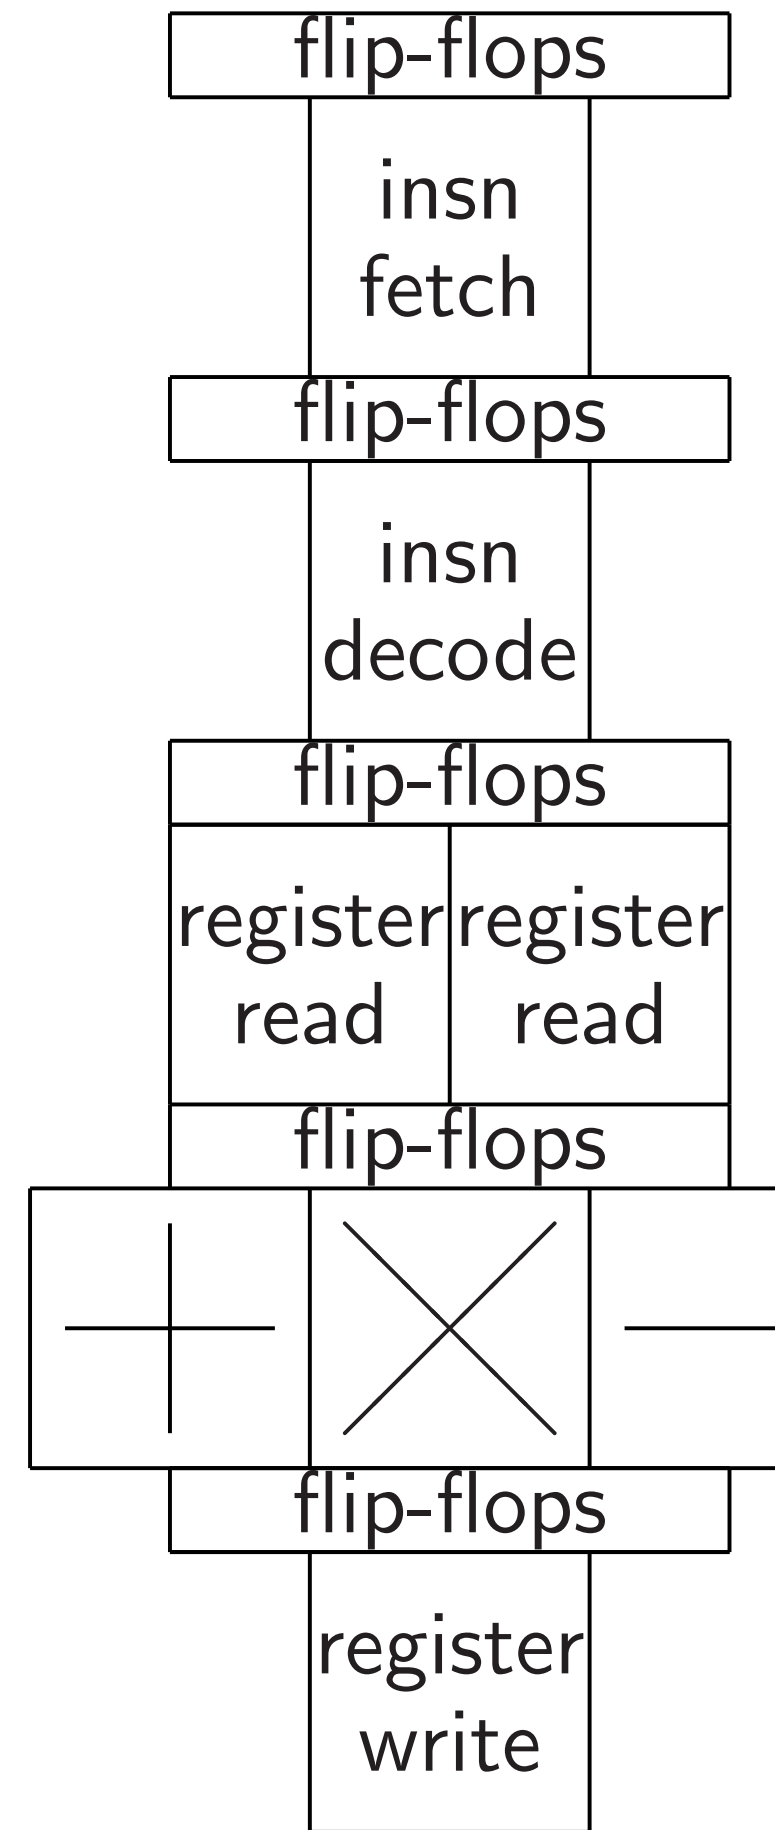


Now have semi-flexible CPU:

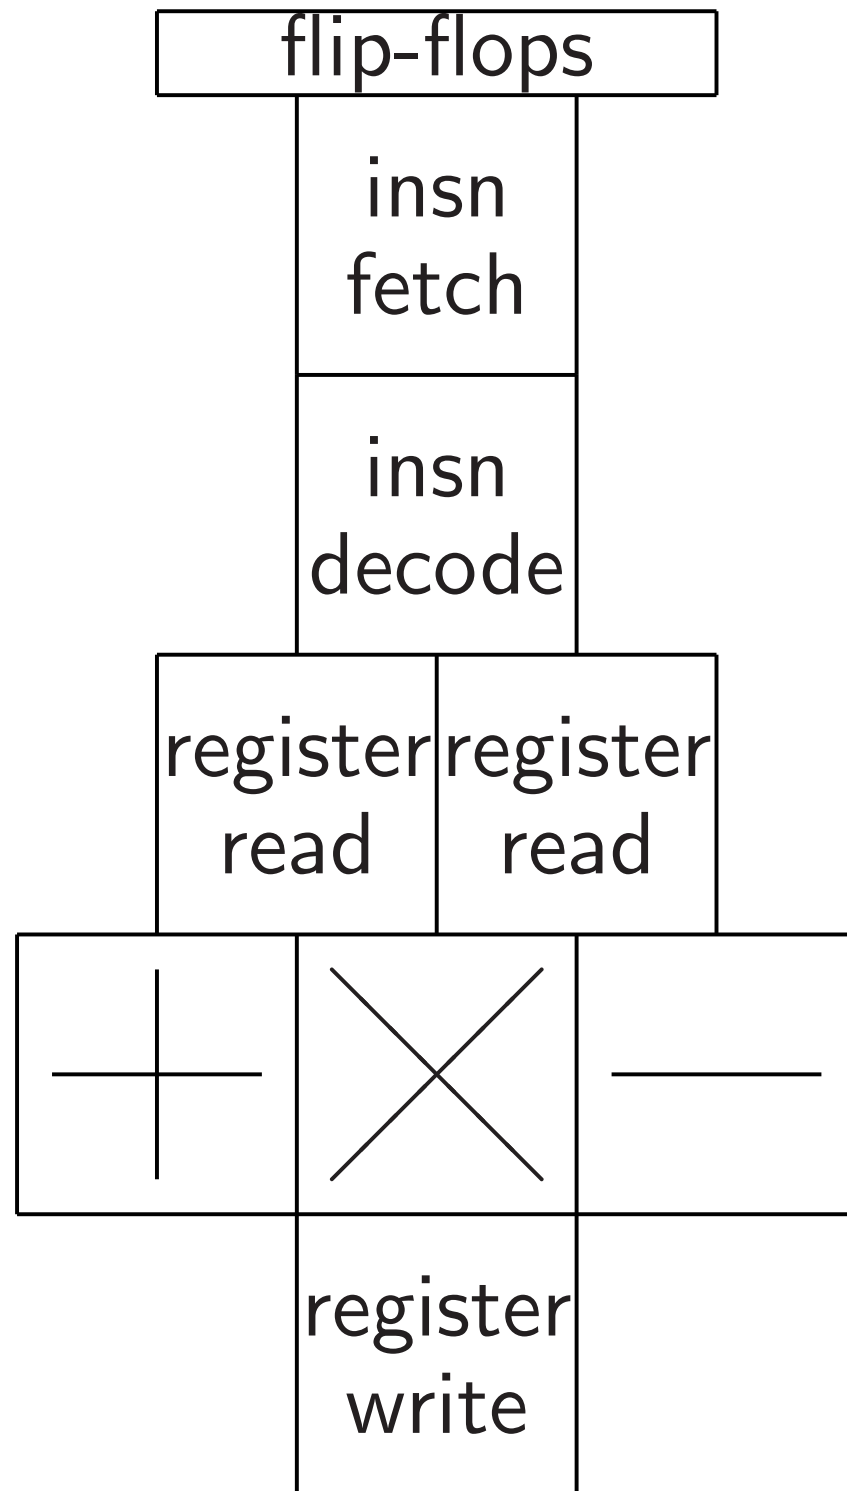


Further flexibility is useful:
e.g., rotation instructions.

“Pipelining” allow

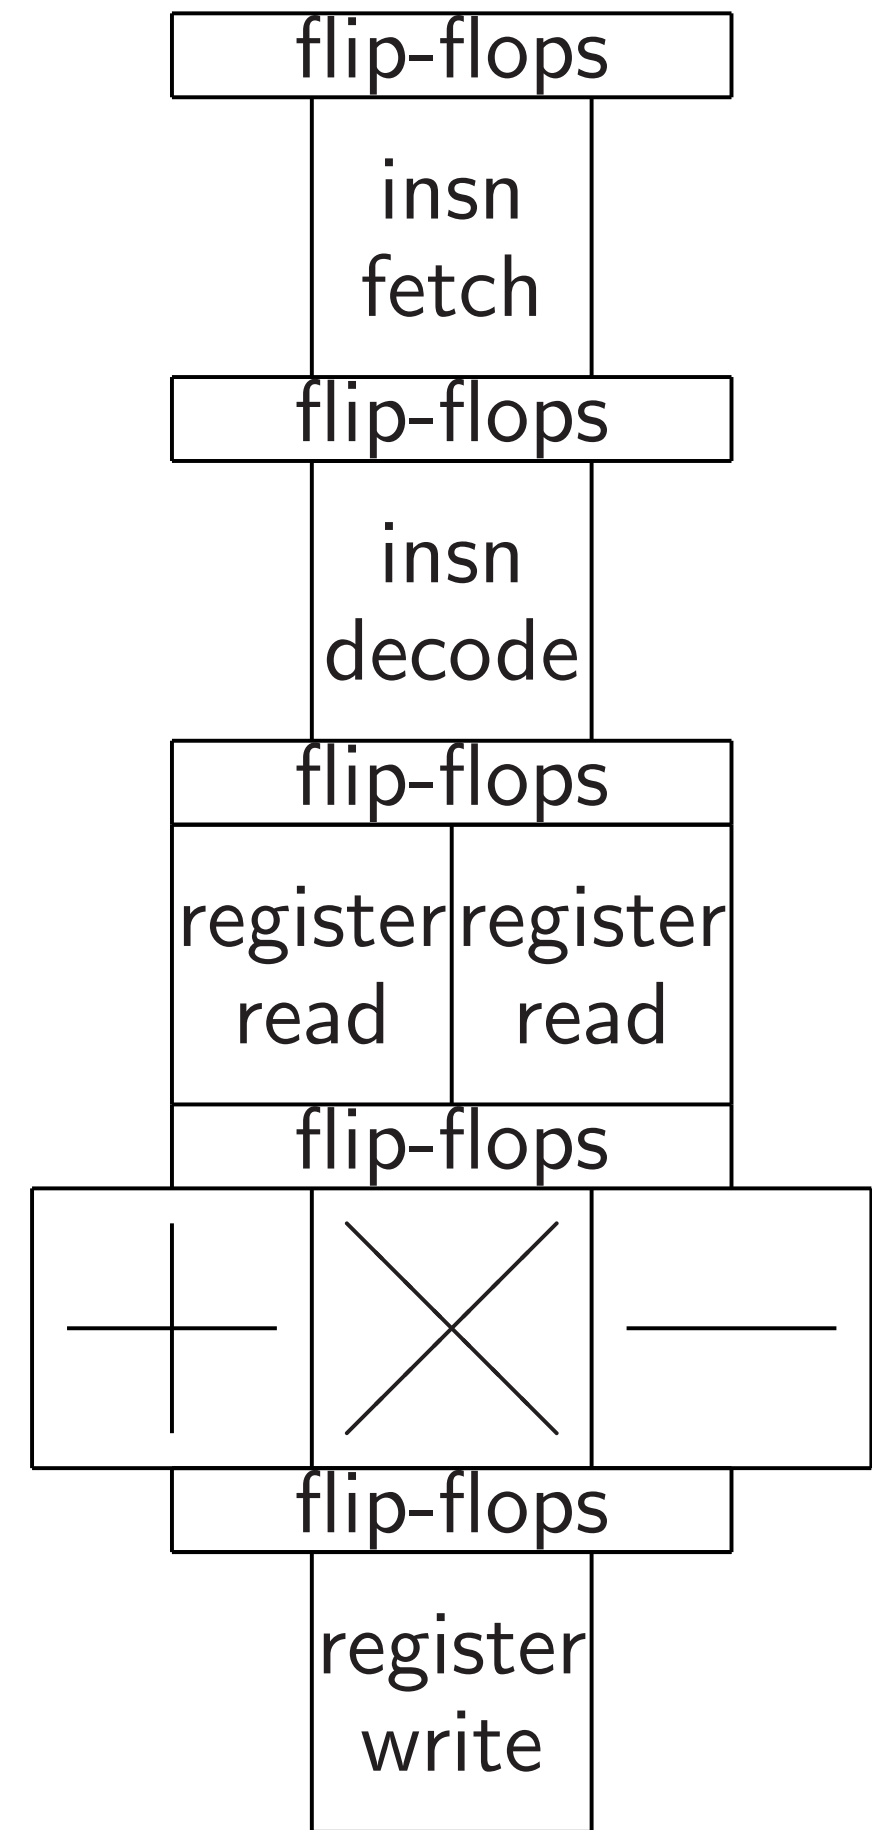


Now have semi-flexible CPU:



Further flexibility is useful:
e.g., rotation instructions.

“Pipelining” allows faster clock



stage

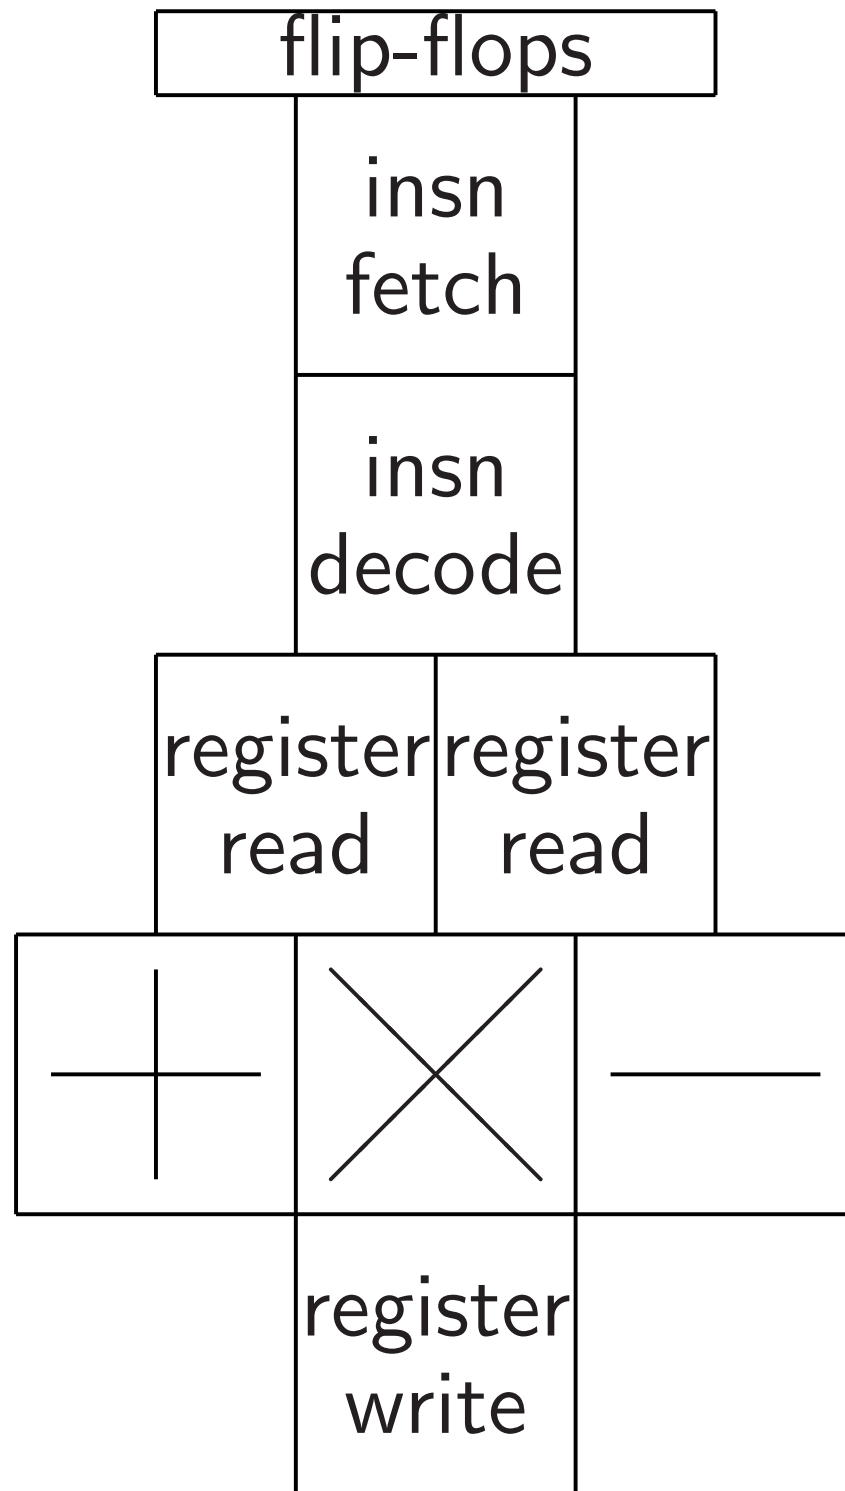
stage

stage

stage

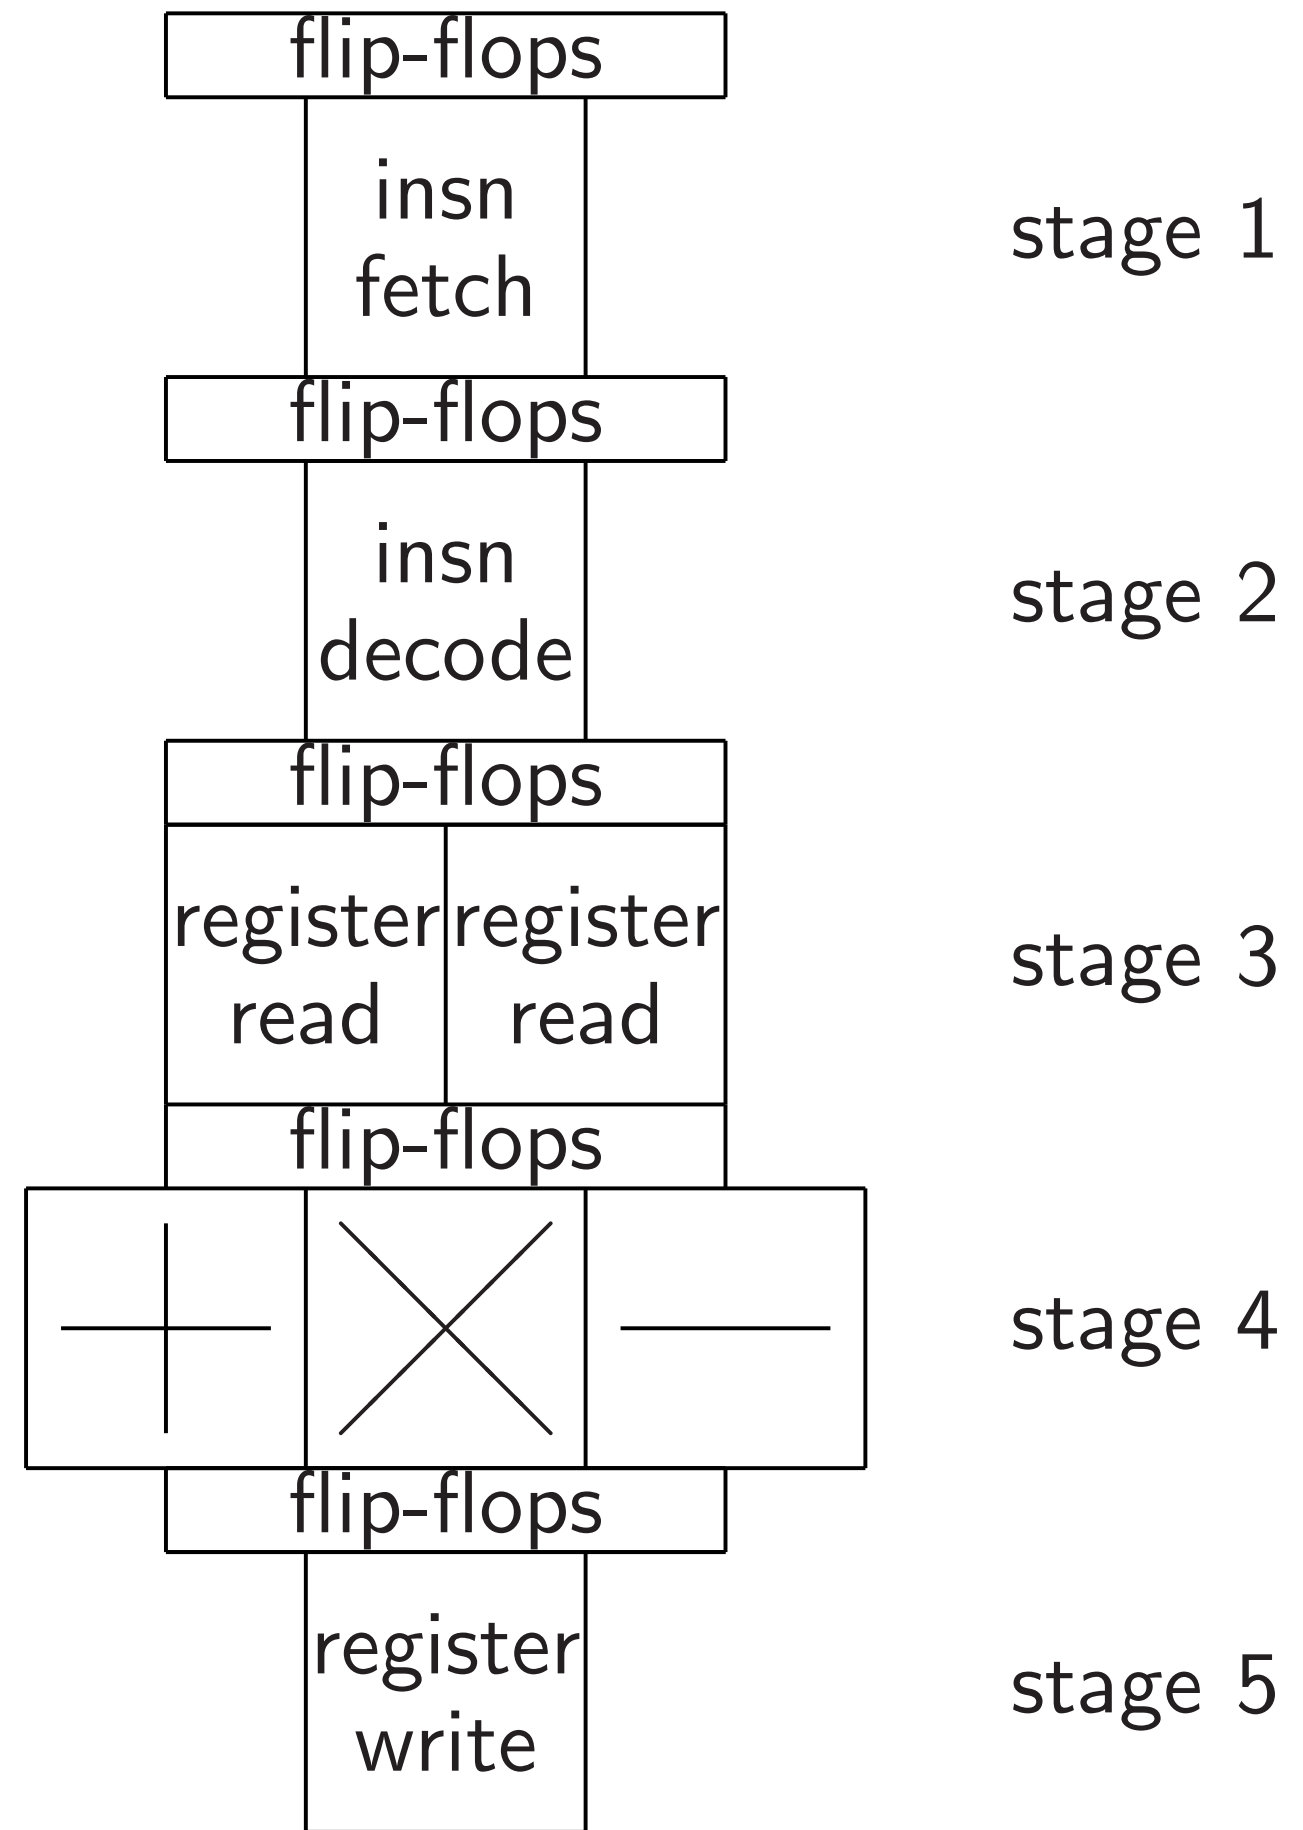
stage

Now have semi-flexible CPU:

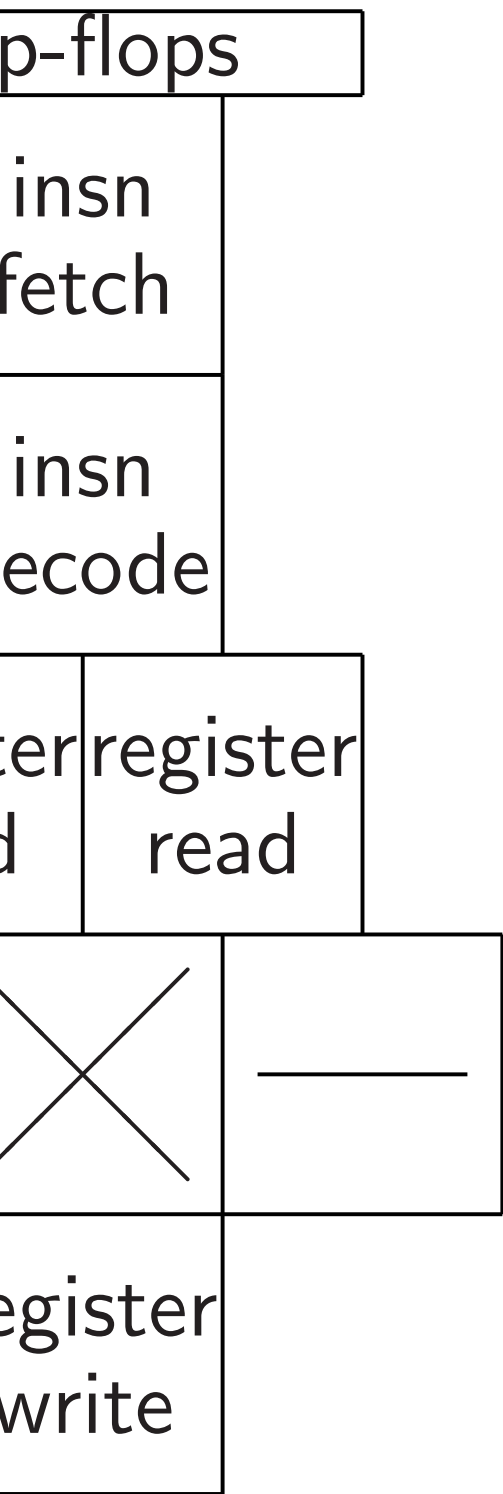


Further flexibility is useful:
e.g., rotation instructions.

“Pipelining” allows faster clock:

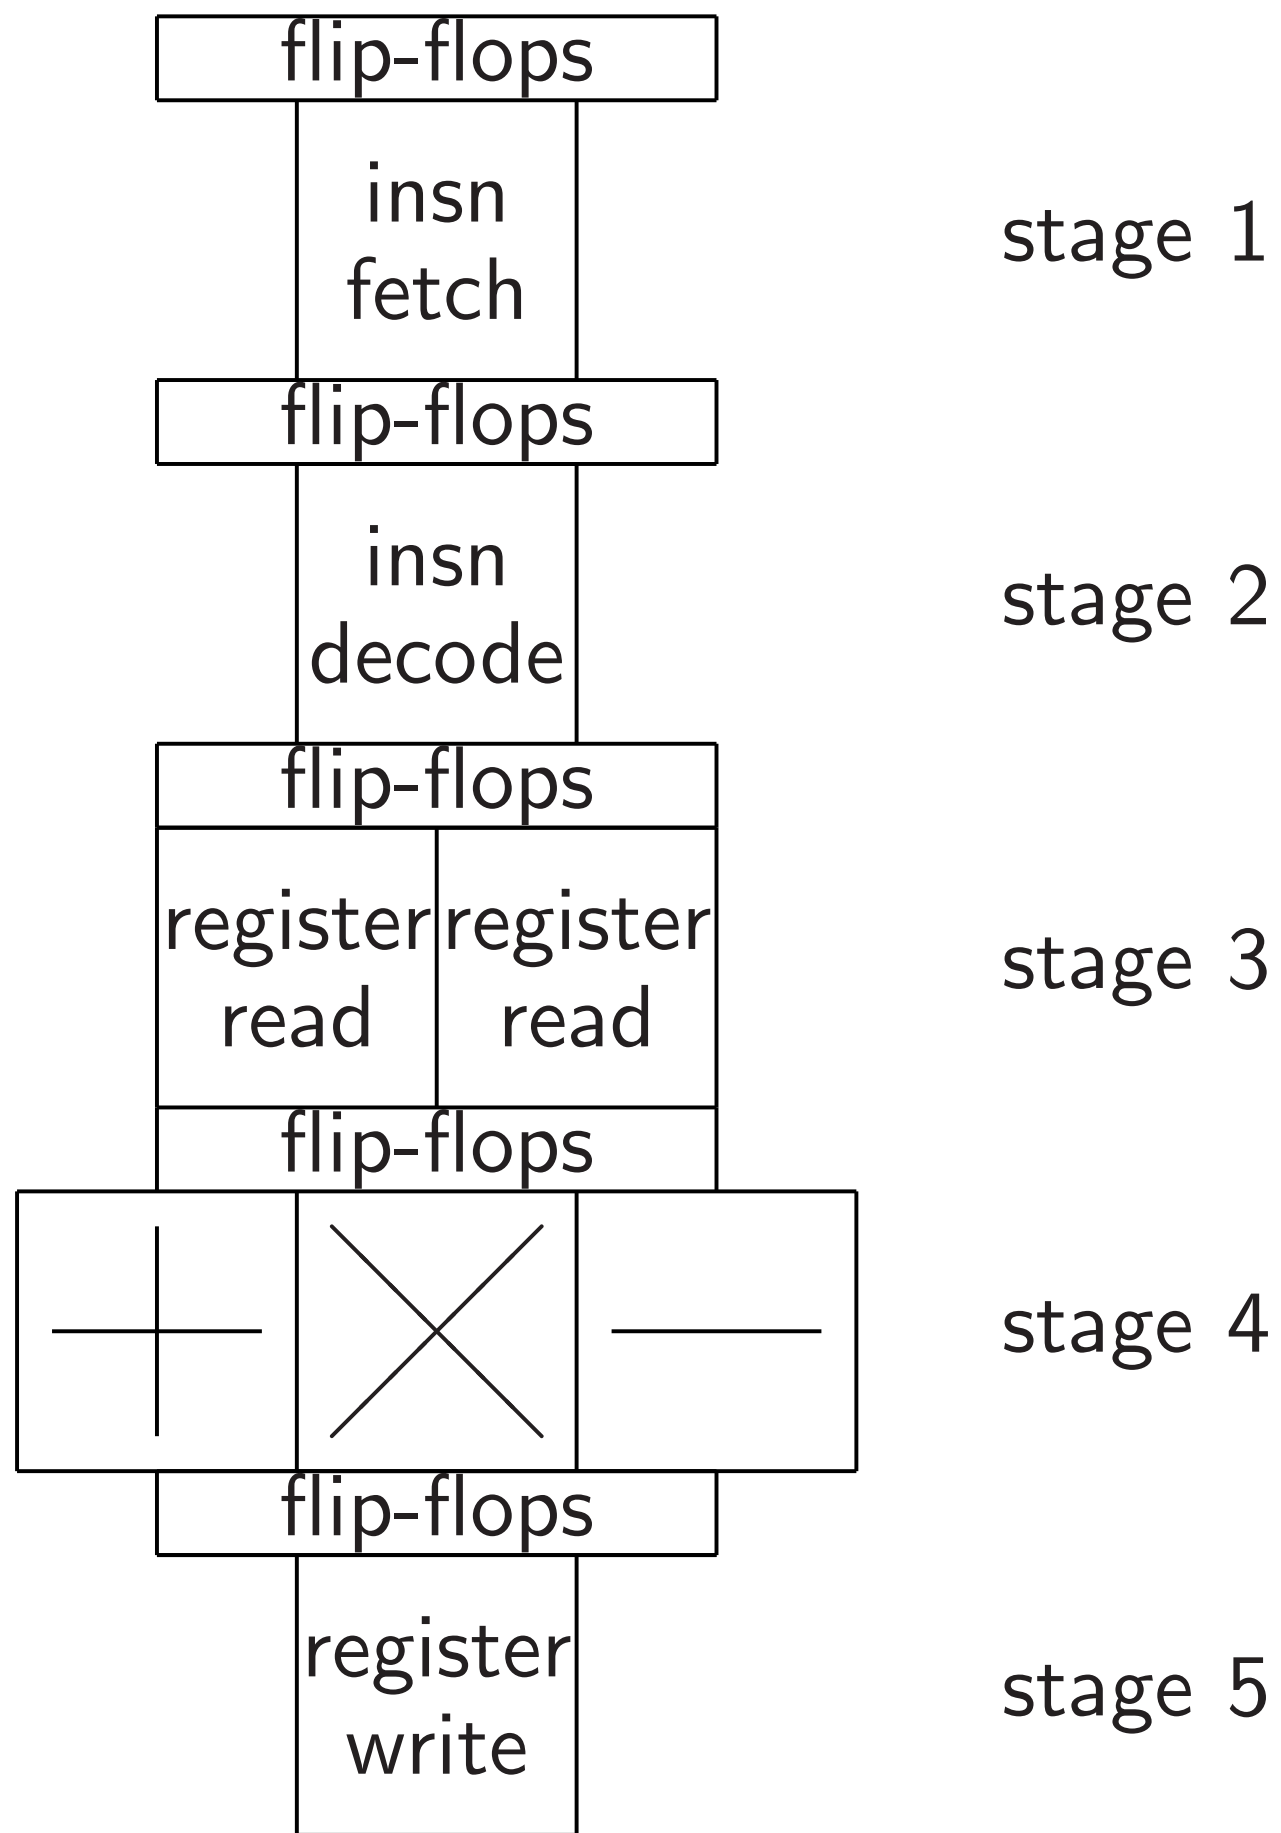


ve semi-flexible CPU:



flexibility is useful:
ation instructions.

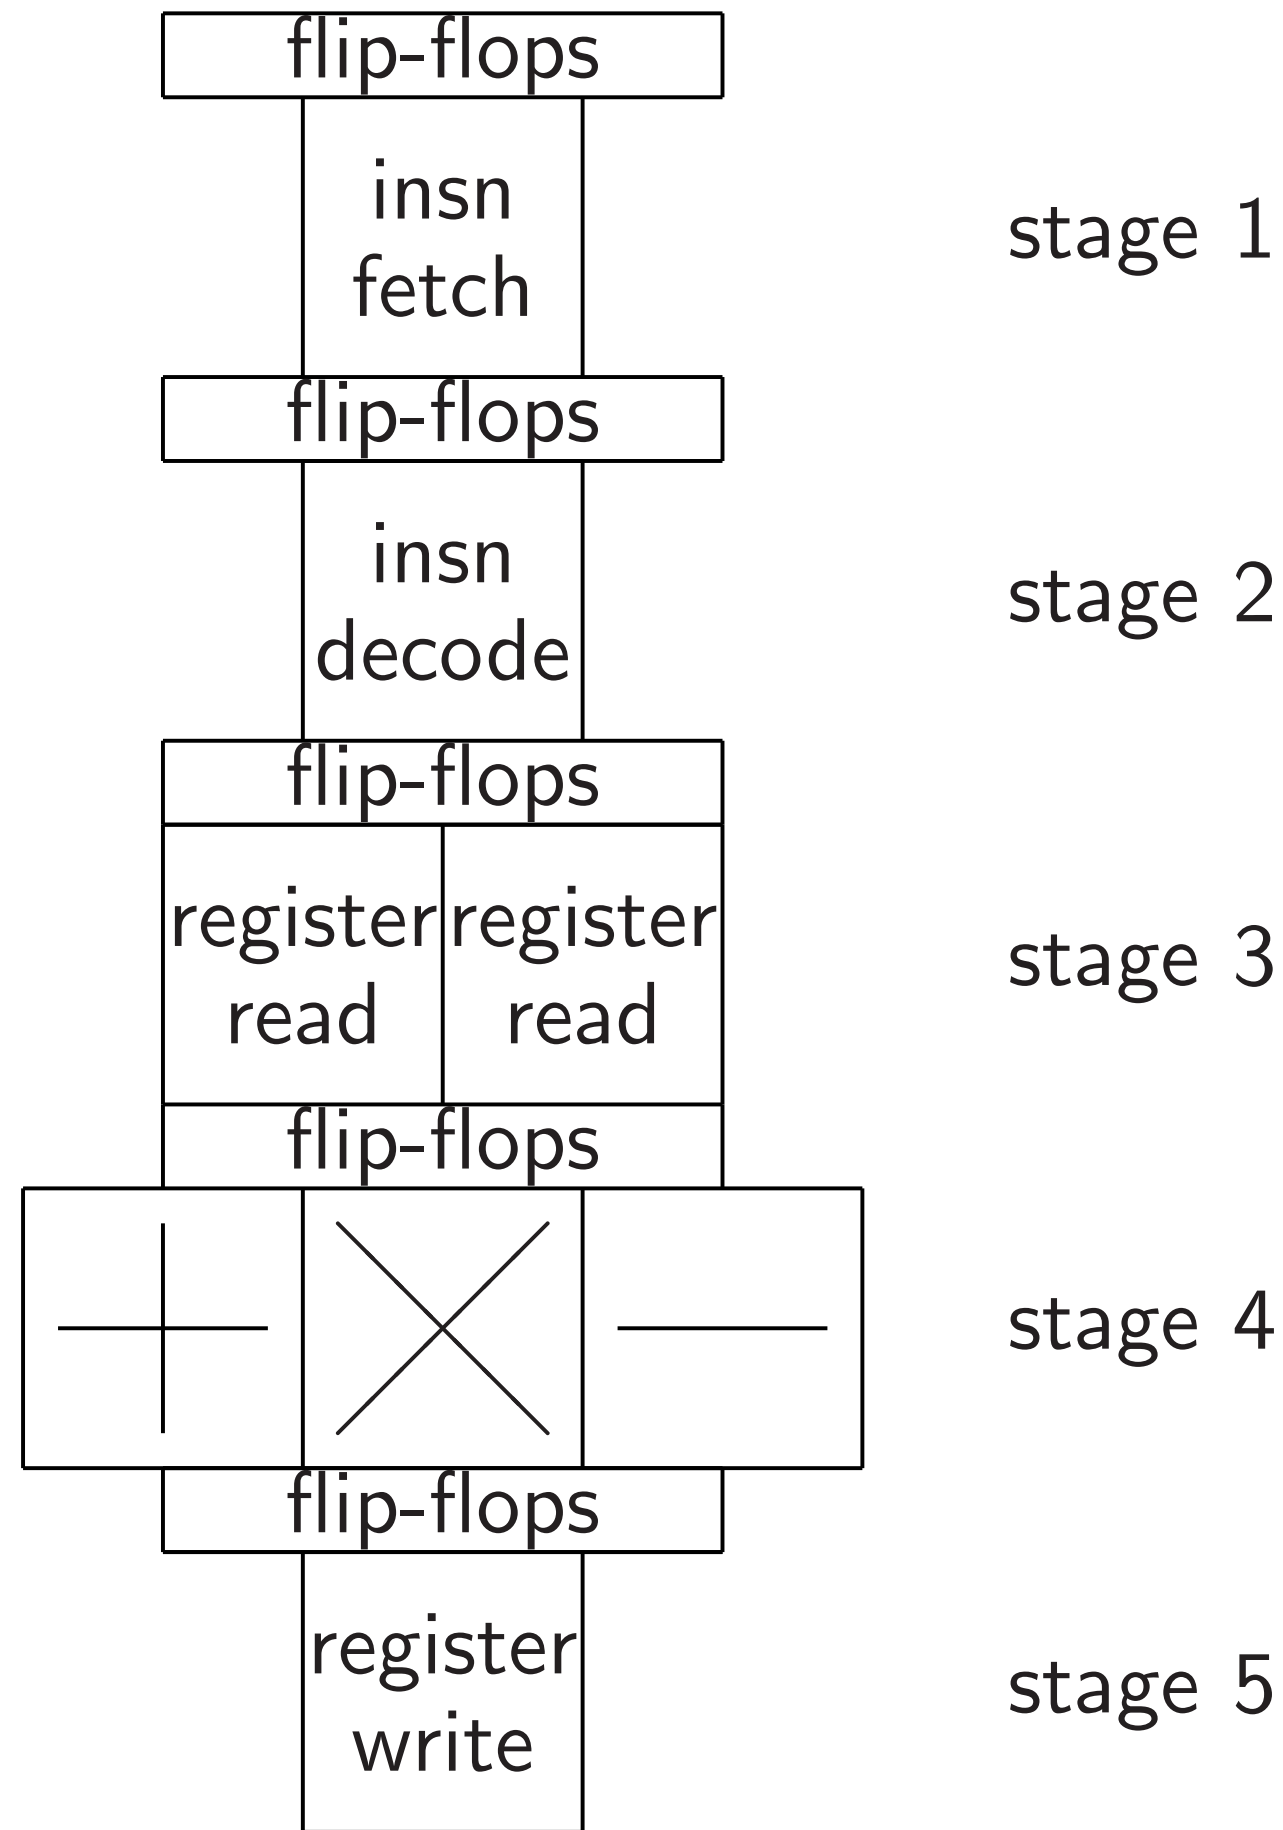
“Pipelining” allows faster clock:



Goal: St
one tick
Instructi
reads ne
feeds p'
After ne
instructi
uncompr
while ins
reads an
Some ex
Also ext
preserve
e.g., sta

flexible CPU:

“Pipelining” allows faster clock:



Goal: Stage n handles one tick after stage $n-1$

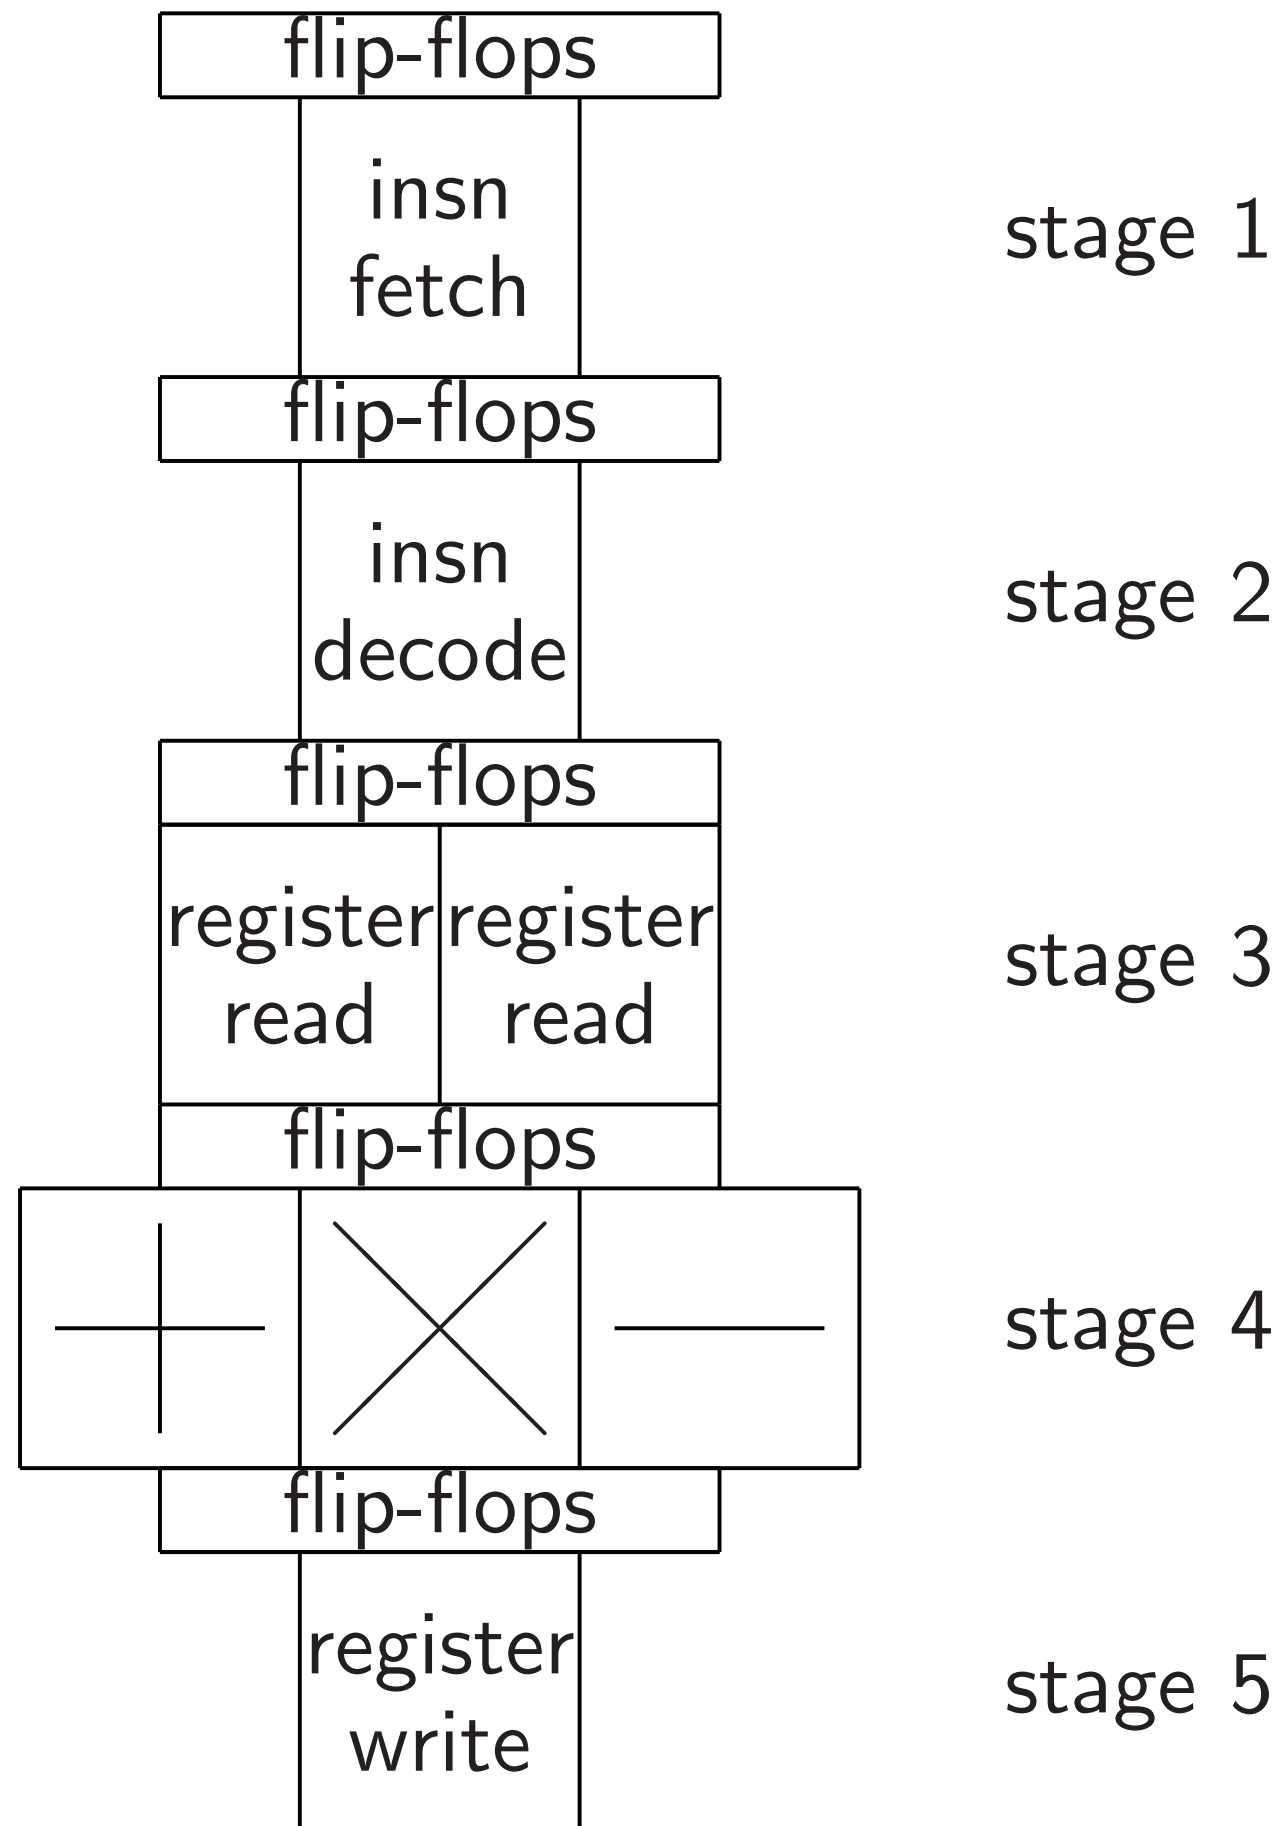
Instruction fetch reads next instruction, feeds p' back, sends p' to stage 2

After next clock tick, instruction decode reads next instruction, uncompresses this instruction, while instruction fetch reads another instruction, feeds p' back, sends p' to stage 2

Some extra flip-flops. Also extra area to preserve instructions. e.g., stall on read-

is useful:
instructions.

“Pipelining” allows faster clock:



Goal: Stage n handles instruction one tick after stage $n - 1$.

Instruction fetch

reads next instruction, feeds p' back, sends instruction

After next clock tick,

instruction decode

uncompresses this instruction

while instruction fetch

reads another instruction.

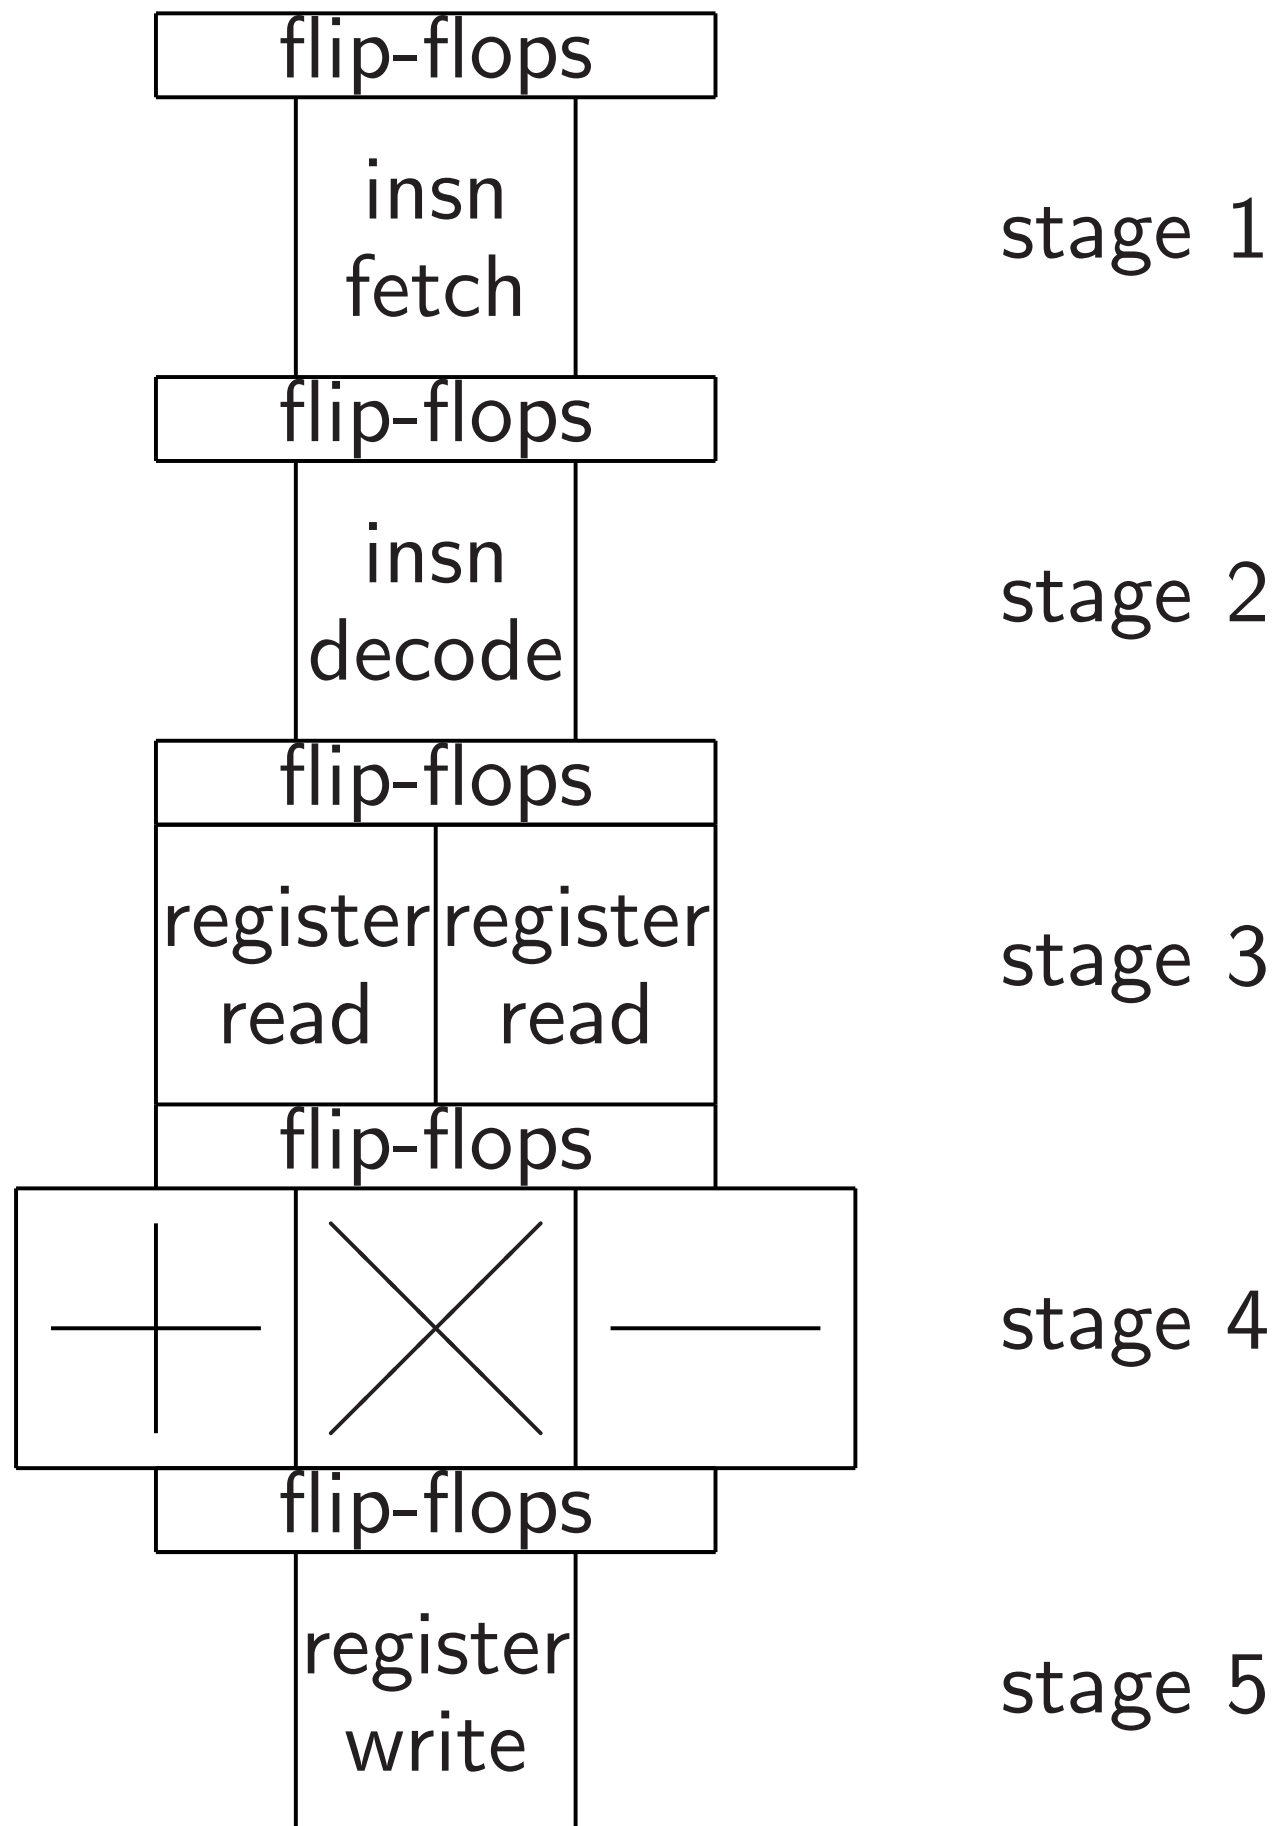
Some extra flip-flop area.

Also extra area to

preserve instruction semantics

e.g., stall on read-after-write

“Pipelining” allows faster clock:



Goal: Stage n handles instruction one tick after stage $n - 1$.

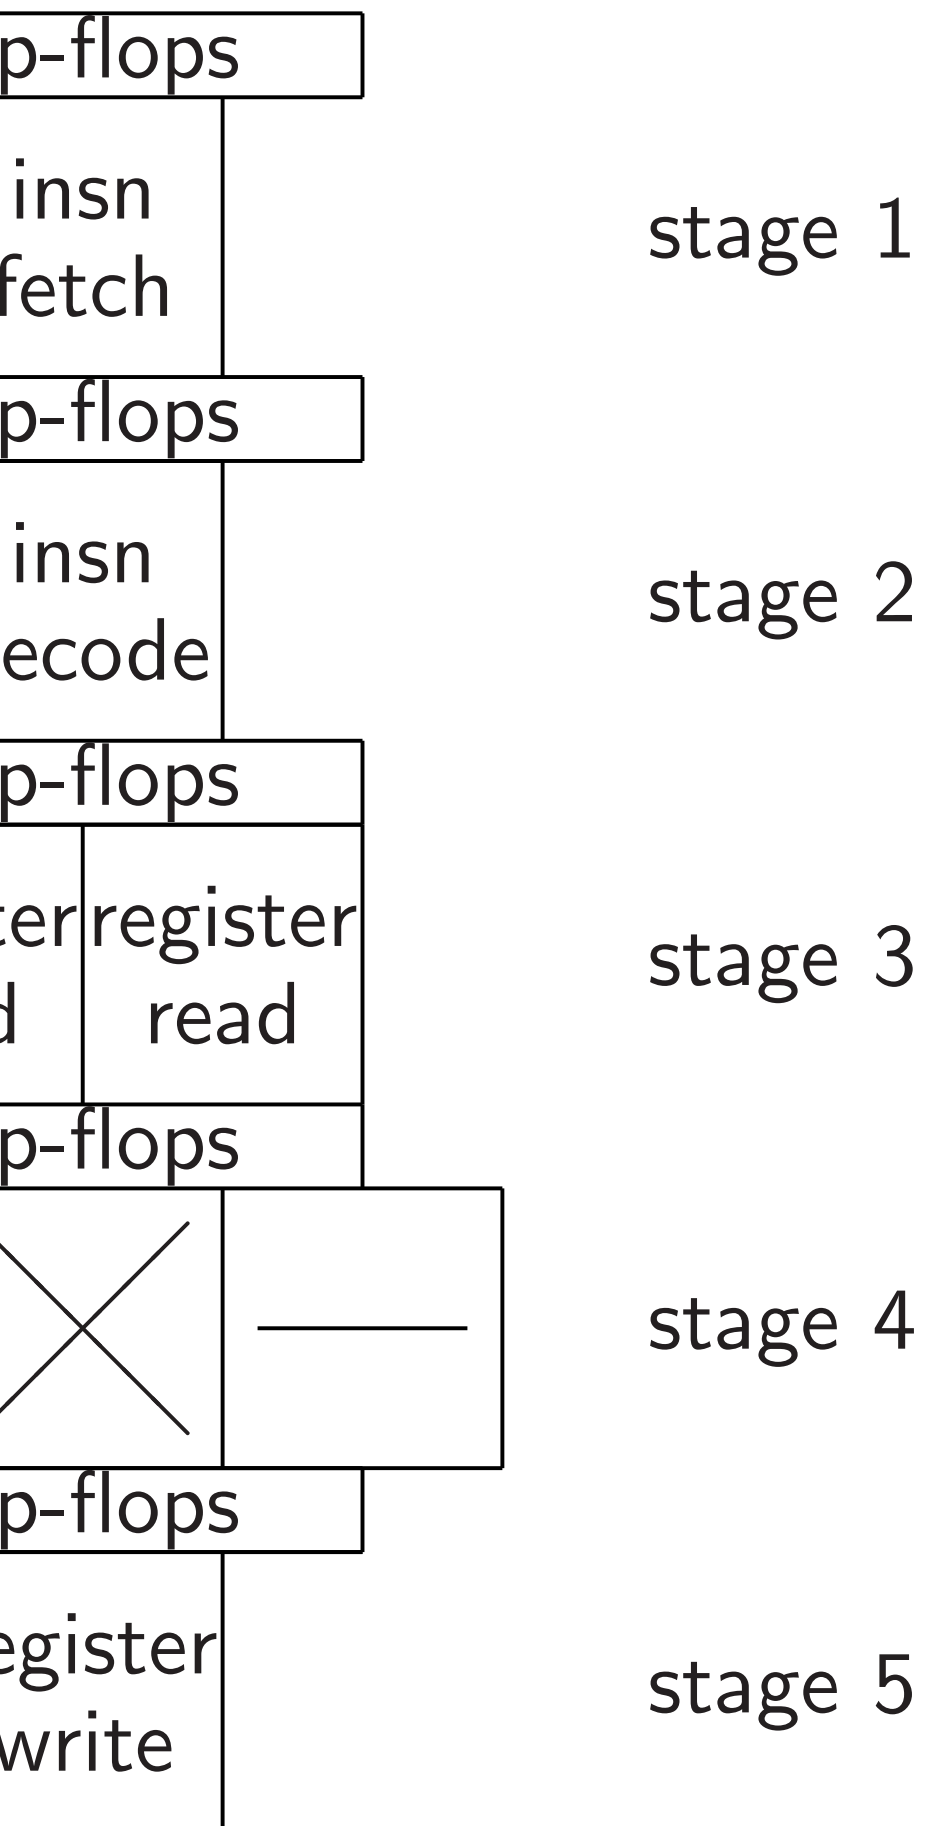
Instruction fetch reads next instruction, feeds p' back, sends instruction.

After next clock tick, instruction decode uncompresses this instruction, while instruction fetch reads another instruction.

Some extra flip-flop area.

Also extra area to preserve instruction semantics: e.g., stall on read-after-write.

“pipelining” allows faster clock:



Goal: Stage n handles instruction one tick after stage $n - 1$.

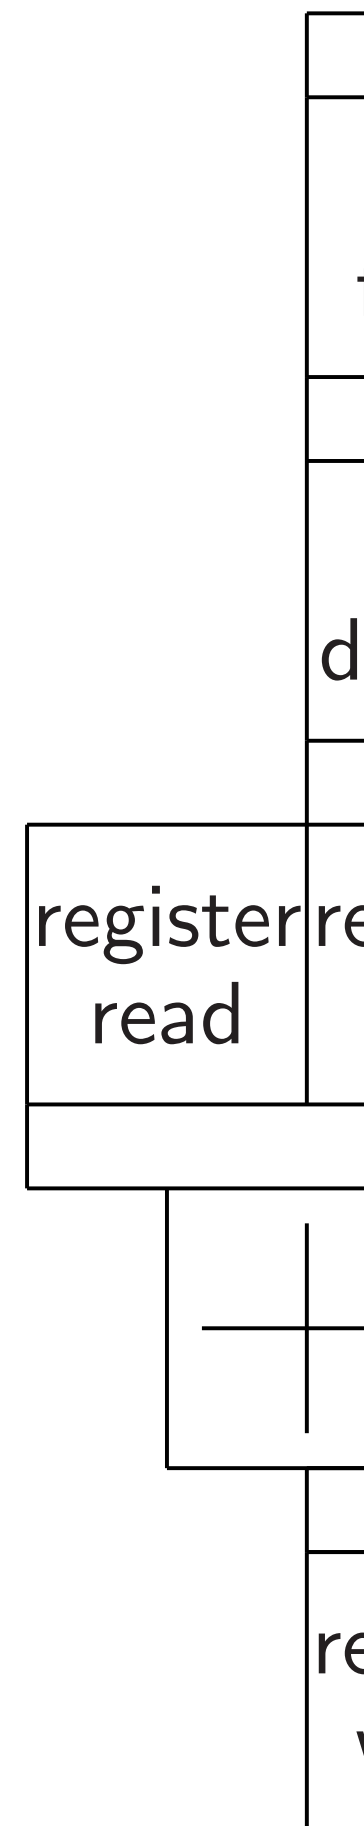
Instruction fetch
reads next instruction,
feeds p' back, sends instruction.

After next clock tick,
instruction decode
uncompresses this instruction,
while instruction fetch
reads another instruction.

Some extra flip-flop area.

Also extra area to
preserve instruction semantics:
e.g., stall on read-after-write.

“Superscalar”



s faster clock:

stage 1

Goal: Stage n handles instruction one tick after stage $n - 1$.

Instruction fetch
reads next instruction,
feeds p' back, sends instruction.

stage 2

After next clock tick,
instruction decode

stage 3

uncompresses this instruction,
while instruction fetch
reads another instruction.

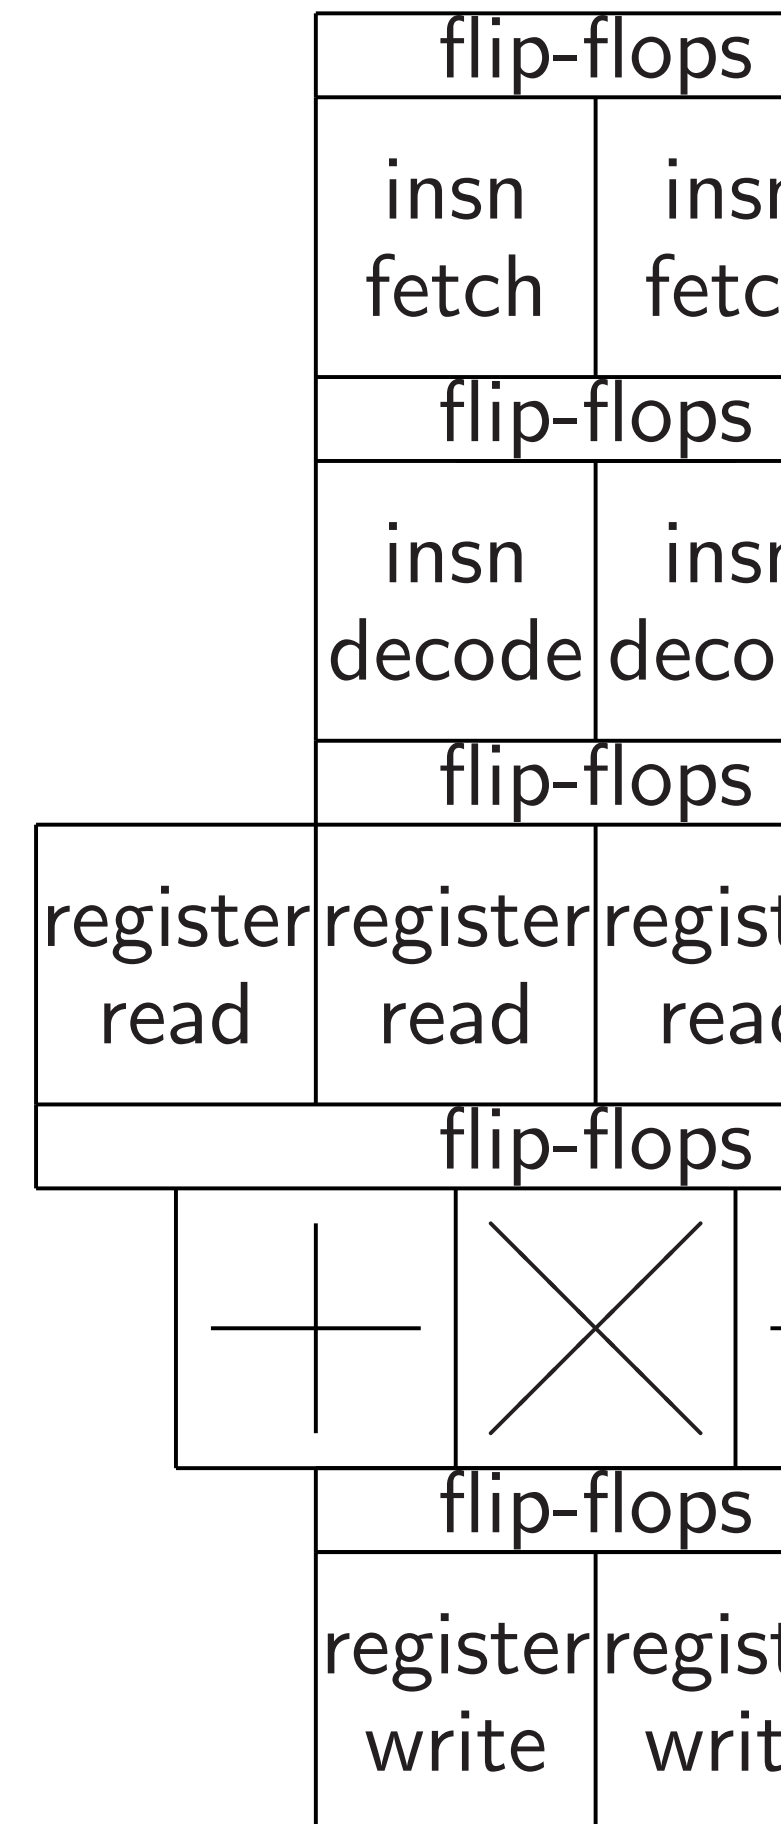
stage 4

Some extra flip-flop area.

stage 5

Also extra area to
preserve instruction semantics:
e.g., stall on read-after-write.

“Superscalar” proc



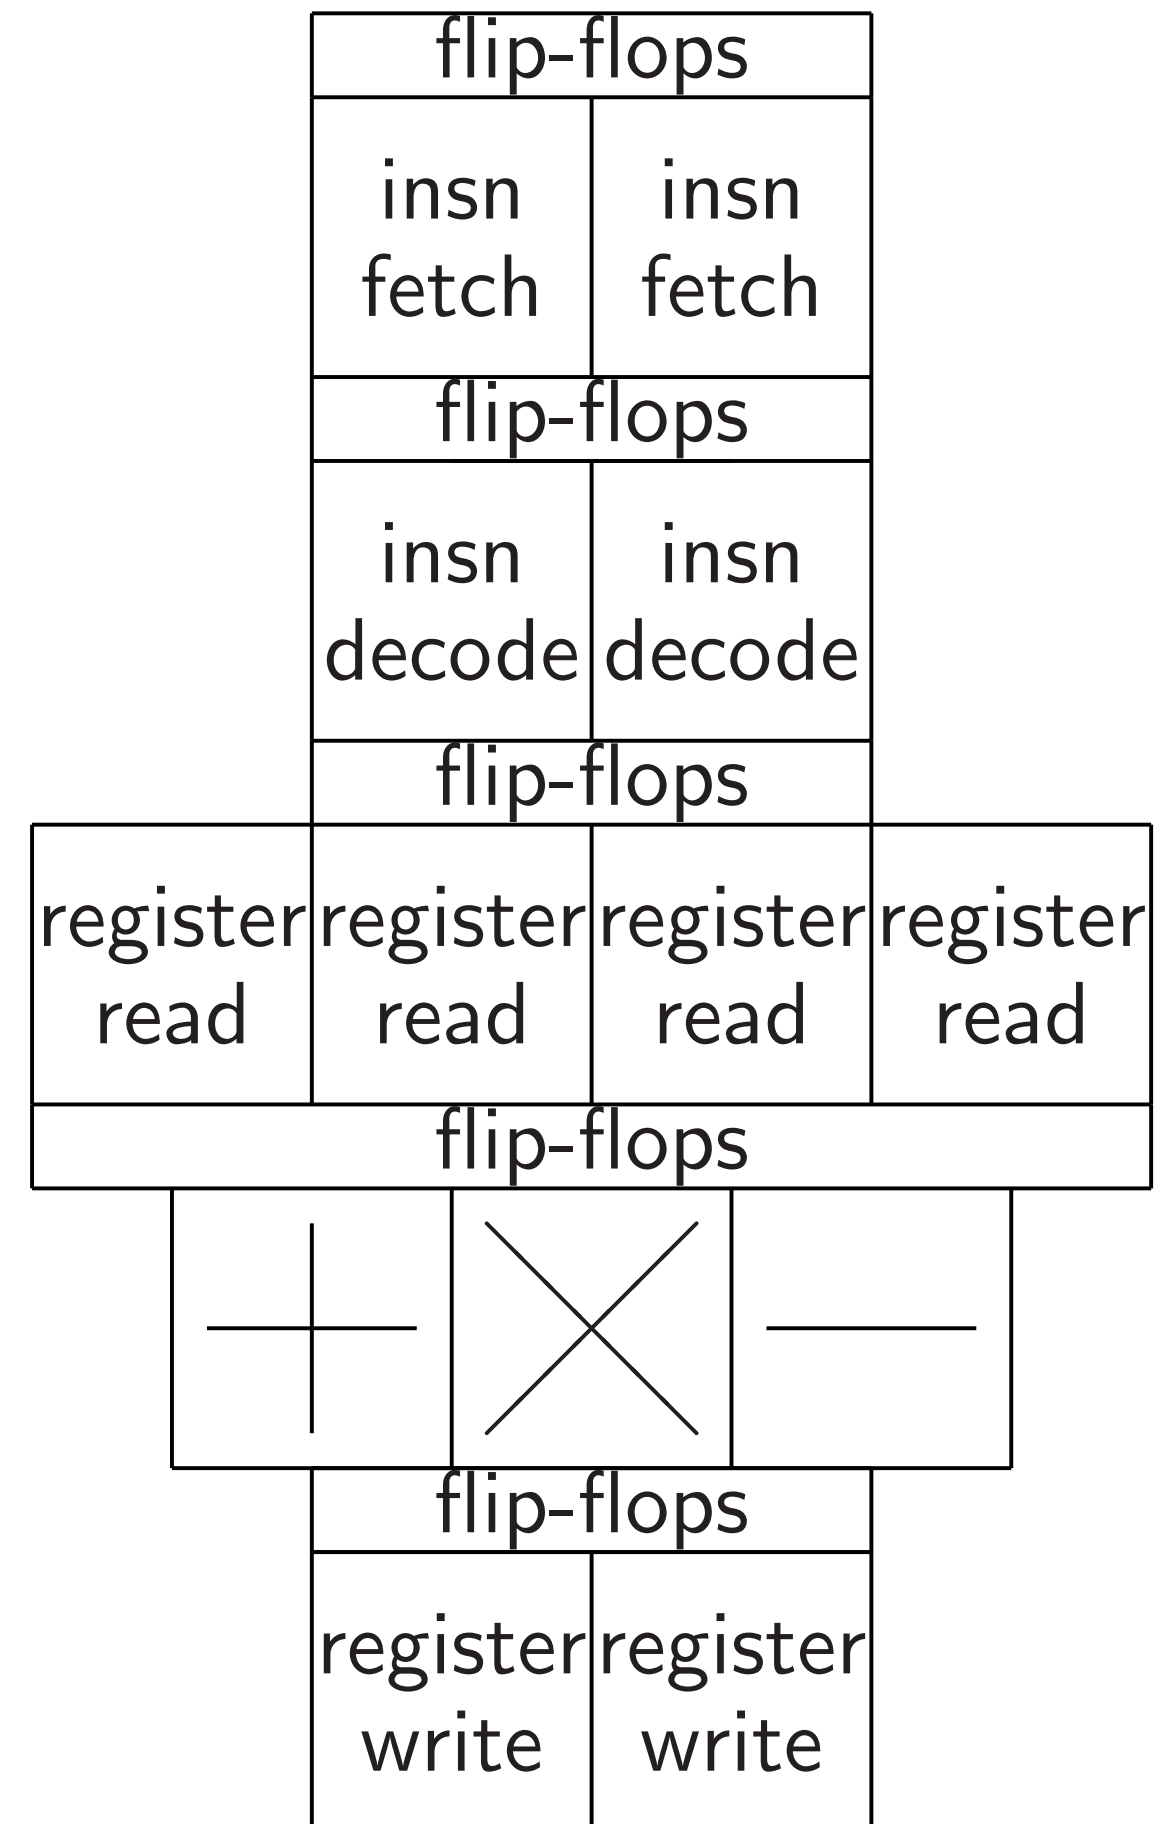
clock: Goal: Stage n handles instruction one tick after stage $n - 1$.

e 1 Instruction fetch
reads next instruction,
feeds p' back, sends instruction.

e 2 After next clock tick,
instruction decode
uncompresses this instruction,
while instruction fetch
reads another instruction.

e 3 Some extra flip-flop area.
Also extra area to
preserve instruction semantics:
e.g., stall on read-after-write.

“Superscalar” processing:



Goal: Stage n handles instruction one tick after stage $n - 1$.

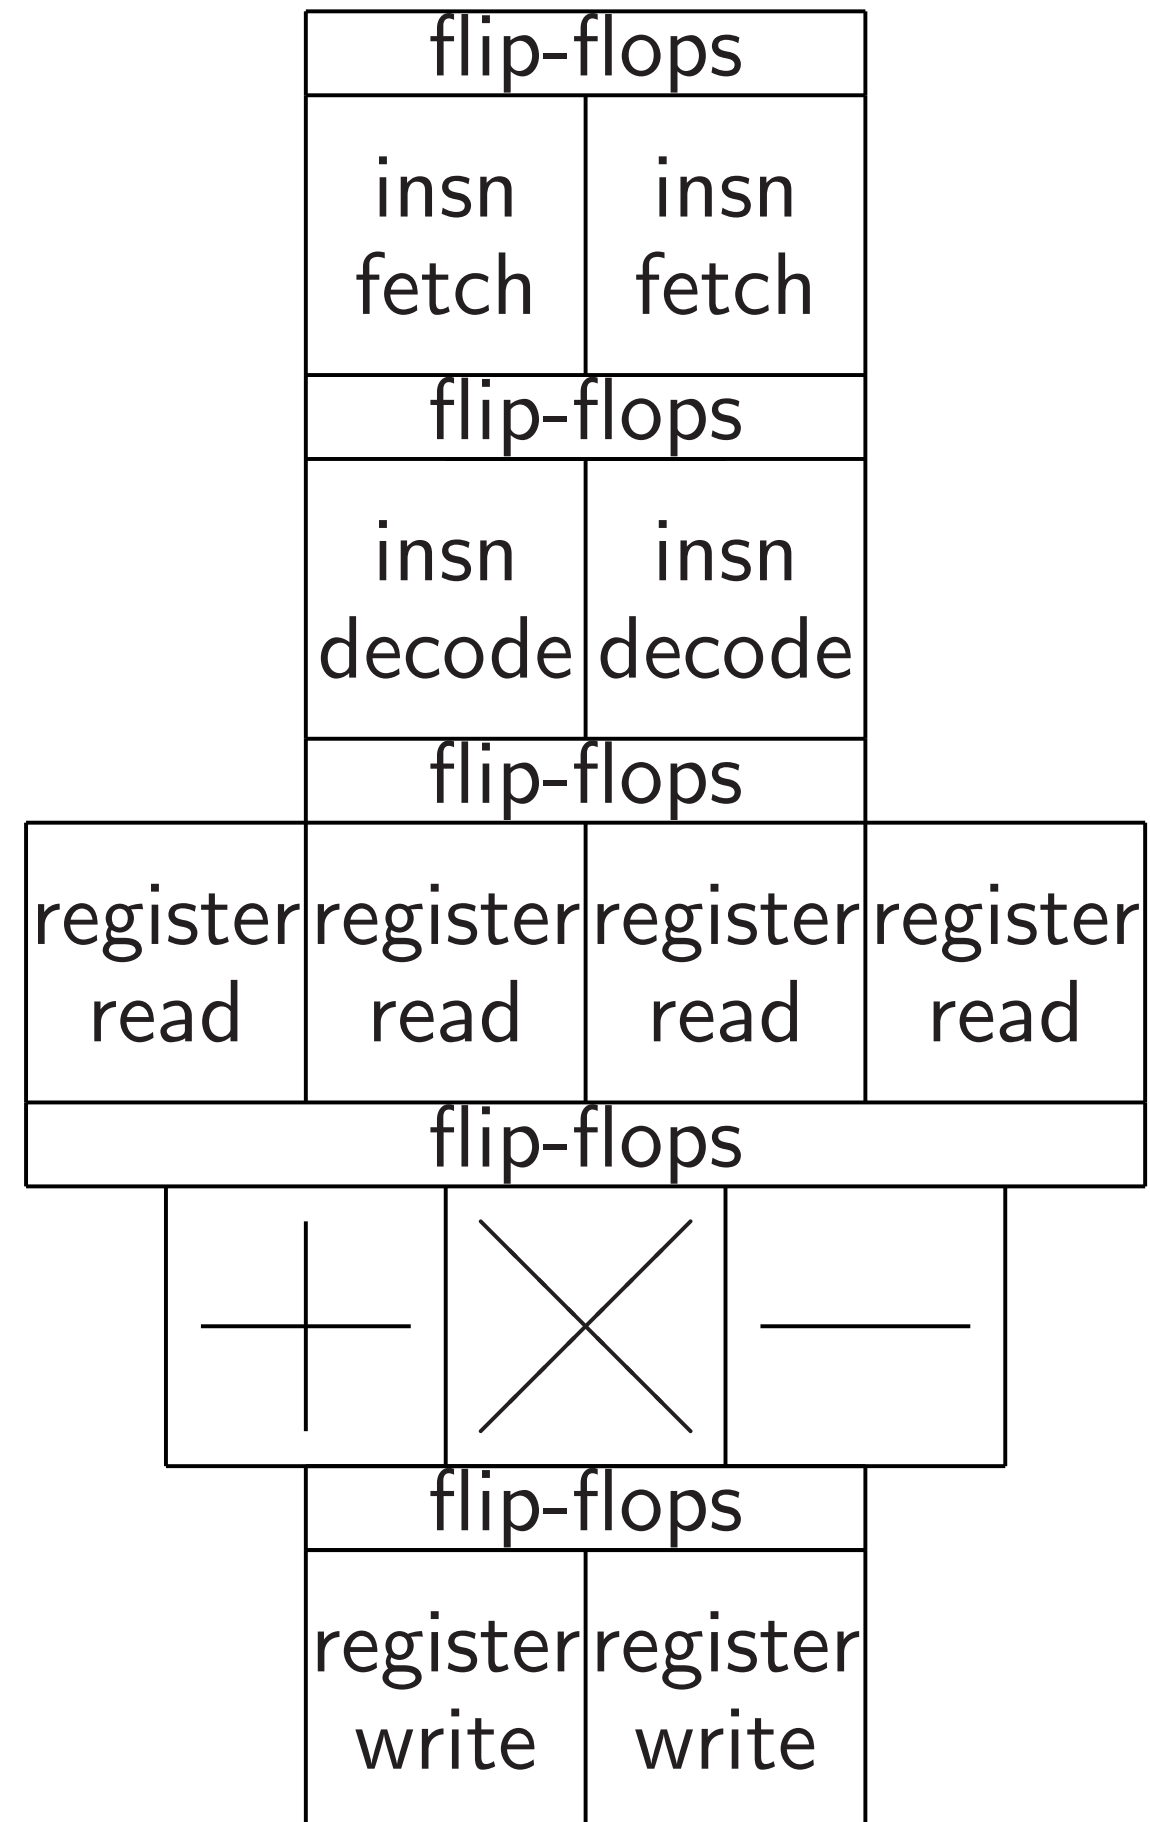
Instruction fetch
reads next instruction,
feeds p' back, sends instruction.

After next clock tick,
instruction decode
uncompresses this instruction,
while instruction fetch
reads another instruction.

Some extra flip-flop area.

Also extra area to
preserve instruction semantics:
e.g., stall on read-after-write.

“Superscalar” processing:



stage n handles instruction
after stage $n - 1$.

on fetch

xt instruction,

back, sends instruction.

xt clock tick,

on decode

resses this instruction,

struction fetch

other instruction.

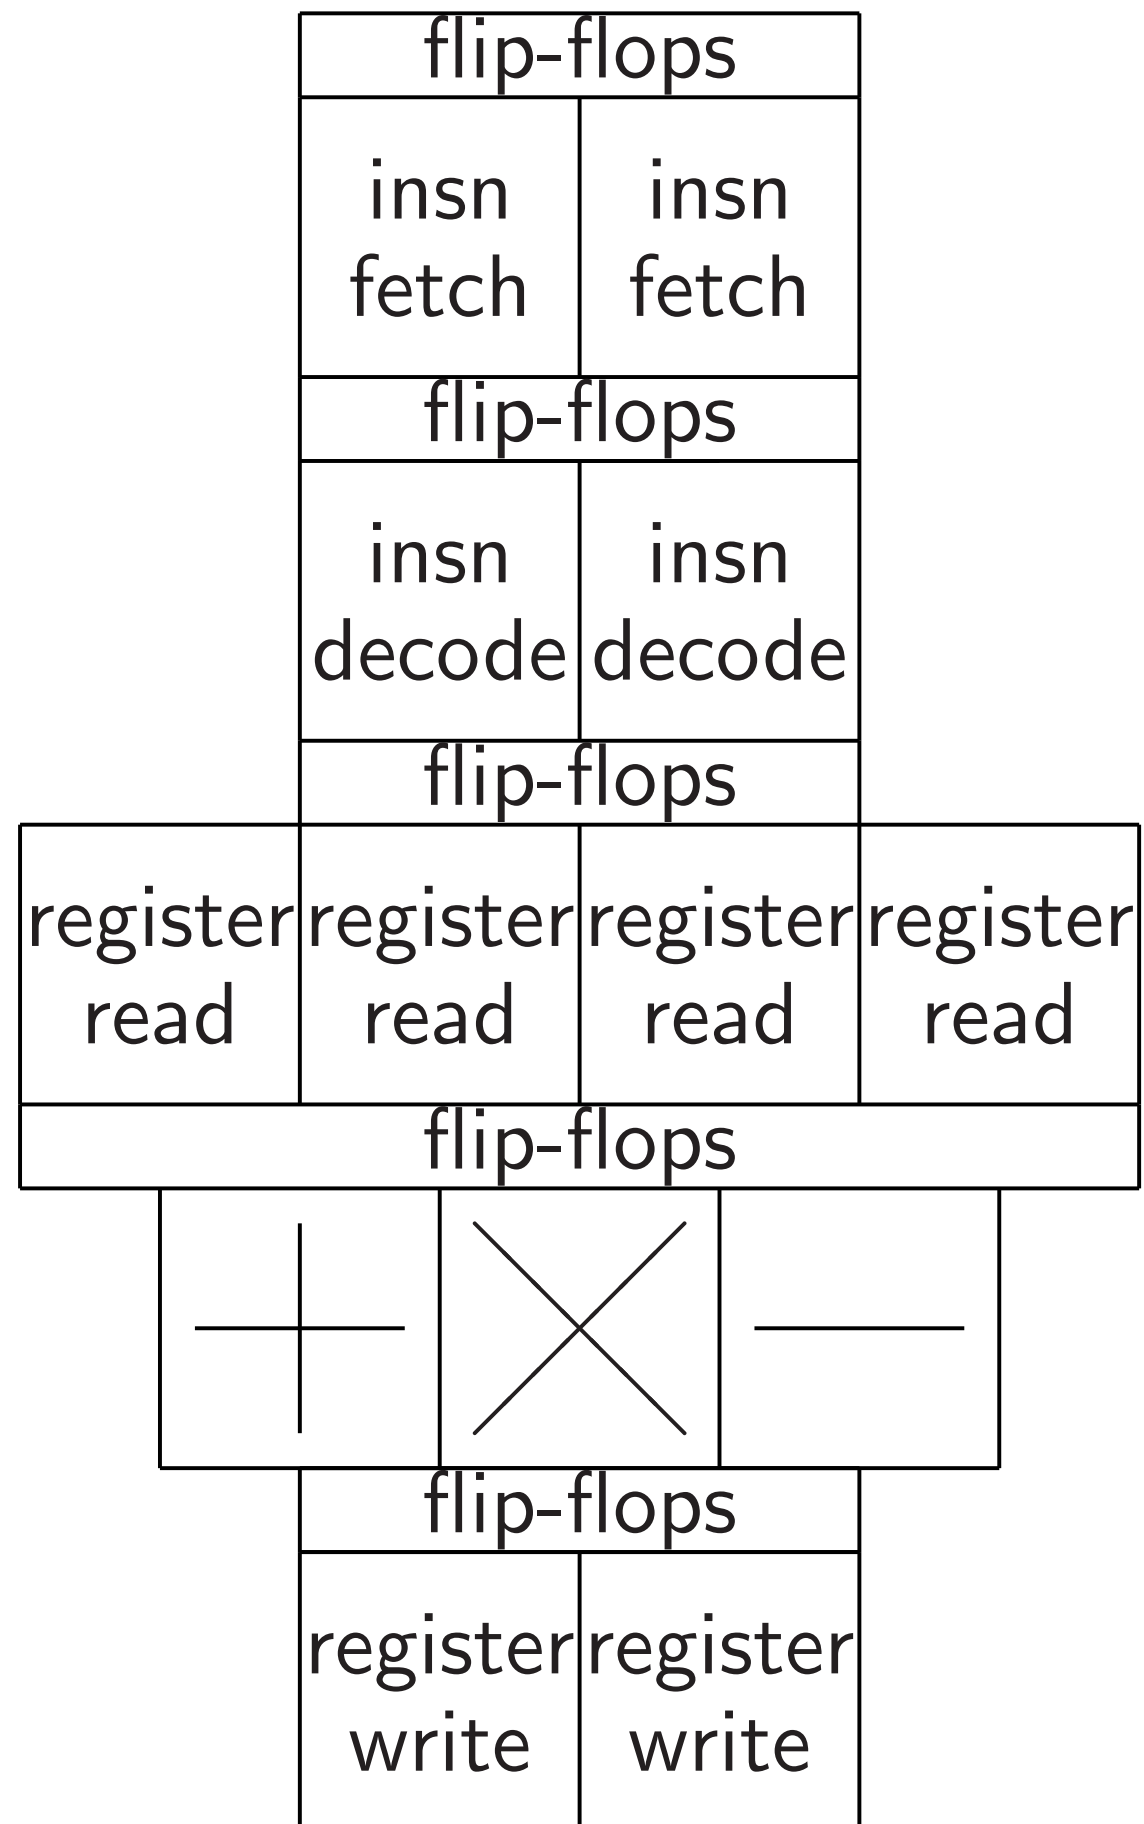
tra flip-flop area.

ra area to

instruction semantics:

ll on read-after-write.

“Superscalar” processing:



“Vector”

Expand

into n -ve

ARM “M

Intel “AV

Intel “AV

GPUs ha

handles instruction
stage $n - 1$.

tion,
ds instruction.

ck,

instruction,

etch

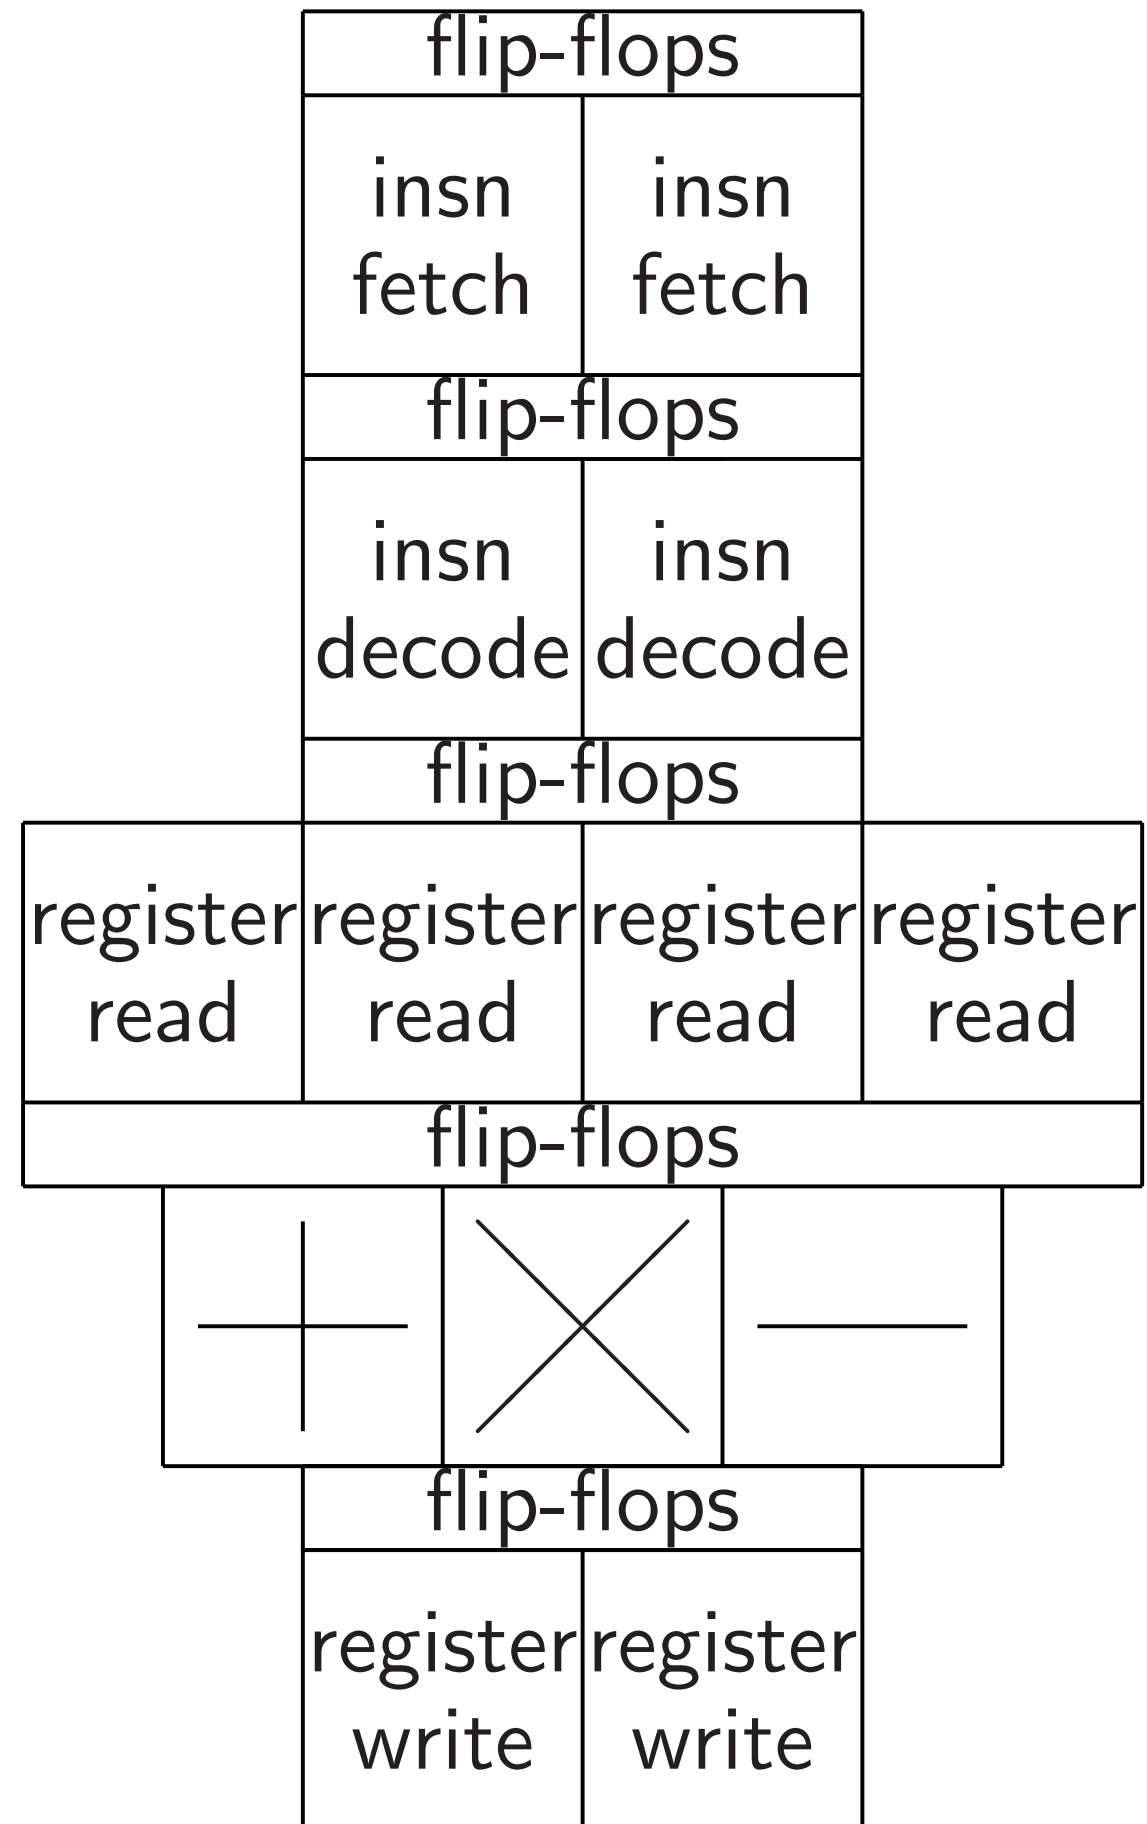
unction.

op area.

n semantics:

after-write.

“Superscalar” processing:



“Vector” processing

Expand each 32-bit
into n -vector of 32

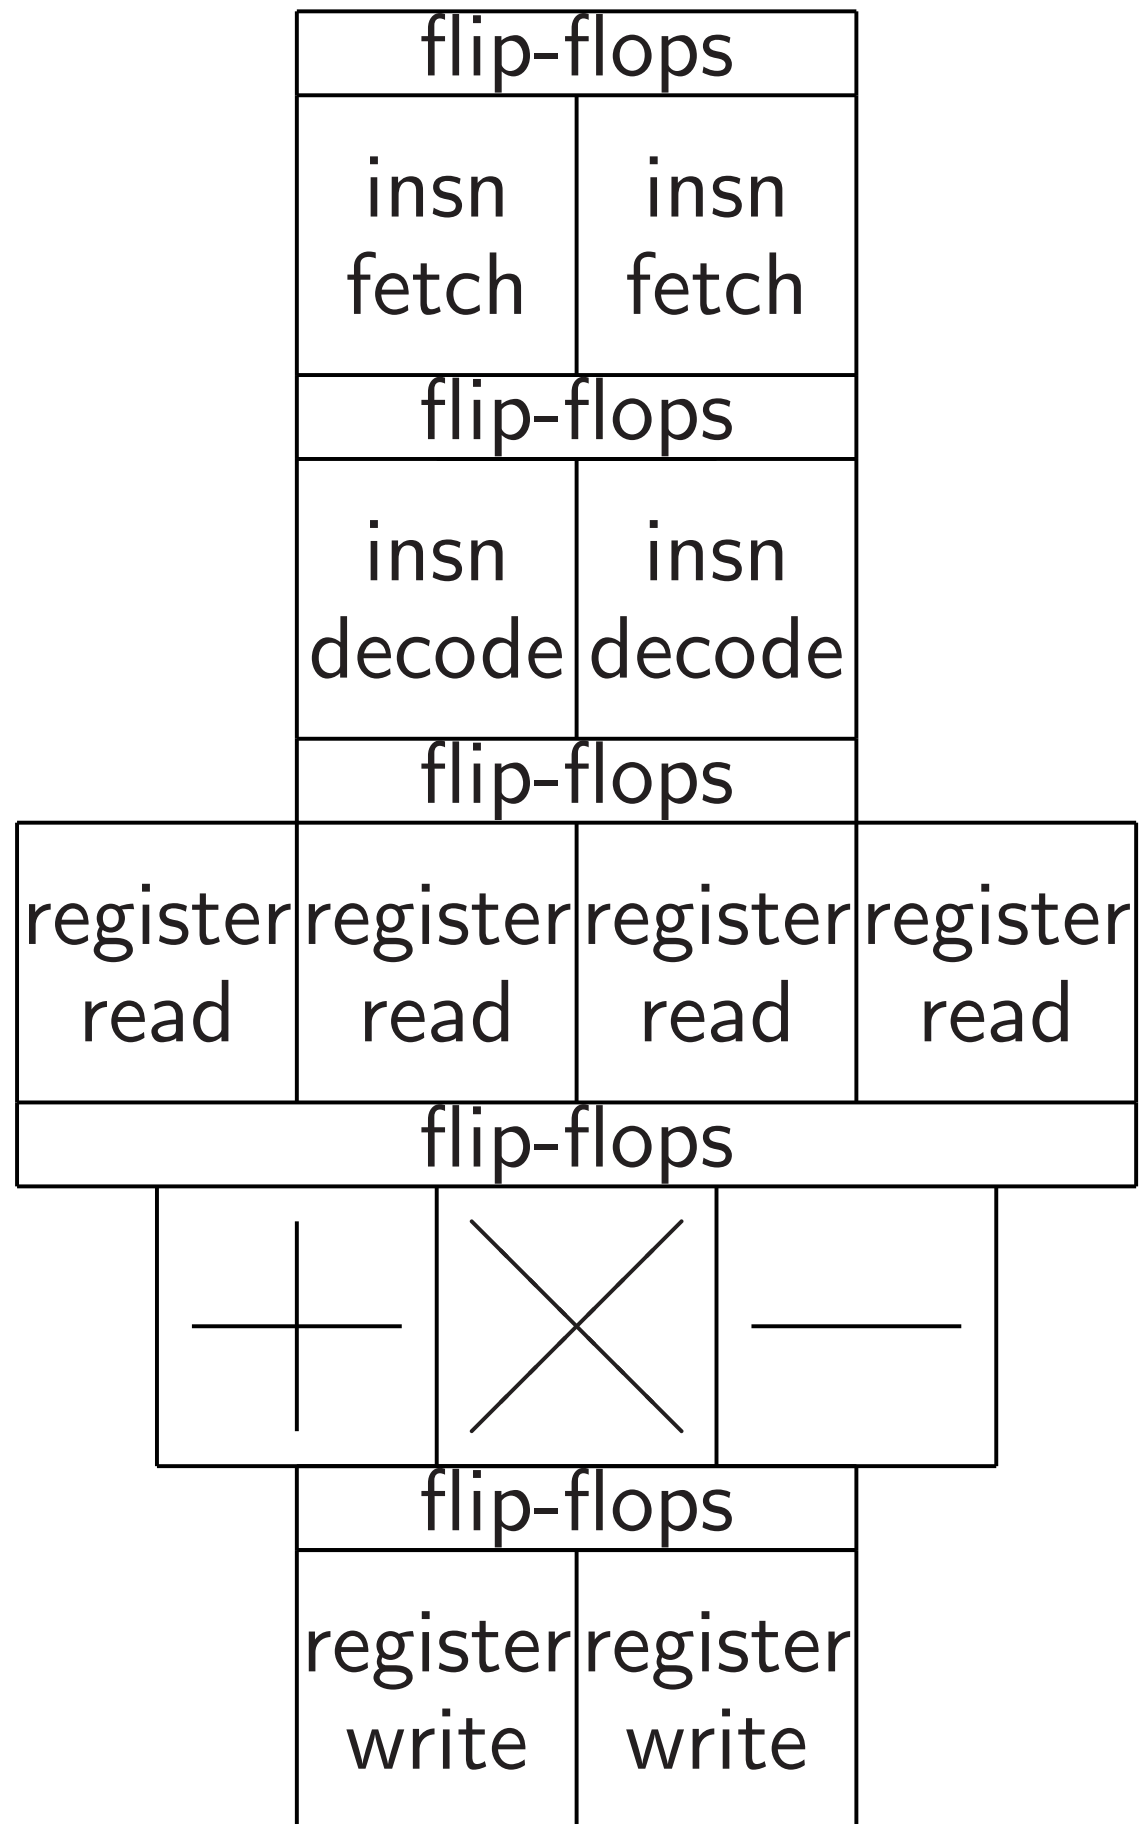
ARM “NEON” has

Intel “AVX2” has

Intel “AVX-512” h

GPUs have larger

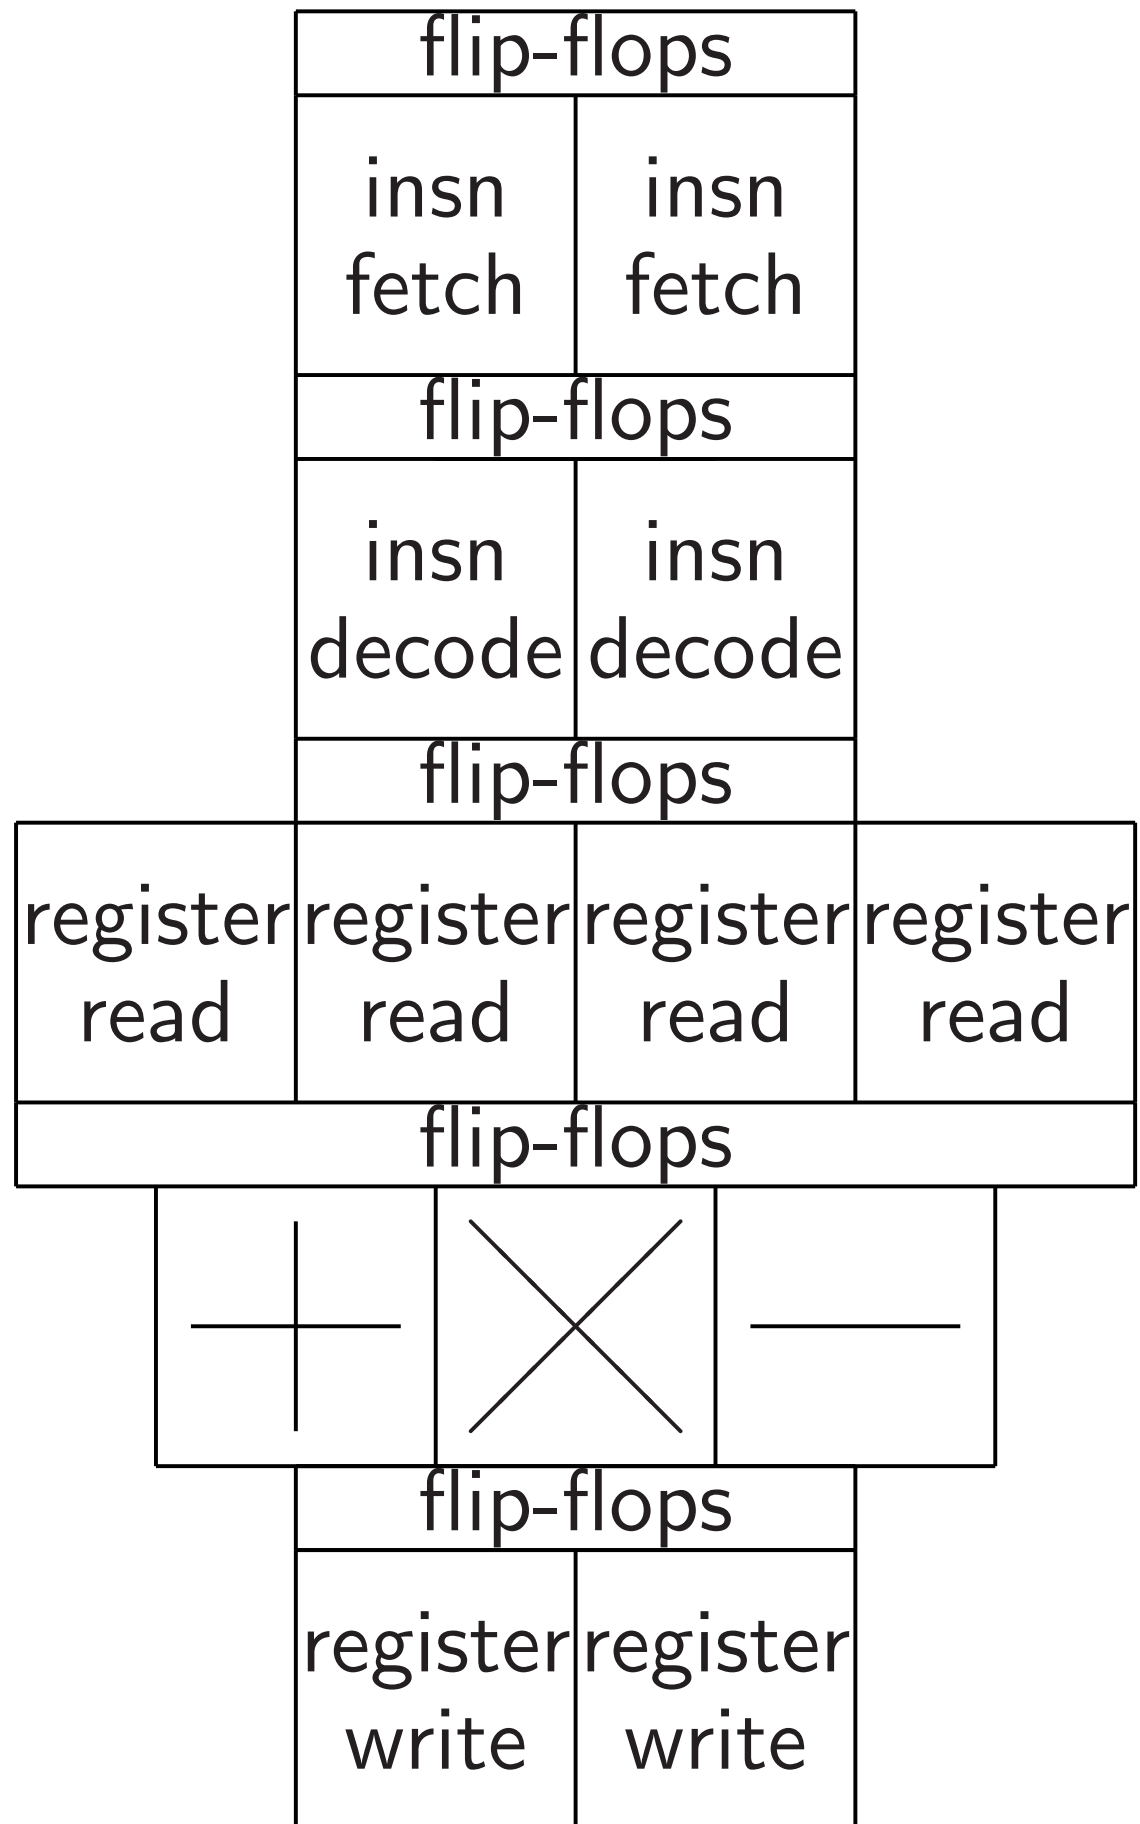
“Superscalar” processing:



“Vector” processing:

Expand each 32-bit integer into n -vector of 32-bit integers
 ARM “NEON” has $n = 4$;
 Intel “AVX2” has $n = 8$;
 Intel “AVX-512” has $n = 16$
 GPUs have larger n .

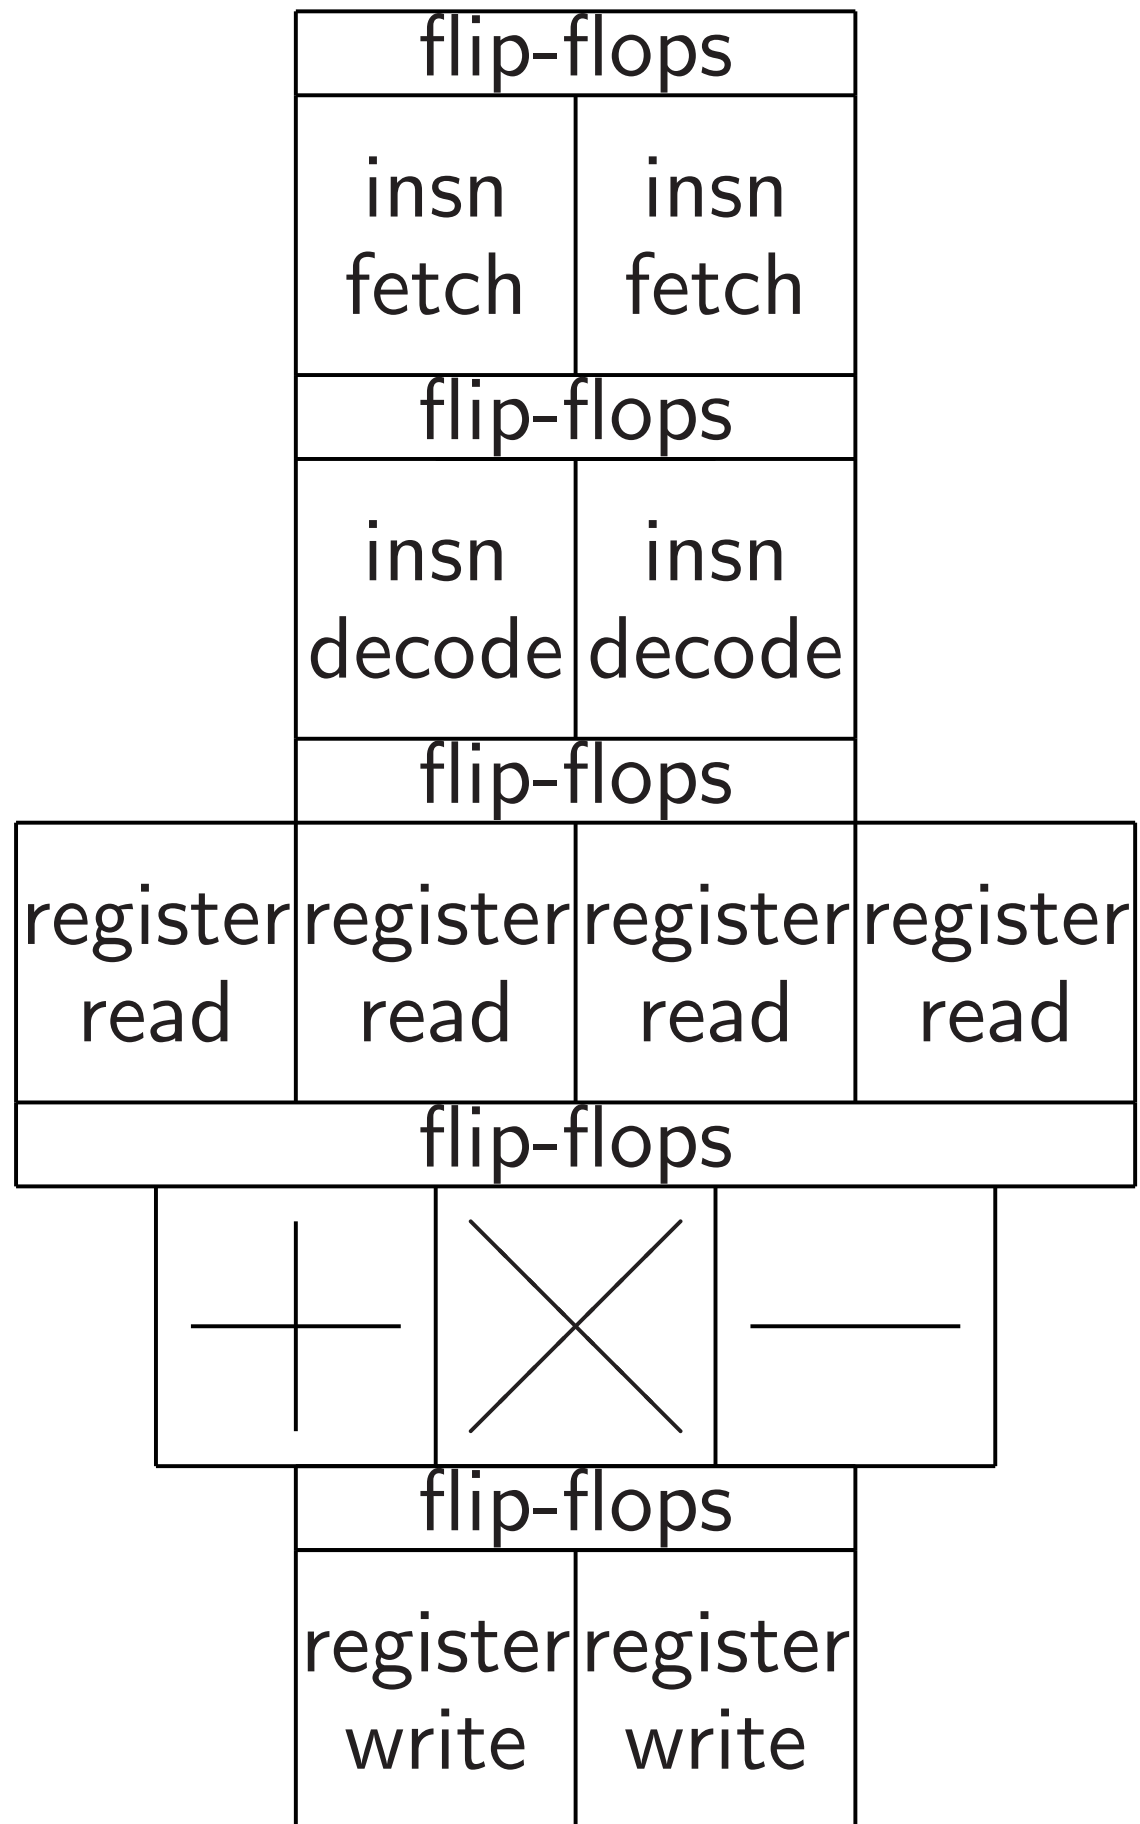
“Superscalar” processing:



“Vector” processing:

Expand each 32-bit integer into n -vector of 32-bit integers.
 ARM “NEON” has $n = 4$;
 Intel “AVX2” has $n = 8$;
 Intel “AVX-512” has $n = 16$;
 GPUs have larger n .

“Superscalar” processing:



“Vector” processing:

Expand each 32-bit integer into n -vector of 32-bit integers.
 ARM “NEON” has $n = 4$;
 Intel “AVX2” has $n = 8$;
 Intel “AVX-512” has $n = 16$;
 GPUs have larger n .

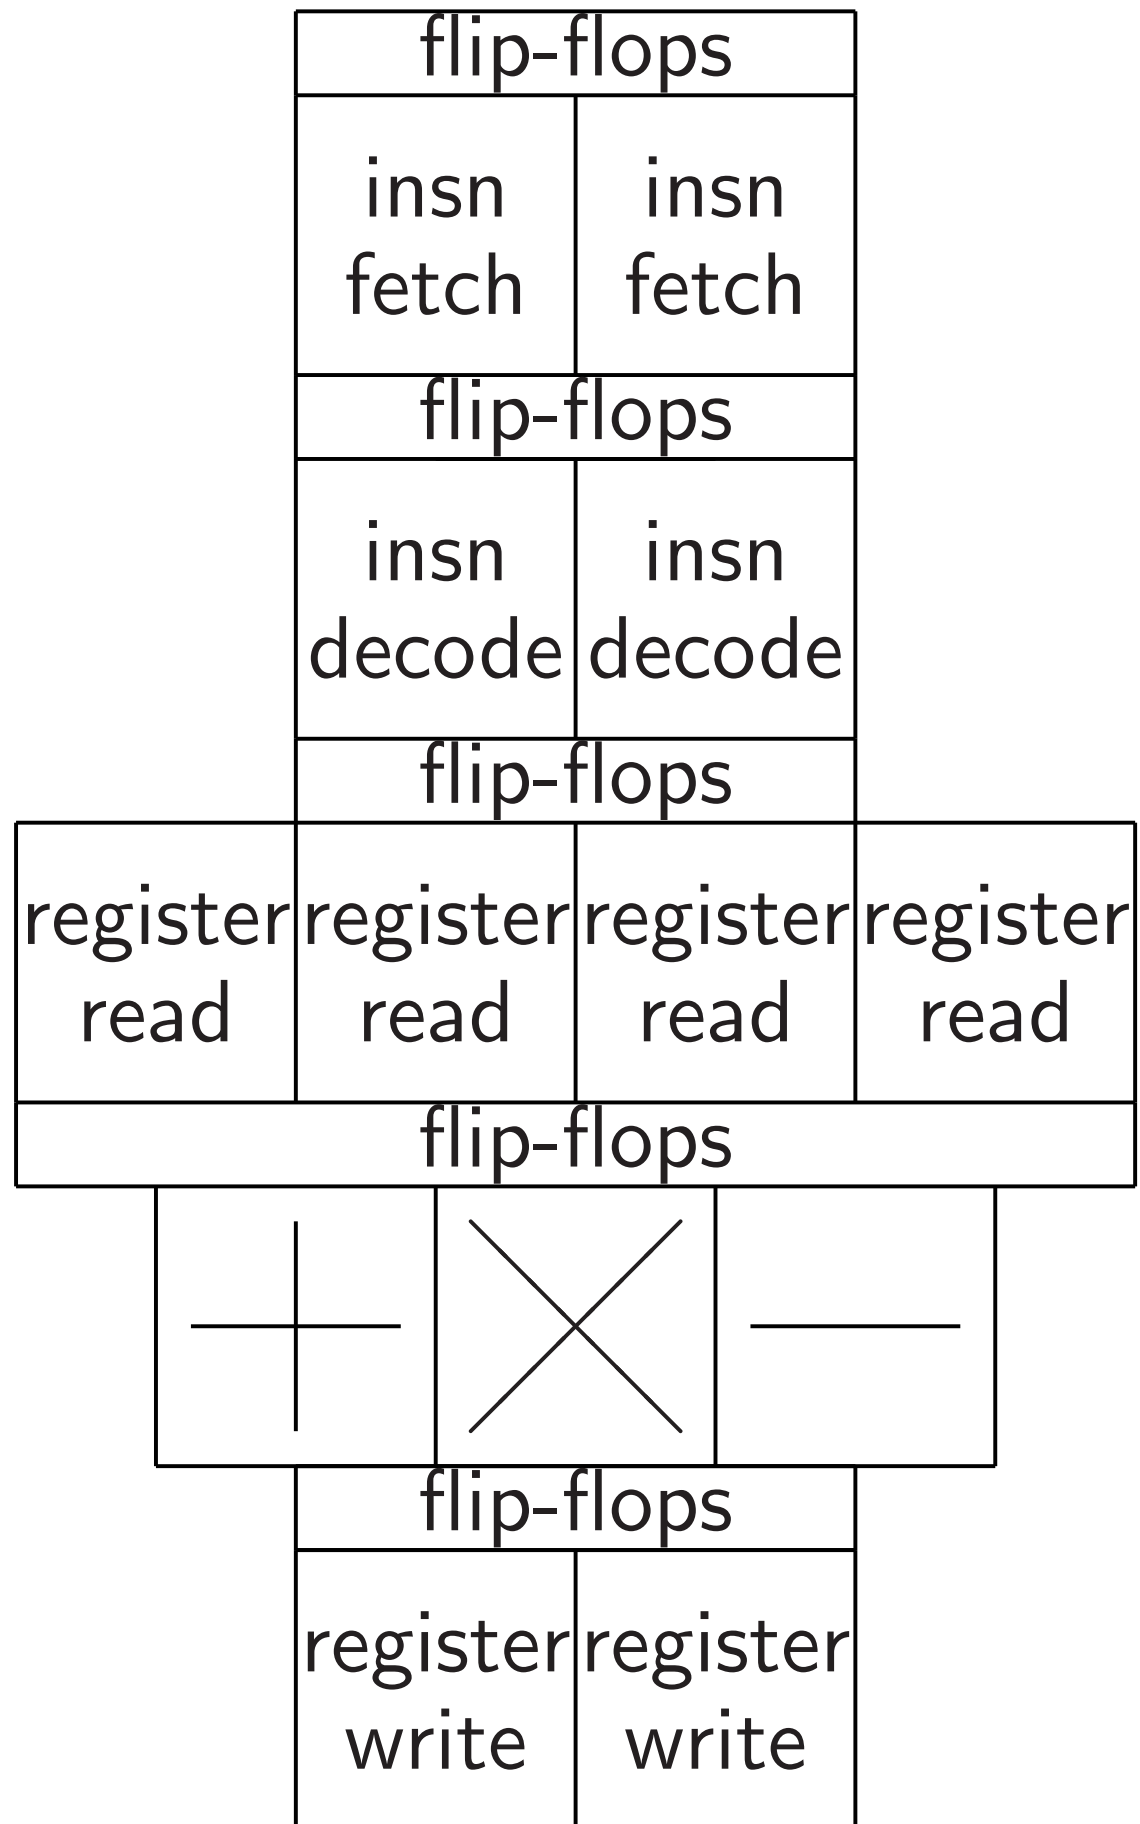
$n \times$ speedup if

$n \times$ arithmetic circuits,

$n \times$ read/write circuits.

Benefit: Amortizes insn circuits.

“Superscalar” processing:



“Vector” processing:

Expand each 32-bit integer into n -vector of 32-bit integers.
 ARM “NEON” has $n = 4$;
 Intel “AVX2” has $n = 8$;
 Intel “AVX-512” has $n = 16$;
 GPUs have larger n .

$n \times$ speedup if

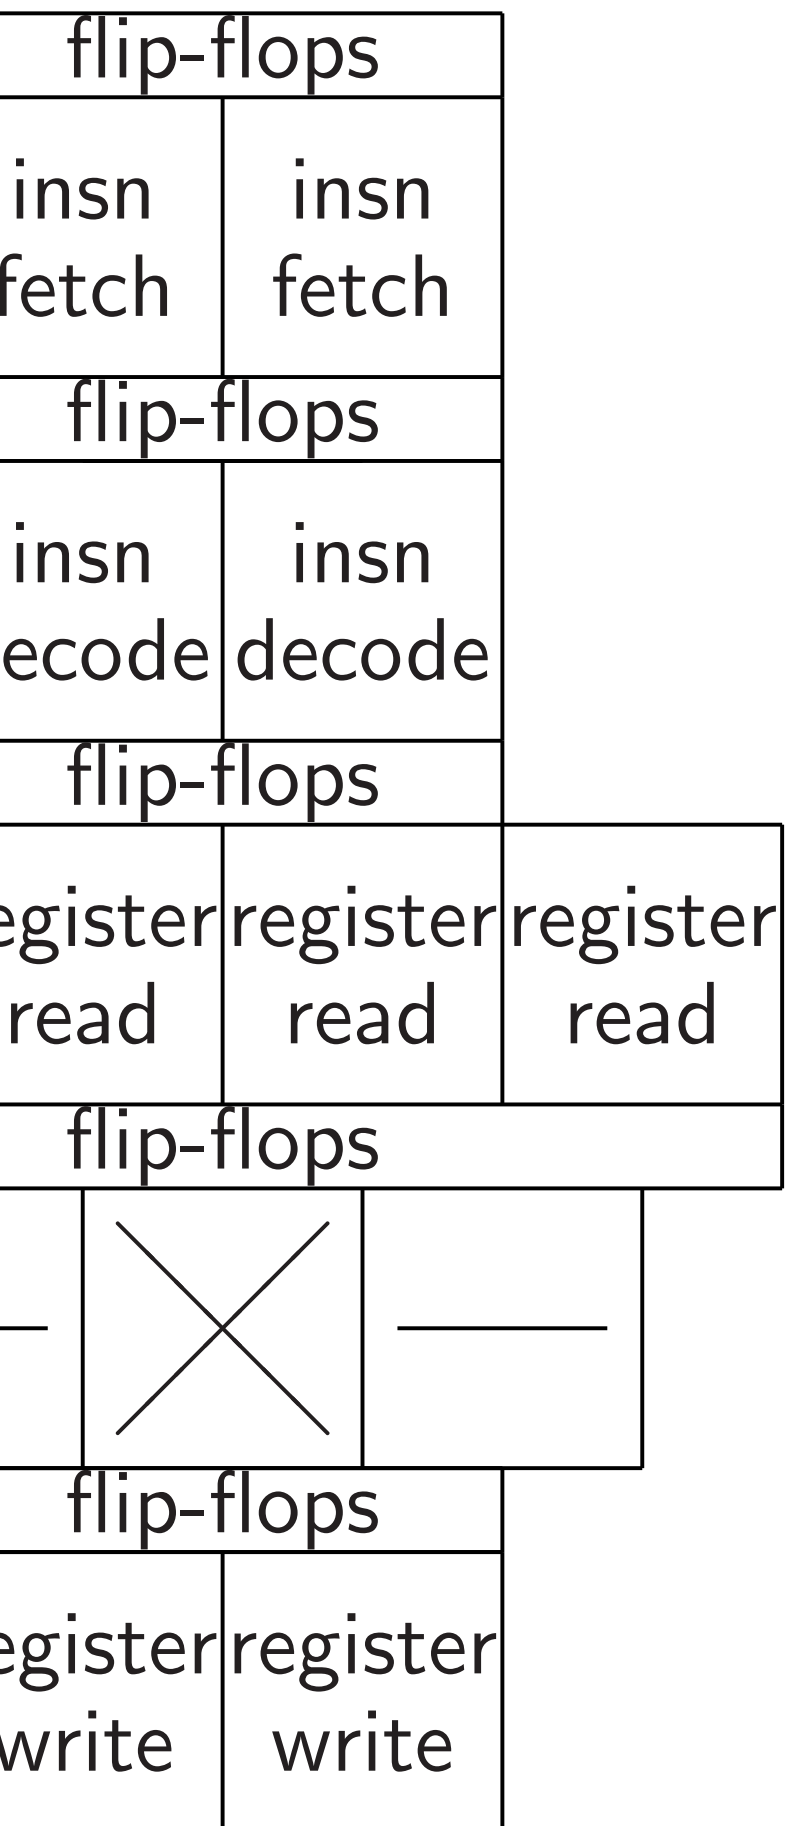
$n \times$ arithmetic circuits,

$n \times$ read/write circuits.

Benefit: Amortizes insn circuits.

Huge effect on higher-level algorithms and data structures.

“Scalar” processing:



“Vector” processing:

Expand each 32-bit integer into n -vector of 32-bit integers.

ARM “NEON” has $n = 4$;

Intel “AVX2” has $n = 8$;

Intel “AVX-512” has $n = 16$;

GPUs have larger n .

$n \times$ speedup if

$n \times$ arithmetic circuits,

$n \times$ read/write circuits.

Benefit: Amortizes insn circuits.

Huge effect on higher-level algorithms and data structures.

Network

How exp

Input: a

Each nu

represen

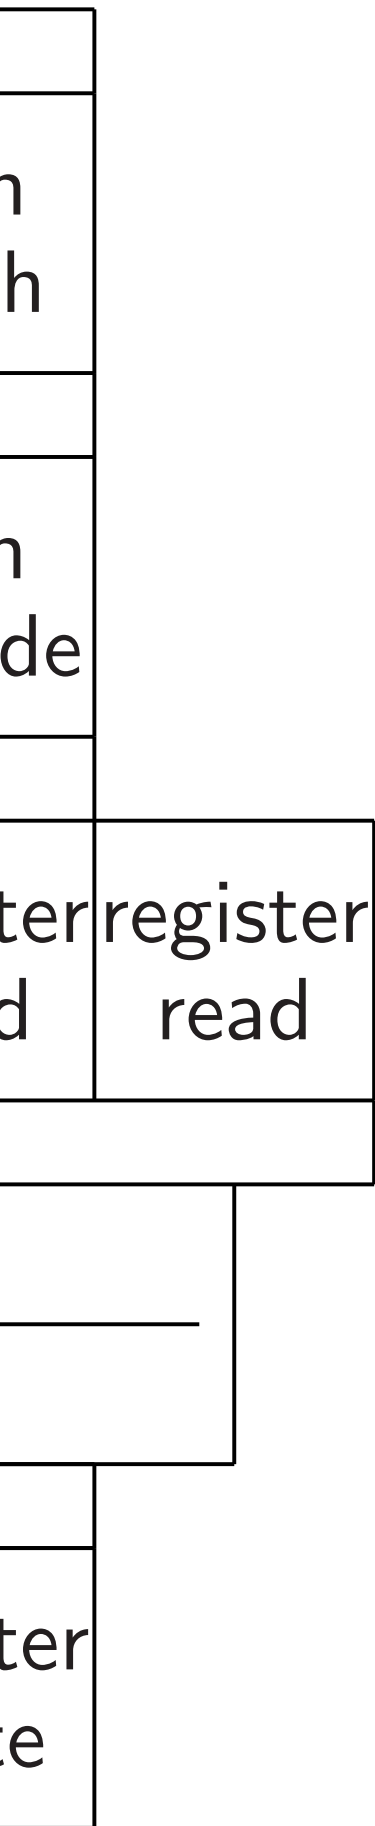
Output:

in increa

represen

same mu

processing:



“Vector” processing:

Expand each 32-bit integer into n -vector of 32-bit integers.

ARM “NEON” has $n = 4$;

Intel “AVX2” has $n = 8$;

Intel “AVX-512” has $n = 16$;

GPUs have larger n .

$n \times$ speedup if

$n \times$ arithmetic circuits,

$n \times$ read/write circuits.

Benefit: Amortizes insn circuits.

Huge effect on higher-level algorithms and data structures.

Network on chip:

How expensive is s

Input: array of n r

Each number in $\{$
represented in bina

Output: array of m

in increasing order

represented in bina

same multiset as i

“Vector” processing:

Expand each 32-bit integer
into n -vector of 32-bit integers.

ARM “NEON” has $n = 4$;

Intel “AVX2” has $n = 8$;

Intel “AVX-512” has $n = 16$;

GPUs have larger n .

$n \times$ speedup if

$n \times$ arithmetic circuits,

$n \times$ read/write circuits.

Benefit: Amortizes insn circuits.

Huge effect on higher-level
algorithms and data structures.

Network on chip: the mesh

How expensive is sorting?

Input: array of n numbers.

Each number in $\{1, 2, \dots, n\}$
represented in binary.

Output: array of n numbers
in increasing order,

represented in binary;
same multiset as input.

“Vector” processing:

Expand each 32-bit integer
into n -vector of 32-bit integers.

ARM “NEON” has $n = 4$;

Intel “AVX2” has $n = 8$;

Intel “AVX-512” has $n = 16$;

GPUs have larger n .

$n \times$ speedup if

$n \times$ arithmetic circuits,

$n \times$ read/write circuits.

Benefit: Amortizes insn circuits.

Huge effect on higher-level
algorithms and data structures.

Network on chip: the mesh

How expensive is sorting?

Input: array of n numbers.

Each number in $\{1, 2, \dots, n^2\}$,
represented in binary.

Output: array of n numbers,
in increasing order,
represented in binary;
same multiset as input.

“Vector” processing:

Expand each 32-bit integer into n -vector of 32-bit integers.

ARM “NEON” has $n = 4$;

Intel “AVX2” has $n = 8$;

Intel “AVX-512” has $n = 16$;

GPUs have larger n .

$n\times$ speedup if

$n\times$ arithmetic circuits,

$n\times$ read/write circuits.

Benefit: Amortizes insn circuits.

Huge effect on higher-level algorithms and data structures.

Network on chip: the mesh

How expensive is sorting?

Input: array of n numbers.

Each number in $\{1, 2, \dots, n^2\}$, represented in binary.

Output: array of n numbers, in increasing order, represented in binary; same multiset as input.

Metric: seconds used by circuit of area $n^{1+o(1)}$.

For simplicity assume $n = 4^k$.

' processing:

each 32-bit integer
ector of 32-bit integers.

"NEON" has $n = 4$;

"VX2" has $n = 8$;

"VX-512" has $n = 16$;

ave larger n .

dup if

arithmetic circuits,

/write circuits.

Amortizes insn circuits.

fect on higher-level

ms and data structures.

Network on chip: the mesh

How expensive is sorting?

Input: array of n numbers.

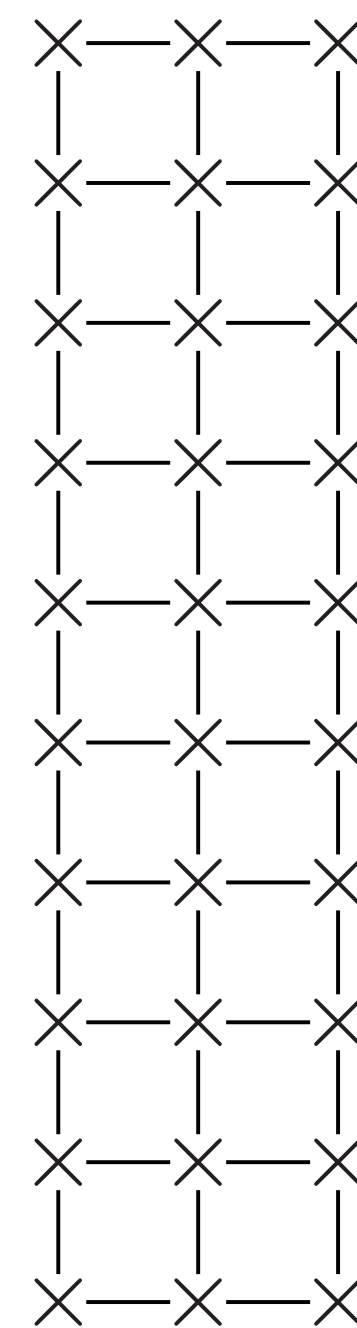
Each number in $\{1, 2, \dots, n^2\}$,
represented in binary.

Output: array of n numbers,
in increasing order,
represented in binary;
same multiset as input.

Metric: seconds used by
circuit of area $n^{1+o(1)}$.

For simplicity assume $n = 4^k$.

Spread a
square m
each of
with nea



Network on chip: the mesh

How expensive is sorting?

Input: array of n numbers.

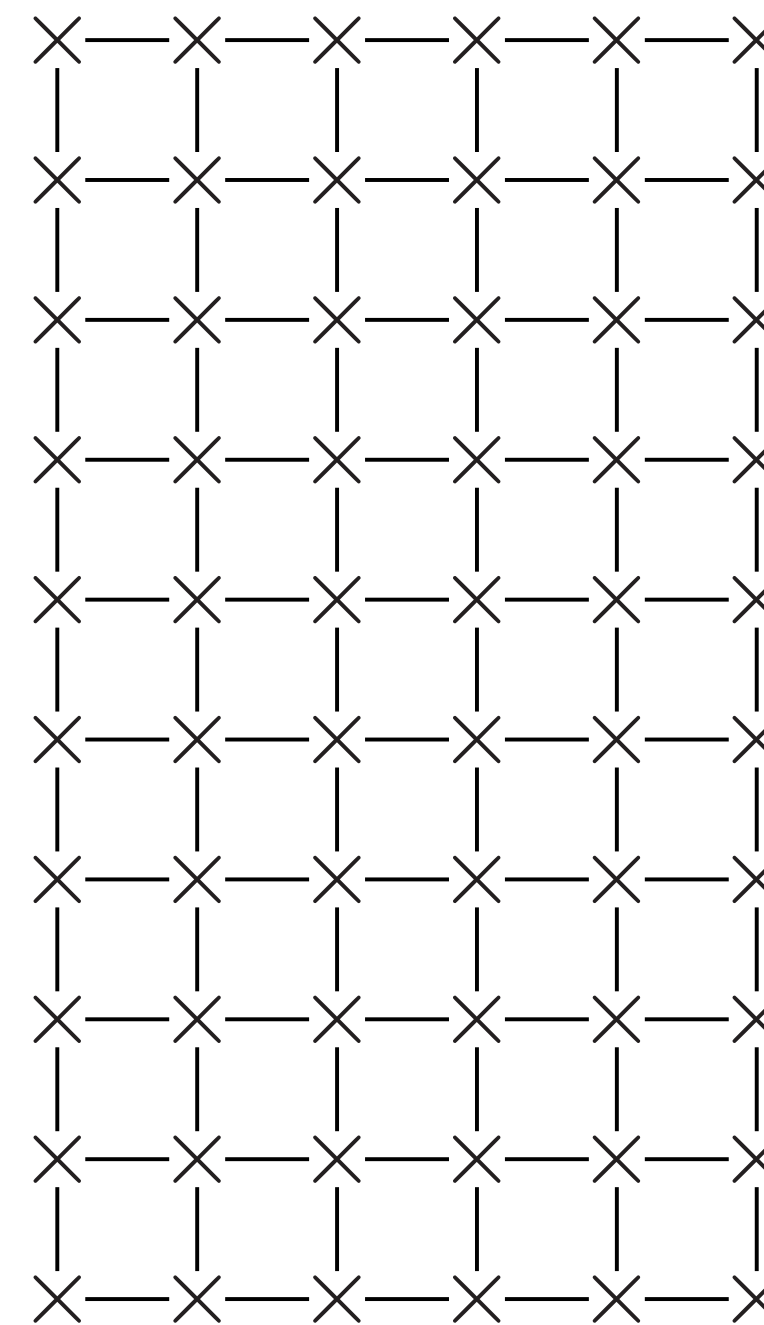
Each number in $\{1, 2, \dots, n^2\}$,
represented in binary.

Output: array of n numbers,
in increasing order,
represented in binary;
same multiset as input.

Metric: seconds used by
circuit of area $n^{1+o(1)}$.

For simplicity assume $n = 4^k$.

Spread array across
square mesh of n nodes,
each of area $n^{o(1)}$
with near-neighborhood



Network on chip: the mesh

How expensive is sorting?

Input: array of n numbers.

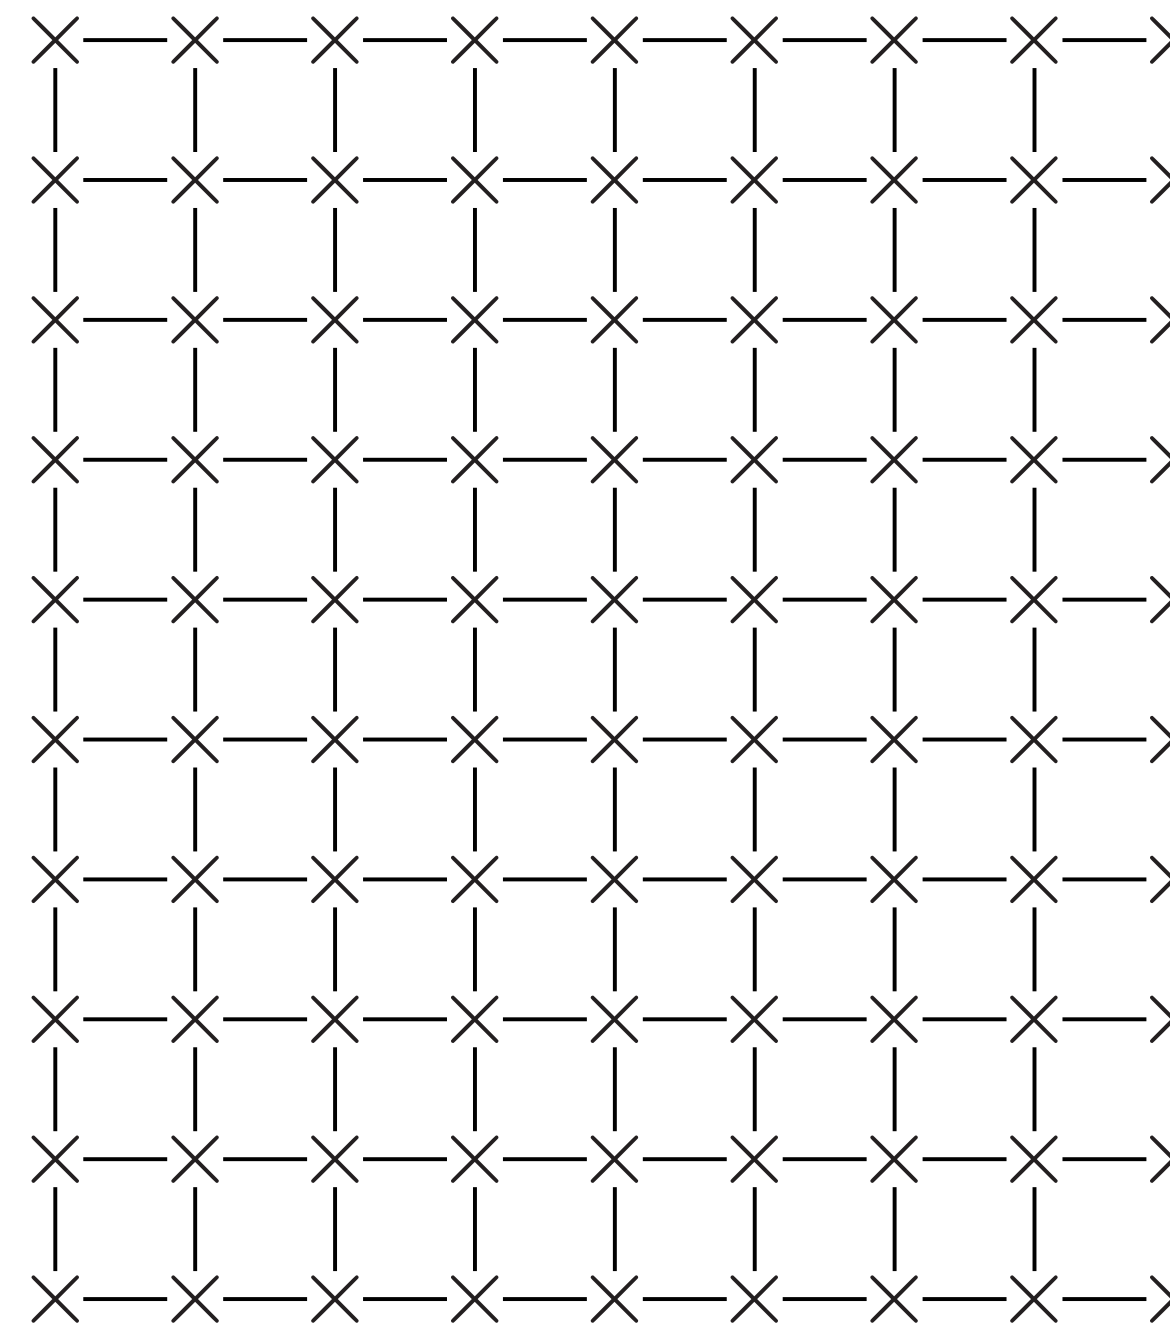
Each number in $\{1, 2, \dots, n^2\}$,
represented in binary.

Output: array of n numbers,
in increasing order,
represented in binary;
same multiset as input.

Metric: seconds used by
circuit of area $n^{1+o(1)}$.

For simplicity assume $n = 4^k$.

Spread array across
square mesh of n small cells
each of area $n^{o(1)}$,
with near-neighbor wiring:



Network on chip: the mesh

How expensive is sorting?

Input: array of n numbers.

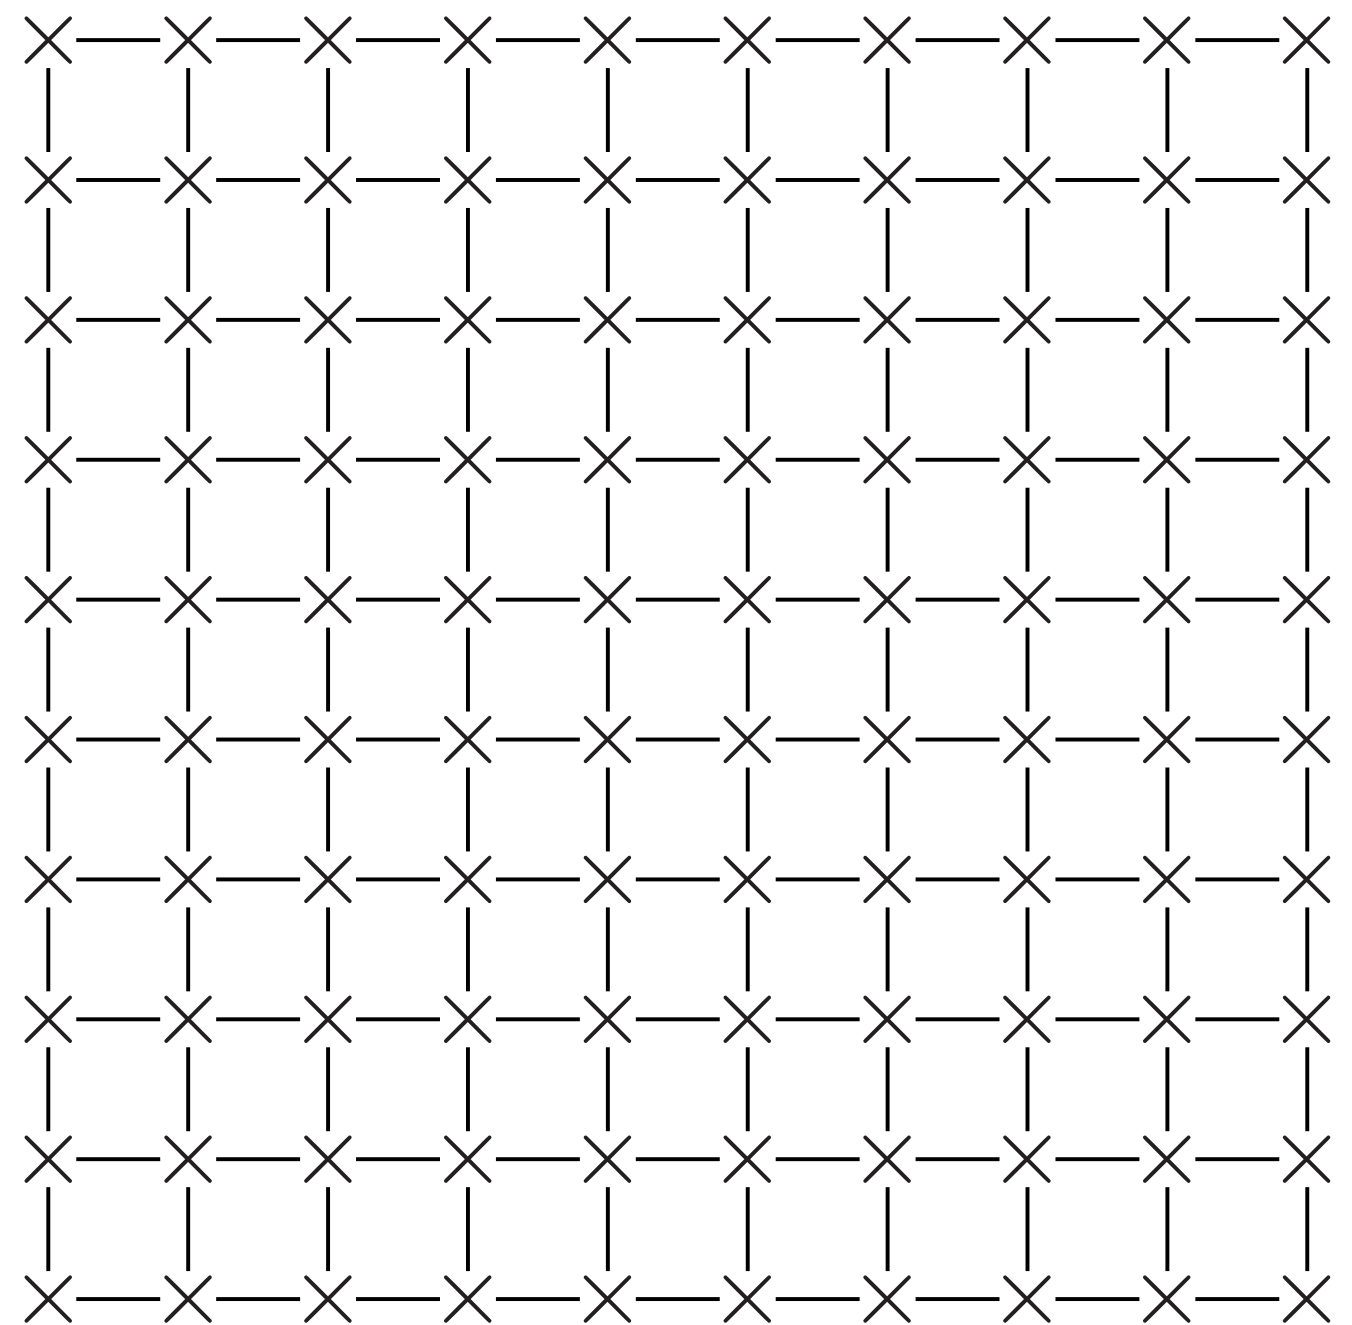
Each number in $\{1, 2, \dots, n^2\}$,
represented in binary.

Output: array of n numbers,
in increasing order,
represented in binary;
same multiset as input.

Metric: seconds used by
circuit of area $n^{1+o(1)}$.

For simplicity assume $n = 4^k$.

Spread array across
square mesh of n small cells,
each of area $n^{o(1)}$,
with near-neighbor wiring:



on chip: the mesh

ensive is sorting?

array of n numbers.

number in $\{1, 2, \dots, n^2\}$,

ted in binary.

array of n numbers,

using order,

ted in binary;

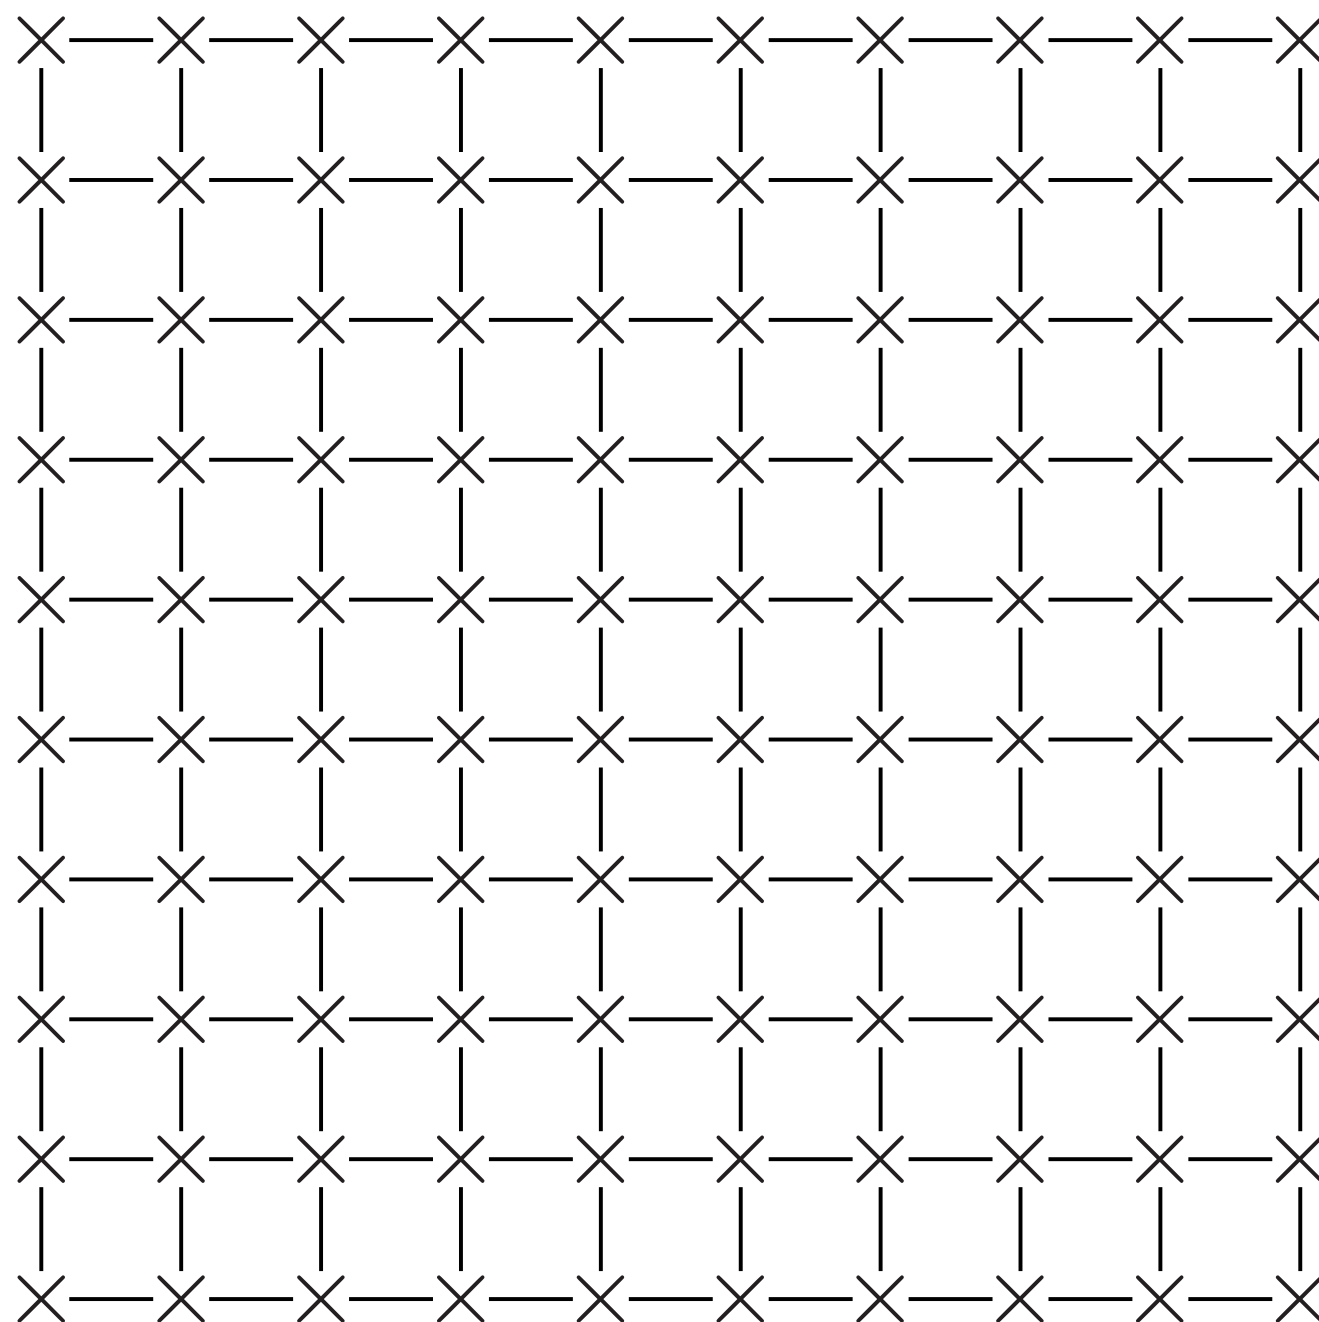
multiset as input.

seconds used by

f area $n^{1+o(1)}$.

licity assume $n = 4^k$.

Spread array across
square mesh of n small cells,
each of area $n^{o(1)}$,
with near-neighbor wiring:



Sort row

in $n^{0.5+o(1)}$

- Sort e

3 1 4

1 3 1 4

- Sort a

1 3 1 4

1 1 3 4

- Repea

equals

the mesh

sorting?

numbers.

$\{1, 2, \dots, n^2\}$,

ary.

n numbers,

,

ary;

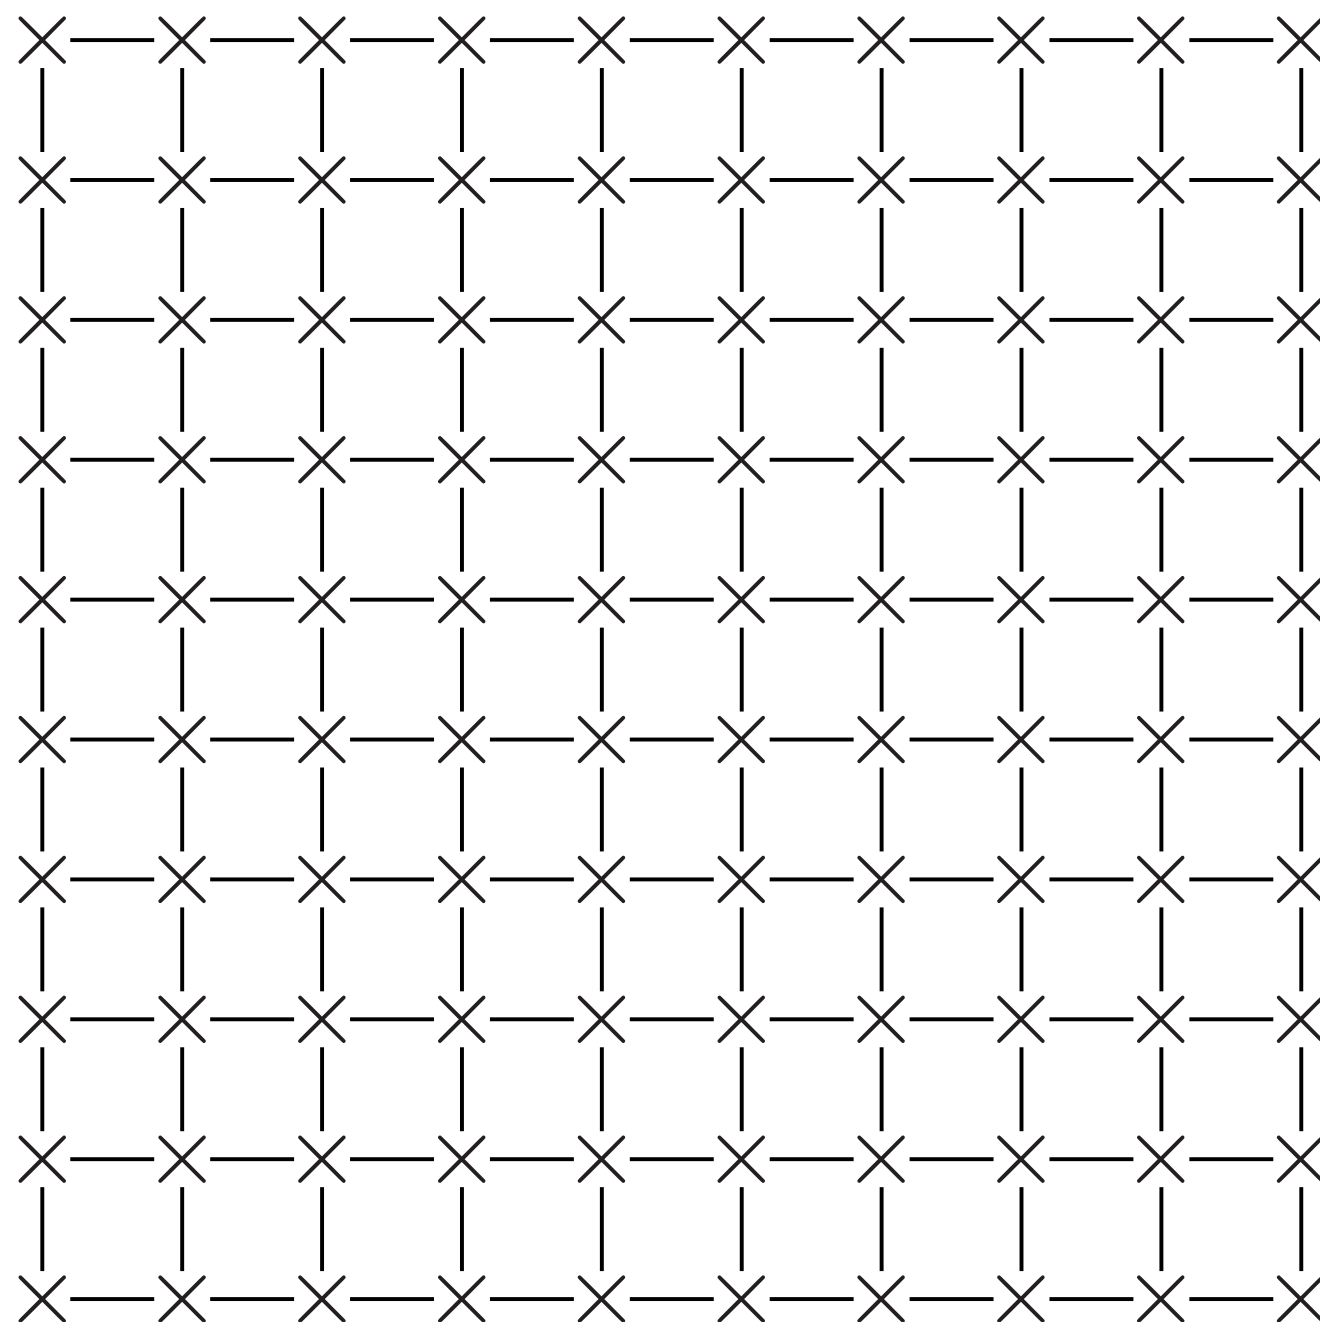
input.

sed by

$o(1)$.

me $n = 4^k$.

Spread array across
square mesh of n small cells,
each of area $n^{o(1)}$,
with near-neighbor wiring:



Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds

- Sort each pair in

3 1 4 1 5 9 2 6

1 3 1 4 5 9 2 6

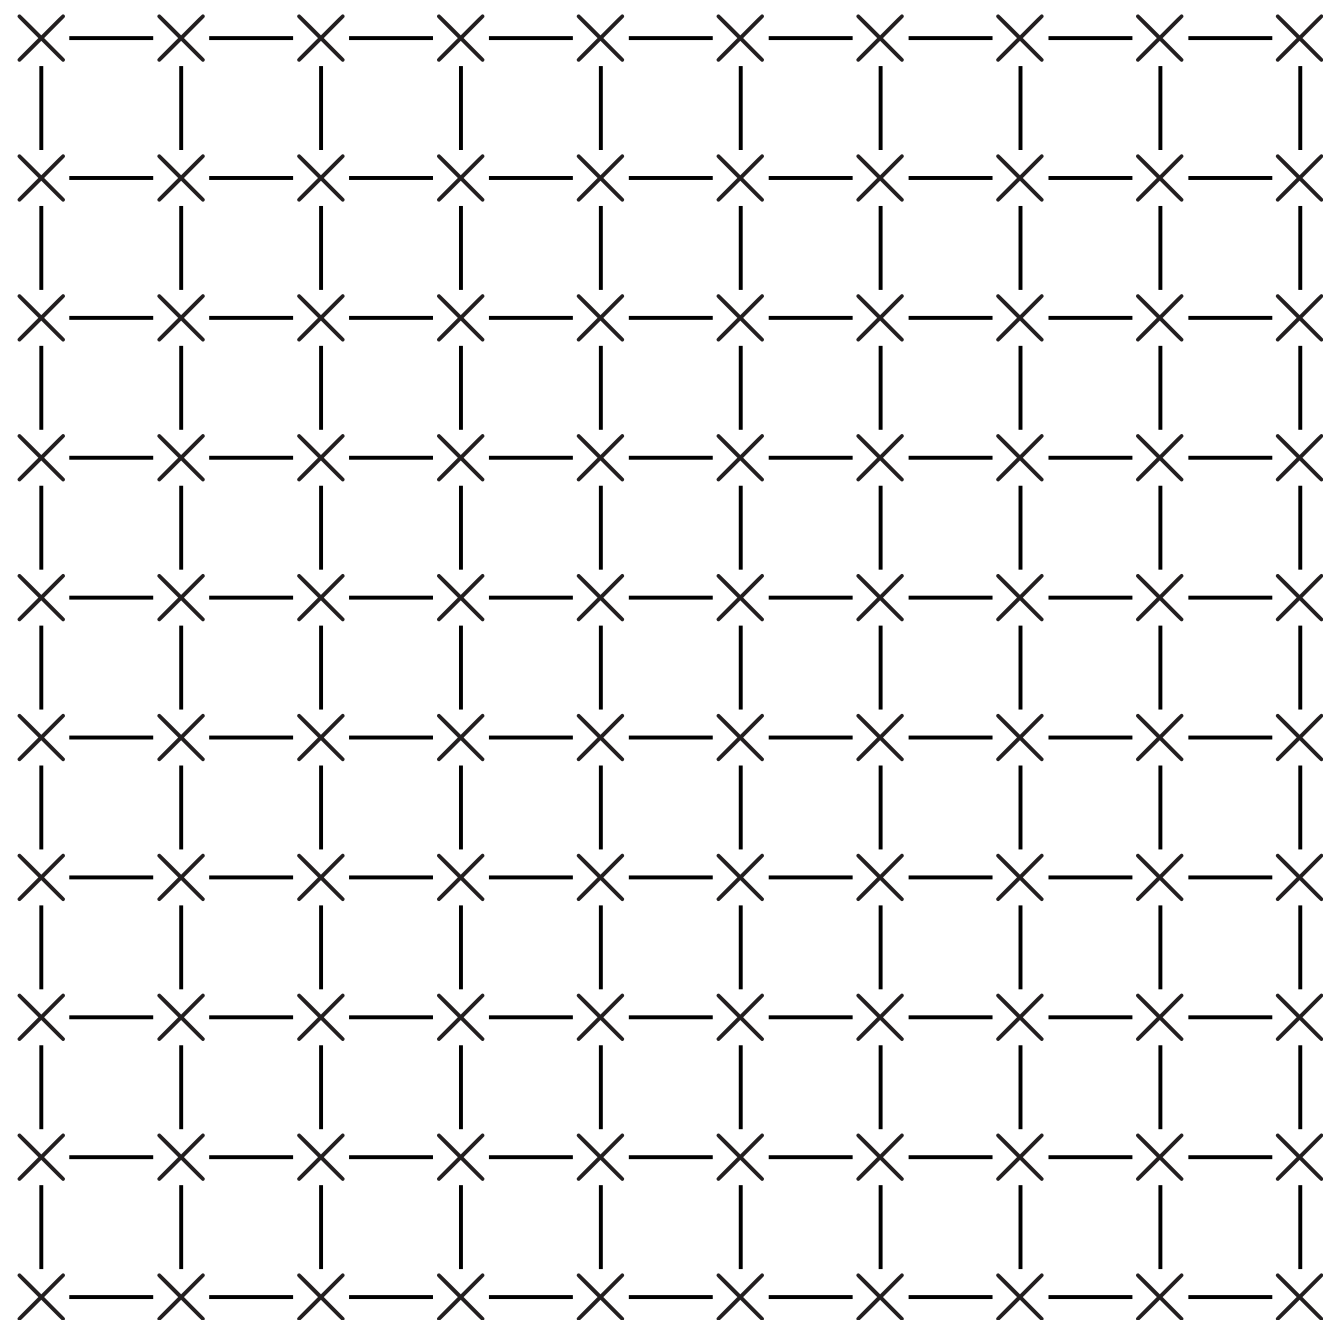
- Sort alternate pairs

1 3 1 4 5 9 2 6

1 1 3 4 5 2 9 6

- Repeat until number of nodes equals row length

Spread array across
square mesh of n small cells,
each of area $n^{o(1)}$,
with near-neighbor wiring:



Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds:

- Sort each pair in parallel.

3 1 4 1 5 9 2 6 \mapsto

1 3 1 4 5 9 2 6

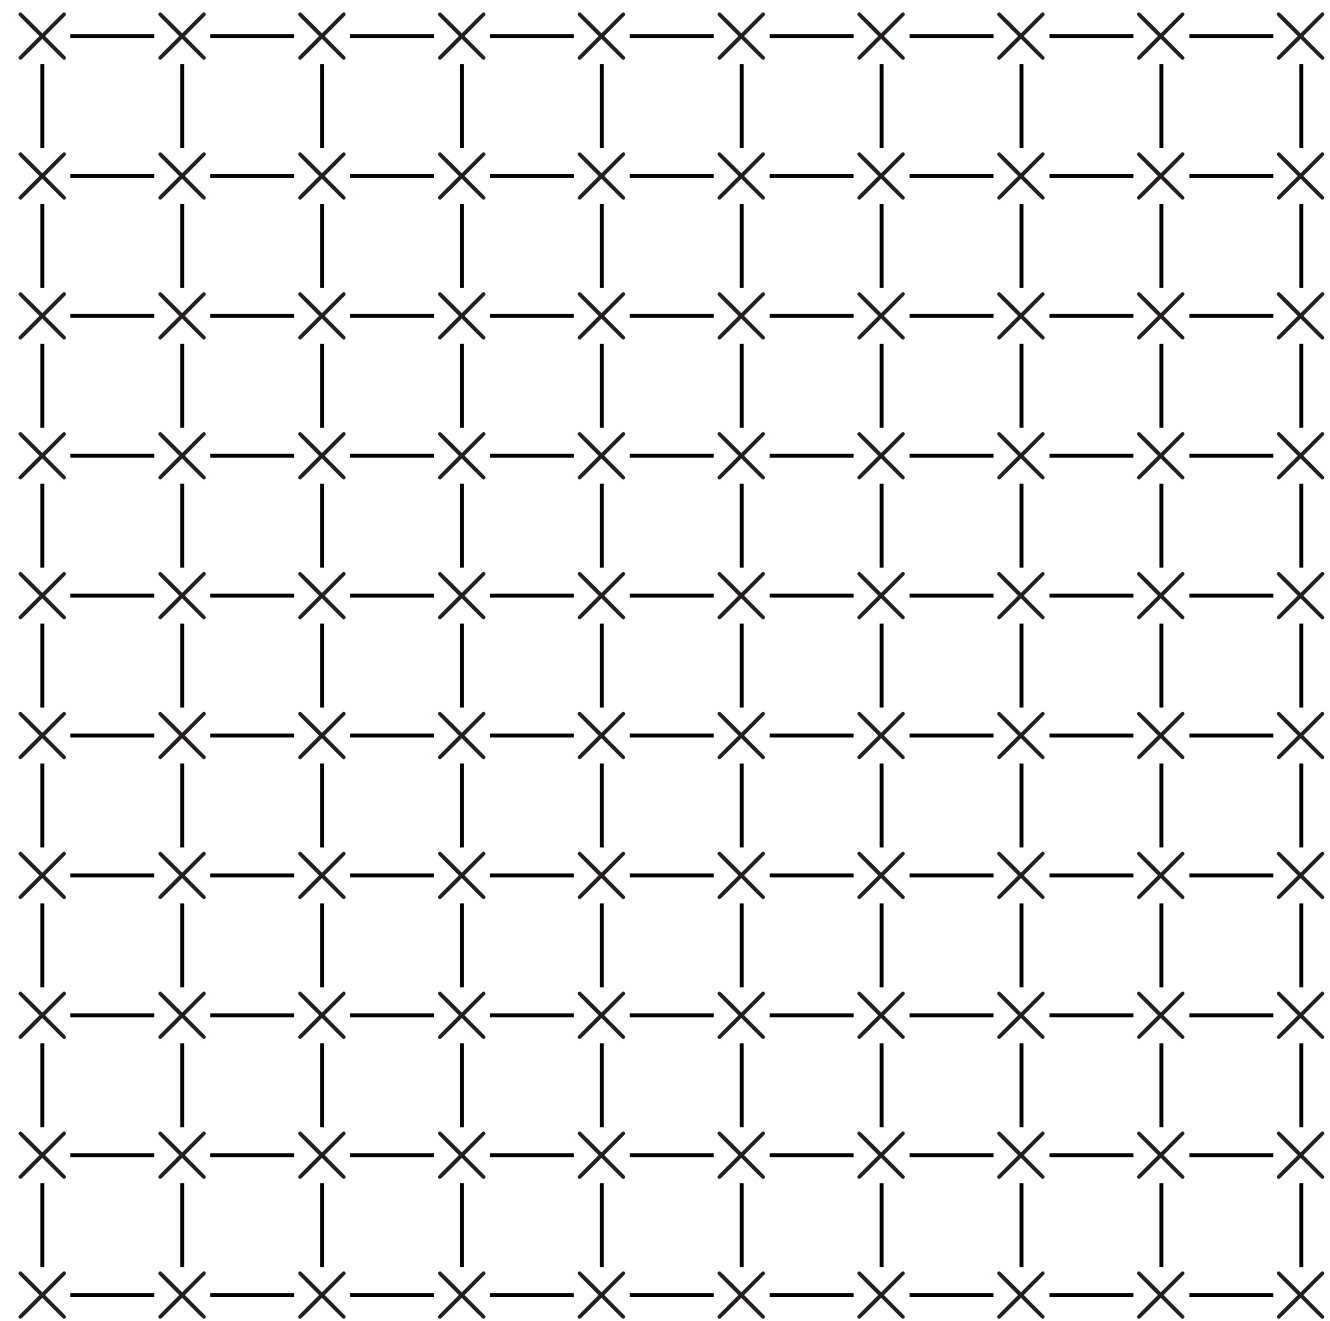
- Sort alternate pairs in parallel.

1 3 1 4 5 9 2 6 \mapsto

1 1 3 4 5 2 9 6

- Repeat until number of stages equals row length.

Spread array across
square mesh of n small cells,
each of area $n^{o(1)}$,
with near-neighbor wiring:



Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds:

- Sort each pair in parallel.

3 1 4 1 5 9 2 6 \mapsto

1 3 1 4 5 9 2 6

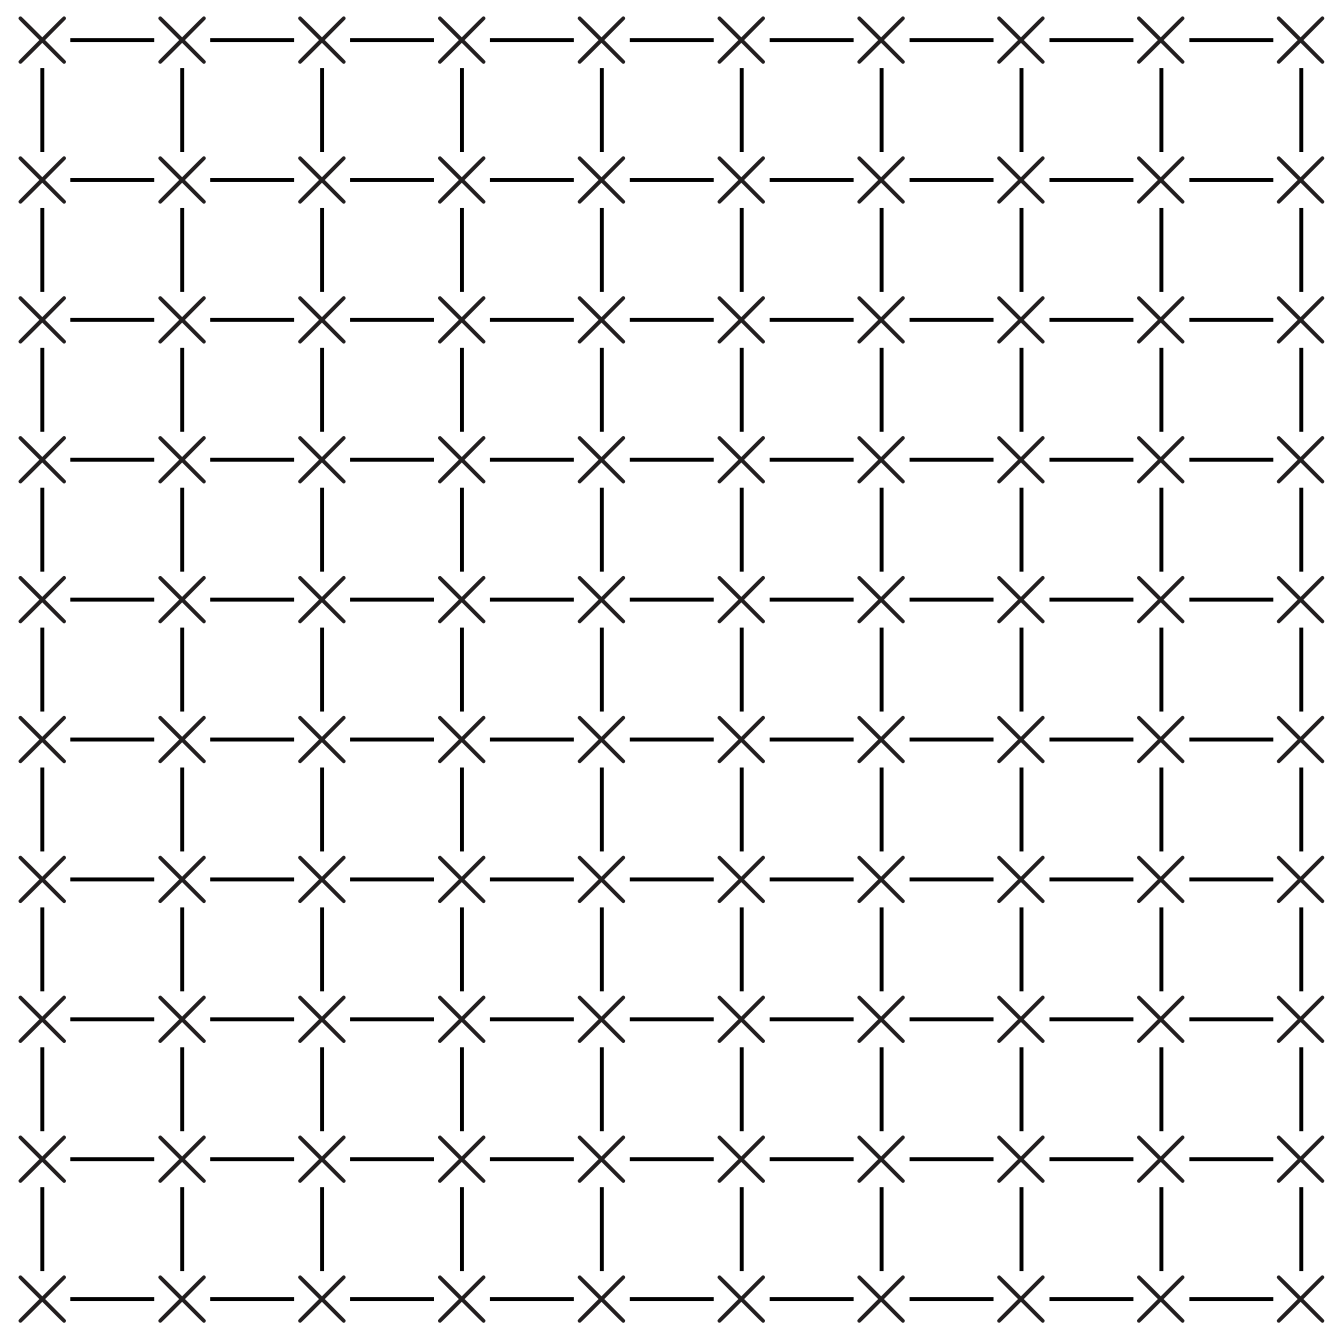
- Sort alternate pairs in parallel.

1 3 1 4 5 9 2 6 \mapsto

1 1 3 4 5 2 9 6

- Repeat until number of steps equals row length.

Spread array across
square mesh of n small cells,
each of area $n^{o(1)}$,
with near-neighbor wiring:



Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds:

- Sort each pair in parallel.

3 1 4 1 5 9 2 6 \mapsto

1 3 1 4 5 9 2 6

- Sort alternate pairs in parallel.

1 3 1 4 5 9 2 6 \mapsto

1 1 3 4 5 2 9 6

- Repeat until number of steps equals row length.

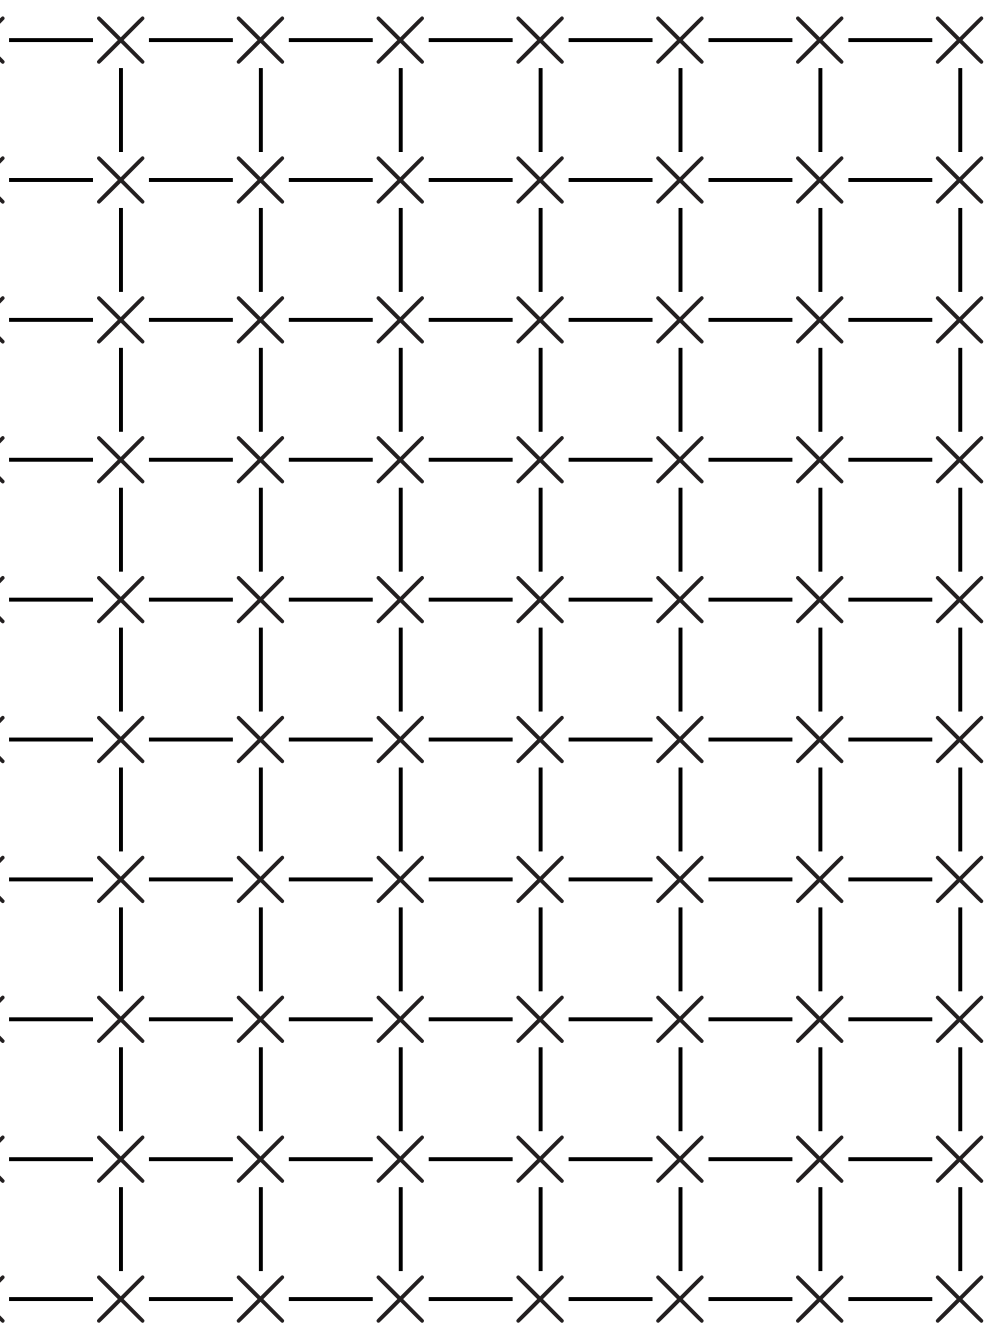
Sort *each* row, in parallel,
in a *total* of $n^{0.5+o(1)}$ seconds.

array across

mesh of n small cells,

area $n^{o(1)}$,

near-neighbor wiring:



Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds:

- Sort each pair in parallel.

3 1 4 1 5 9 2 6 \mapsto

1 3 1 4 5 9 2 6

- Sort alternate pairs in parallel.

1 3 1 4 5 9 2 6 \mapsto

1 1 3 4 5 2 9 6

- Repeat until number of steps equals row length.

Sort *each* row, in parallel,

in a *total* of $n^{0.5+o(1)}$ seconds.

Sort all

in $n^{0.5+o(1)}$

- Recurs

in para

- Sort e

- Sort e

- Sort e

- Sort e

With pro

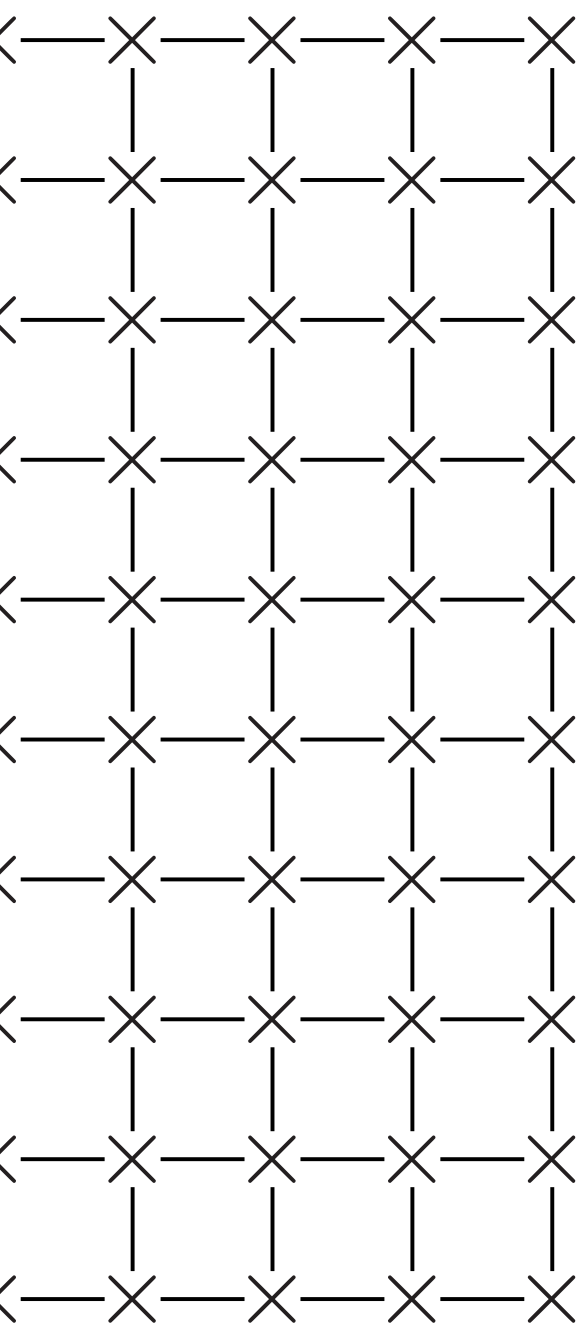
left-to-ri

for each

that this

ss
small cells,

wiring:



Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds:

- Sort each pair in parallel.

3 1 4 1 5 9 2 6 \mapsto

1 3 1 4 5 9 2 6

- Sort alternate pairs in parallel.

1 3 1 4 5 9 2 6 \mapsto

1 1 3 4 5 2 9 6

- Repeat until number of steps equals row length.

Sort *each* row, in parallel,
in a *total* of $n^{0.5+o(1)}$ seconds.

Sort all n cells
in $n^{0.5+o(1)}$ seconds

- Recursively sort $n^{0.5}$ cells in parallel, if $n > n^{0.5}$
- Sort each column in parallel
- Sort each row in parallel
- Sort each column in parallel
- Sort each row in parallel

With proper choice of n ,
left-to-right/right-to-left
for each row, can be done
that this sorts whole array

Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds:

- Sort each pair in parallel.

3 1 4 1 5 9 2 6 \mapsto

1 3 1 4 5 9 2 6

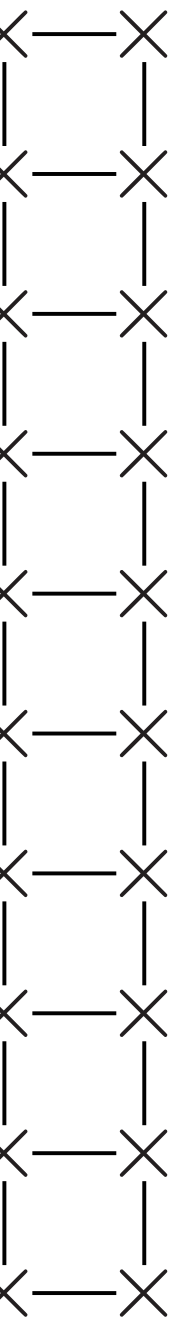
- Sort alternate pairs in parallel.

1 3 1 4 5 9 2 6 \mapsto

1 1 3 4 5 2 9 6

- Repeat until number of steps equals row length.

Sort *each* row, in parallel,
in a *total* of $n^{0.5+o(1)}$ seconds.



Sort all n cells
in $n^{0.5+o(1)}$ seconds:

- Recursively sort quadrants in parallel, if $n > 1$.
- Sort each column in parallel.
- Sort each row in parallel.
- Sort each column in parallel.
- Sort each row in parallel.

With proper choice of left-to-right/right-to-left for each row, can prove that this sorts whole array.

Sort row of $n^{0.5}$ cells
in $n^{0.5+o(1)}$ seconds:

- Sort each pair in parallel.

3 1 4 1 5 9 2 6 \mapsto

1 3 1 4 5 9 2 6

- Sort alternate pairs in parallel.

1 3 1 4 5 9 2 6 \mapsto

1 1 3 4 5 2 9 6

- Repeat until number of steps equals row length.

Sort *each* row, in parallel,
in a *total* of $n^{0.5+o(1)}$ seconds.

Sort all n cells
in $n^{0.5+o(1)}$ seconds:

- Recursively sort quadrants in parallel, if $n > 1$.
- Sort each column in parallel.
- Sort each row in parallel.
- Sort each column in parallel.
- Sort each row in parallel.

With proper choice of left-to-right/right-to-left for each row, can prove that this sorts whole array.

of $n^{0.5}$ cells
 in $\Theta(n^{0.5})$ seconds:

Sort each pair in parallel.

1 5 9 2 6 \mapsto

4 5 9 2 6

Sort alternate pairs in parallel.

4 5 9 2 6 \mapsto

4 5 2 9 6

Continue until number of steps
 is proportional to
 row length.

Sort each row, in parallel,

in total of $n^{0.5+o(1)}$ seconds.

Sort all n cells
 in $n^{0.5+o(1)}$ seconds:

- Recursively sort quadrants in parallel, if $n > 1$.
- Sort each column in parallel.
- Sort each row in parallel.
- Sort each column in parallel.
- Sort each row in parallel.

With proper choice of
 left-to-right/right-to-left
 for each row, can prove
 that this sorts whole array.

For example,
 this 8×3 array

3	1	4
5	3	5
2	3	8
3	3	8
0	2	8
1	6	9
5	1	0
7	4	9

Sort all n cells
in $n^{0.5+o(1)}$ seconds:

- Recursively sort quadrants in parallel, if $n > 1$.
- Sort each column in parallel.
- Sort each row in parallel.
- Sort each column in parallel.
- Sort each row in parallel.

With proper choice of left-to-right/right-to-left for each row, can prove that this sorts whole array.

For example, assume
this 8×8 array is

3	1	4	1	5	9		
5	3	5	8	9	7		
2	3	8	4	6	2		
3	3	8	3	2	7		
0	2	8	8	4	1		
1	6	9	3	9	9		
5	1	0	5	8	2		
7	4	9	4	4	5		

Sort all n cells
in $n^{0.5+o(1)}$ seconds:

- Recursively sort quadrants in parallel, if $n > 1$.
- Sort each column in parallel.
- Sort each row in parallel.
- Sort each column in parallel.
- Sort each row in parallel.

With proper choice of left-to-right/right-to-left for each row, can prove that this sorts whole array.

For example, assume that this 8×8 array is in cells:

3	1	4	1	5	9	2	6
5	3	5	8	9	7	9	3
2	3	8	4	6	2	6	4
3	3	8	3	2	7	9	5
0	2	8	8	4	1	9	7
1	6	9	3	9	9	3	7
5	1	0	5	8	2	0	9
7	4	9	4	4	5	9	2

Sort all n cells
in $n^{0.5+o(1)}$ seconds:

- Recursively sort quadrants in parallel, if $n > 1$.
- Sort each column in parallel.
- Sort each row in parallel.
- Sort each column in parallel.
- Sort each row in parallel.

With proper choice of left-to-right/right-to-left for each row, can prove that this sorts whole array.

For example, assume that this 8×8 array is in cells:

3	1	4	1	5	9	2	6
5	3	5	8	9	7	9	3
2	3	8	4	6	2	6	4
3	3	8	3	2	7	9	5
0	2	8	8	4	1	9	7
1	6	9	3	9	9	3	7
5	1	0	5	8	2	0	9
7	4	9	4	4	5	9	2

n cells
 $p(1)$ seconds:

recursively sort quadrants

in parallel, if $n > 1$.

sort each column in parallel.

sort each row in parallel.

sort each column in parallel.

sort each row in parallel.

proper choice of

left/right/right-to-left

row, can prove

algorithm sorts whole array.

For example, assume that
 this 8×8 array is in cells:

3	1	4	1	5	9	2	6
5	3	5	8	9	7	9	3
2	3	8	4	6	2	6	4
3	3	8	3	2	7	9	5
0	2	8	8	4	1	9	7
1	6	9	3	9	9	3	7
5	1	0	5	8	2	0	9
7	4	9	4	4	5	9	2

Recursive

top \rightarrow , l

1	1	2
3	3	3
3	4	4
5	8	8
1	1	0
4	4	3
7	6	5
9	9	8

For example, assume that
this 8×8 array is in cells:

3	1	4	1	5	9	2	6
5	3	5	8	9	7	9	3
2	3	8	4	6	2	6	4
3	3	8	3	2	7	9	5
0	2	8	8	4	1	9	7
1	6	9	3	9	9	3	7
5	1	0	5	8	2	0	9
7	4	9	4	4	5	9	2

Recursively sort qu
top \rightarrow , bottom \leftarrow

1	1	2	3	2	2
3	3	3	3	4	5
3	4	4	5	6	6
5	8	8	8	9	9
1	1	0	0	2	2
4	4	3	2	5	4
7	6	5	5	9	8
9	9	8	8	9	9

For example, assume that this 8×8 array is in cells:

3	1	4	1	5	9	2	6
5	3	5	8	9	7	9	3
2	3	8	4	6	2	6	4
3	3	8	3	2	7	9	5
0	2	8	8	4	1	9	7
1	6	9	3	9	9	3	7
5	1	0	5	8	2	0	9
7	4	9	4	4	5	9	2

el.

el.

Recursively sort quadrants, top \rightarrow , bottom \leftarrow :

1	1	2	3	2	2	2	3
3	3	3	3	4	5	5	6
3	4	4	5	6	6	7	7
5	8	8	8	9	9	9	9
1	1	0	0	2	2	1	0
4	4	3	2	5	4	4	3
7	6	5	5	9	8	7	7
9	9	8	8	9	9	9	9

For example, assume that this 8×8 array is in cells:

3	1	4	1	5	9	2	6
5	3	5	8	9	7	9	3
2	3	8	4	6	2	6	4
3	3	8	3	2	7	9	5
0	2	8	8	4	1	9	7
1	6	9	3	9	9	3	7
5	1	0	5	8	2	0	9
7	4	9	4	4	5	9	2

Recursively sort quadrants,
top \rightarrow , bottom \leftarrow :

1	1	2	3	2	2	2	3
3	3	3	3	4	5	5	6
3	4	4	5	6	6	7	7
5	8	8	8	9	9	9	9
1	1	0	0	2	2	1	0
4	4	3	2	5	4	4	3
7	6	5	5	9	8	7	7
9	9	8	8	9	9	9	9

Example, assume that
8 array is in cells:

1	5	9	2	6
8	9	7	9	3
4	6	2	6	4
3	2	7	9	5
8	4	1	9	7
3	9	9	3	7
5	8	2	0	9
4	4	5	9	2

Recursively sort quadrants,
top \rightarrow , bottom \leftarrow :

1	1	2	3	2	2	2	3
3	3	3	3	4	5	5	6
3	4	4	5	6	6	7	7
5	8	8	8	9	9	9	9
1	1	0	0	2	2	1	0
4	4	3	2	5	4	4	3
7	6	5	5	9	8	7	7
9	9	8	8	9	9	9	9

Sort each
in parallel

1	1	0
1	1	2
3	3	3
3	4	3
4	4	4
5	6	5
7	8	8
9	9	8

me that
in cells:

2	6
9	3
6	4
9	5
9	7
3	7
0	9
9	2

Recursively sort quadrants,
top \rightarrow , bottom \leftarrow :

1	1	2	3	2	2	2	3
3	3	3	3	4	5	5	6
3	4	4	5	6	6	7	7
5	8	8	8	9	9	9	9
1	1	0	0	2	2	1	0
4	4	3	2	5	4	4	3
7	6	5	5	9	8	7	7
9	9	8	8	9	9	9	9

Sort each column
in parallel:

1	1	0	0	2	2
1	1	2	2	2	2
3	3	3	3	4	4
3	4	3	3	5	5
4	4	4	5	6	6
5	6	5	5	9	8
7	8	8	8	9	9
9	9	8	8	9	9

Recursively sort quadrants,
top \rightarrow , bottom \leftarrow :

1	1	2	3	2	2	2	3
3	3	3	3	4	5	5	6
3	4	4	5	6	6	7	7
5	8	8	8	9	9	9	9
1	1	0	0	2	2	1	0
4	4	3	2	5	4	4	3
7	6	5	5	9	8	7	7
9	9	8	8	9	9	9	9

Sort each column
in parallel:

1	1	0	0	2	2	1	0
1	1	2	2	2	2	2	3
3	3	3	3	4	4	4	3
3	4	3	3	5	5	5	6
4	4	4	5	6	6	7	7
5	6	5	5	9	8	7	7
7	8	8	8	9	9	9	9
9	9	8	8	9	9	9	9

Recursively sort quadrants,
top \rightarrow , bottom \leftarrow :

1	1	2	3	2	2	2	3
3	3	3	3	4	5	5	6
3	4	4	5	6	6	7	7
5	8	8	8	9	9	9	9
1	1	0	0	2	2	1	0
4	4	3	2	5	4	4	3
7	6	5	5	9	8	7	7
9	9	8	8	9	9	9	9

Sort each column
in parallel:

1	1	0	0	2	2	1	0
1	1	2	2	2	2	2	3
3	3	3	3	4	4	4	3
3	4	3	3	5	5	5	6
4	4	4	5	6	6	7	7
5	6	5	5	9	8	7	7
7	8	8	8	9	9	9	9
9	9	8	8	9	9	9	9

ely sort quadrants,
bottom \leftarrow :

2	3	2	2	2	3
3	3	4	5	5	6
4	5	6	6	7	7
8	8	9	9	9	9
0	0	2	2	1	0
3	2	5	4	4	3
5	5	9	8	7	7
3	8	9	9	9	9

Sort each column
in parallel:

1	1	0	0	2	2	1	0
1	1	2	2	2	2	2	3
3	3	3	3	4	4	4	3
3	4	3	3	5	5	5	6
4	4	4	5	6	6	7	7
5	6	5	5	9	8	7	7
7	8	8	8	9	9	9	9
9	9	8	8	9	9	9	9

Sort each
alternate

0	0	0
3	2	2
3	3	3
6	5	5
4	4	4
9	8	7
7	8	8
9	9	9

quadrants,
:

2	3
5	6
7	7
9	9
1	0
4	3
7	7
9	9

Sort each column
in parallel:

1	1	0	0	2	2	1	0
1	1	2	2	2	2	2	3
3	3	3	3	4	4	4	3
3	4	3	3	5	5	5	6
4	4	4	5	6	6	7	7
5	6	5	5	9	8	7	7
7	8	8	8	9	9	9	9
9	9	8	8	9	9	9	9

Sort each row in p
alternately \leftarrow , \rightarrow :

0	0	0	1	1	1
3	2	2	2	2	2
3	3	3	3	3	4
6	5	5	5	4	3
4	4	4	5	6	6
9	8	7	7	6	5
7	8	8	8	9	9
9	9	9	9	9	9

h column
el:

0	0	2	2	1	0
2	2	2	2	2	3
3	3	4	4	4	3
3	3	5	5	5	6
4	5	6	6	7	7
5	5	9	8	7	7
3	8	9	9	9	9
3	8	9	9	9	9

47

Sort each row in parallel,
alternately \leftarrow , \rightarrow :

0	0	0	1	1	1	2	2
3	2	2	2	2	2	1	1
3	3	3	3	3	4	4	4
6	5	5	5	4	3	3	3
4	4	4	5	6	6	7	7
9	8	7	7	6	5	5	5
7	8	8	8	9	9	9	9
9	9	9	9	9	9	8	8

48

Sort each
in parallel

0	0	0
3	2	2
3	3	3
4	4	4
6	5	5
7	8	7
9	8	8
9	9	9

47

Sort each row in parallel,
alternately \leftarrow , \rightarrow :

1	0
2	3
4	3
5	6
7	7
7	7
9	9
9	9

0	0	0	1	1	1	2	2
3	2	2	2	2	2	1	1
3	3	3	3	3	4	4	4
6	5	5	5	4	3	3	3
4	4	4	5	6	6	7	7
9	8	7	7	6	5	5	5
7	8	8	8	9	9	9	9
9	9	9	9	9	9	8	8

48

Sort each column
in parallel:

0	0	0	1	1	1
3	2	2	2	2	2
3	3	3	3	3	3
4	4	4	5	4	4
6	5	5	5	6	5
7	8	7	7	6	6
9	8	8	8	9	9
9	9	9	9	9	9

h row in parallel,
ely \leftarrow , \rightarrow :

0	1	1	1	2	2
2	2	2	2	1	1
3	3	3	4	4	4
5	5	4	3	3	3
4	5	6	6	7	7
7	7	6	5	5	5
8	8	9	9	9	9
9	9	9	9	8	8

48

Sort each column
in parallel:

0	0	0	1	1	1	1	1
3	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3
4	4	4	5	4	4	4	4
6	5	5	5	6	5	5	5
7	8	7	7	6	6	7	7
9	8	8	8	9	9	8	8
9	9	9	9	9	9	9	9

49

Sort each
 \leftarrow or \rightarrow

0	0	0
2	2	2
3	3	3
4	4	4
5	5	5
6	6	7
8	8	8
9	9	9

parallel,

2	2
1	1
4	4
3	3
7	7
5	5
9	9
8	8

Sort each column
in parallel:

0	0	0	1	1	1	1	1
3	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3
4	4	4	5	4	4	4	4
6	5	5	5	6	5	5	5
7	8	7	7	6	6	7	7
9	8	8	8	9	9	8	8
9	9	9	9	9	9	9	9

Sort each row in p

← or → as desired

0	0	0	1	1	1
2	2	2	2	2	2
3	3	3	3	3	3
4	4	4	4	4	4
5	5	5	5	5	5
6	6	7	7	7	7
8	8	8	8	8	9
9	9	9	9	9	9

h column
el:

0	1	1	1	1	1
2	2	2	2	2	2
3	3	3	3	3	3
4	5	4	4	4	4
5	5	6	5	5	5
7	7	6	6	7	7
8	8	9	9	8	8
9	9	9	9	9	9

Sort each row in parallel,
← or → as desired:

0	0	0	1	1	1	1	1
2	2	2	2	2	2	2	3
3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	5
5	5	5	5	5	5	6	6
6	6	7	7	7	7	7	8
8	8	8	8	8	9	9	9
9	9	9	9	9	9	9	9

Chips ar
towards
parallelis
GPUs: p
Old Xeo
New Xeo

Sort each row in parallel,
 \leftarrow or \rightarrow as desired:

1	1
2	2
3	3
4	4
5	5
7	7
8	8
9	9

0	0	0	1	1	1	1	1
2	2	2	2	2	2	2	3
3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	5
5	5	5	5	5	5	6	6
6	6	7	7	7	7	7	8
8	8	8	8	8	9	9	9
9	9	9	9	9	9	9	9

Chips are in fact e
 towards having thi
 parallelism and co

GPUs: parallel +

Old Xeon Phi: pa

New Xeon Phi: pa

Sort each row in parallel,
 \leftarrow or \rightarrow as desired:

0	0	0	1	1	1	1	1
2	2	2	2	2	2	2	3
3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	5
5	5	5	5	5	5	6	6
6	6	7	7	7	7	7	8
8	8	8	8	8	9	9	9
9	9	9	9	9	9	9	9

Chips are in fact evolving
 towards having this much
 parallelism and communicat

GPUs: parallel + global RA

Old Xeon Phi: parallel + rin

New Xeon Phi: parallel + m

Sort each row in parallel,
 \leftarrow or \rightarrow as desired:

0	0	0	1	1	1	1	1
2	2	2	2	2	2	2	3
3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	5
5	5	5	5	5	5	6	6
6	6	7	7	7	7	7	8
8	8	8	8	8	9	9	9
9	9	9	9	9	9	9	9

Chips are in fact evolving
 towards having this much
 parallelism and communication.

GPUs: parallel + global RAM.

Old Xeon Phi: parallel + ring.

New Xeon Phi: parallel + mesh.

Sort each row in parallel,
 \leftarrow or \rightarrow as desired:

0	0	0	1	1	1	1	1
2	2	2	2	2	2	2	3
3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	5
5	5	5	5	5	5	6	6
6	6	7	7	7	7	7	8
8	8	8	8	8	9	9	9
9	9	9	9	9	9	9	9

Chips are in fact evolving
 towards having this much
 parallelism and communication.

GPUs: parallel + global RAM.

Old Xeon Phi: parallel + ring.

New Xeon Phi: parallel + mesh.

Algorithm designers

don't even get the right exponent
 without taking this into account.

Sort each row in parallel,
 \leftarrow or \rightarrow as desired:

0	0	0	1	1	1	1	1
2	2	2	2	2	2	2	3
3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	5
5	5	5	5	5	5	6	6
6	6	7	7	7	7	7	8
8	8	8	8	8	9	9	9
9	9	9	9	9	9	9	9

Chips are in fact evolving
 towards having this much
 parallelism and communication.

GPUs: parallel + global RAM.

Old Xeon Phi: parallel + ring.

New Xeon Phi: parallel + mesh.

Algorithm designers

don't even get the right exponent
 without taking this into account.

Shock waves from subroutines
 into high-level algorithm design.