

Smartphone/tablet CPUs

iPad 1 (2010) was the first popular tablet: more than 15 million sold.

iPad 1 contains 45nm

Apple A4 system-on-chip.

Apple A4 contains

1GHz ARM Cortex-A8 CPU core
+ PowerVR SGX 535 GPU.

Cortex-A8 CPU core (2005) supports ARMv7-A insn set, including NEON vector insns.

Apple A4 also appeared in iPhone 4 (2010).

45nm 1GHz Samsung Exynos 3110 in Samsung Galaxy S (2010) contains Cortex-A8 CPU core.

45nm 1GHz TI OMAP3630 in Motorola Droid X (2010) contains Cortex-A8 CPU core.

65nm 800MHz Freescale i.MX50 in Amazon Kindle 4 (2011) contains Cortex-A8 CPU core.

Phone/tablet CPUs

(2010) was the

popular tablet:

more than 15 million sold.

It contains 45nm

ARM Cortex-A8 system-on-chip.

The Cortex-A8 contains

ARM Cortex-A8 CPU core

and PowerVR SGX 535 GPU.

The Cortex-A8 CPU core (2005)

implements ARMv7-A insn set,

including NEON vector insns.

1

Apple A4 also appeared
in iPhone 4 (2010).

45nm 1GHz Samsung Exynos
3110 in Samsung Galaxy S (2010)
contains Cortex-A8 CPU core.

45nm 1GHz TI OMAP3630 in
Motorola Droid X (2010)
contains Cortex-A8 CPU core.

65nm 800MHz Freescale i.MX50
in Amazon Kindle 4 (2011)
contains Cortex-A8 CPU core.

2

ARM de
supporti

Cortex-A

Cortex-A

Cortex-A

Cortex-A

Cortex-A

Also som

A9, A15

cores are

tries to

compens

CPUs

the

t:

on sold.

5nm

on-chip.

x-A8 CPU core

535 GPU.

ore (2005)

A insn set,

ector insns.

1

Apple A4 also appeared
in iPhone 4 (2010).

45nm 1GHz Samsung Exynos
3110 in Samsung Galaxy S (2010)
contains Cortex-A8 CPU core.

45nm 1GHz TI OMAP3630 in
Motorola Droid X (2010)
contains Cortex-A8 CPU core.

65nm 800MHz Freescale i.MX50
in Amazon Kindle 4 (2011)
contains Cortex-A8 CPU core.

2

ARM designed mo
supporting same A
Cortex-A9 (2007),
Cortex-A5 (2009),
Cortex-A15 (2010),
Cortex-A7 (2011),
Cortex-A17 (2014)

Also some larger 6

A9, A15, A17, and
cores are “out of c
tries to reorder ins
compensate for du

1

Apple A4 also appeared in iPhone 4 (2010).

45nm 1GHz Samsung Exynos 3110 in Samsung Galaxy S (2010) contains Cortex-A8 CPU core.

45nm 1GHz TI OMAP3630 in Motorola Droid X (2010) contains Cortex-A8 CPU core.

65nm 800MHz Freescale i.MX50 in Amazon Kindle 4 (2011) contains Cortex-A8 CPU core.

2

ARM designed more cores supporting same ARMv7-A
Cortex-A9 (2007),
Cortex-A5 (2009),
Cortex-A15 (2010),
Cortex-A7 (2011),
Cortex-A17 (2014), etc.

Also some larger 64-bit cores

A9, A15, A17, and some 64-bit
cores are "out of order": CPU
tries to reorder instructions to
compensate for dumb compilers

core

S.

Apple A4 also appeared
in iPhone 4 (2010).

45nm 1GHz Samsung Exynos
3110 in Samsung Galaxy S (2010)
contains Cortex-A8 CPU core.

45nm 1GHz TI OMAP3630 in
Motorola Droid X (2010)
contains Cortex-A8 CPU core.

65nm 800MHz Freescale i.MX50
in Amazon Kindle 4 (2011)
contains Cortex-A8 CPU core.

ARM designed more cores
supporting same ARMv7-A insns:
Cortex-A9 (2007),
Cortex-A5 (2009),
Cortex-A15 (2010),
Cortex-A7 (2011),
Cortex-A17 (2014), etc.

Also some larger 64-bit cores.

A9, A15, A17, and some 64-bit
cores are “out of order”: CPU
tries to reorder instructions to
compensate for dumb compilers.

4 also appeared
e 4 (2010).

GHz Samsung Exynos
Samsung Galaxy S (2010)
Cortex-A8 CPU core.

GHz TI OMAP3630 in
a Droid X (2010)
Cortex-A8 CPU core.

00MHz Freescale i.MX50
on Kindle 4 (2011)
Cortex-A8 CPU core.

2

ARM designed more cores
supporting same ARMv7-A insns:
Cortex-A9 (2007),
Cortex-A5 (2009),
Cortex-A15 (2010),
Cortex-A7 (2011),
Cortex-A17 (2014), etc.

Also some larger 64-bit cores.

A9, A15, A17, and some 64-bit
cores are “out of order”: CPU
tries to reorder instructions to
compensate for dumb compilers.

3

A5, A7,
fewer ins

appeared

).

ung Exynos

Galaxy S (2010)

8 CPU core.

MAP3630 in

(2010)

8 CPU core.

eescale i.MX50

4 (2011)

8 CPU core.

2

ARM designed more cores
supporting same ARMv7-A insns:

Cortex-A9 (2007),

Cortex-A5 (2009),

Cortex-A15 (2010),

Cortex-A7 (2011),

Cortex-A17 (2014), etc.

Also some larger 64-bit cores.

A9, A15, A17, and some 64-bit
cores are “out of order”: CPU
tries to reorder instructions to
compensate for dumb compilers.

3

A5, A7, original A

fewer insns at onc

2

ARM designed more cores supporting same ARMv7-A insns:

Cortex-A9 (2007),
Cortex-A5 (2009),
Cortex-A15 (2010),
Cortex-A7 (2011),
Cortex-A17 (2014), etc.

Also some larger 64-bit cores.

A9, A15, A17, and some 64-bit cores are “out of order”: CPU tries to reorder instructions to compensate for dumb compilers.

3

A5, A7, original A8 are in-order, fewer insns at once.

ARM designed more cores supporting same ARMv7-A insns:

Cortex-A9 (2007),

Cortex-A5 (2009),

Cortex-A15 (2010),

Cortex-A7 (2011),

Cortex-A17 (2014), etc.

Also some larger 64-bit cores.

A9, A15, A17, and some 64-bit cores are “out of order”: CPU tries to reorder instructions to compensate for dumb compilers.

A5, A7, original A8 are in-order, fewer insns at once.

ARM designed more cores supporting same ARMv7-A insns:
Cortex-A9 (2007),
Cortex-A5 (2009),
Cortex-A15 (2010),
Cortex-A7 (2011),
Cortex-A17 (2014), etc.

Also some larger 64-bit cores.

A9, A15, A17, and some 64-bit cores are “out of order”: CPU tries to reorder instructions to compensate for dumb compilers.

A5, A7, original A8 are in-order, fewer insns at once. \Rightarrow Simpler, cheaper, more energy-efficient.

ARM designed more cores supporting same ARMv7-A insns:

Cortex-A9 (2007),

Cortex-A5 (2009),

Cortex-A15 (2010),

Cortex-A7 (2011),

Cortex-A17 (2014), etc.

Also some larger 64-bit cores.

A9, A15, A17, and some 64-bit cores are “out of order”: CPU tries to reorder instructions to compensate for dumb compilers.

A5, A7, original A8 are in-order, fewer insns at once. \Rightarrow Simpler, cheaper, more energy-efficient.

More than one billion Cortex-A7 devices have been sold.

Popular in low-cost and mid-range smartphones: Mobiistar Buddy, Mobiistar Kool, Mobiistar LAI Z1, Samsung Galaxy J1 Ace Neo, etc.

Also used in typical TV boxes, Sony SmartWatch 3, Samsung Gear S2, Raspberry Pi 2, etc.

signed more cores
ng same ARMv7-A insns:
A9 (2007),
A5 (2009),
A15 (2010),
A7 (2011),
A17 (2014), etc.

ne larger 64-bit cores.

, A17, and some 64-bit
e “out of order”: CPU
reorder instructions to
sate for dumb compilers.

3

A5, A7, original A8 are in-order,
fewer insns at once. \Rightarrow Simpler,
cheaper, more energy-efficient.

More than one billion Cortex-A7
devices have been sold.

Popular in low-cost and mid-range
smartphones: Mobiistar Buddy,
Mobiistar Kool, Mobiistar LAI Z1,
Samsung Galaxy J1 Ace Neo, etc.

Also used in typical TV boxes,
Sony SmartWatch 3, Samsung
Gear S2, Raspberry Pi 2, etc.

4

NEON c
Basic AF
16 32-bi
Optiona
16 128-b
Cortex-A
(and Co
and Qua
and Qua
always h
Cortex-A
sometim

3

A5, A7, original A8 are in-order, fewer insns at once. \Rightarrow Simpler, cheaper, more energy-efficient.

More than one billion Cortex-A7 devices have been sold.

Popular in low-cost and mid-range smartphones: Mobiistar Buddy, Mobiistar Kool, Mobiistar LAI Z1, Samsung Galaxy J1 Ace Neo, etc.

Also used in typical TV boxes, Sony SmartWatch 3, Samsung Gear S2, Raspberry Pi 2, etc.

4

NEON crypto

Basic ARM insn set
16 32-bit registers

Optional NEON extension
16 128-bit registers

Cortex-A7 and Cortex-A9
(and Cortex-A15 and Cortex-A53)
and Qualcomm Snapdragon
and Qualcomm Kraken
always have NEON

Cortex-A5 and Cortex-A8
sometimes have NEON

3

insns:

A5, A7, original A8 are in-order,
fewer insns at once. \Rightarrow Simpler,
cheaper, more energy-efficient.

More than one billion Cortex-A7
devices have been sold.

Popular in low-cost and mid-range
smartphones: Mobiistar Buddy,
Mobiistar Kool, Mobiistar LAI Z1,
Samsung Galaxy J1 Ace Neo, etc.

Also used in typical TV boxes,
Sony SmartWatch 3, Samsung
Gear S2, Raspberry Pi 2, etc.

4

NEON crypto

Basic ARM insn set uses
16 32-bit registers: 512 bits

Optional NEON extension u
16 128-bit registers: 2048 b

Cortex-A7 and Cortex-A8
(and Cortex-A15 and Cortex
and Qualcomm Scorpion
and Qualcomm Krait)
always have NEON insns.

Cortex-A5 and Cortex-A9
sometimes have NEON insns.

A5, A7, original A8 are in-order, fewer insns at once. \Rightarrow Simpler, cheaper, more energy-efficient.

More than one billion Cortex-A7 devices have been sold.

Popular in low-cost and mid-range smartphones: Mobiistar Buddy, Mobiistar Kool, Mobiistar LAI Z1, Samsung Galaxy J1 Ace Neo, etc.

Also used in typical TV boxes, Sony SmartWatch 3, Samsung Gear S2, Raspberry Pi 2, etc.

NEON crypto

Basic ARM insn set uses 16 32-bit registers: 512 bits.

Optional NEON extension uses 16 128-bit registers: 2048 bits.

Cortex-A7 and Cortex-A8 (and Cortex-A15 and Cortex-A17 and Qualcomm Scorpion and Qualcomm Krait) always have NEON insns.

Cortex-A5 and Cortex-A9 sometimes have NEON insns.

original A8 are in-order,
insns at once. \Rightarrow Simpler,
more energy-efficient.

more than one billion Cortex-A7
processors have been sold.

found in low-cost and mid-range
phones: Mobiistar Buddy,
Mobiistar Kool, Mobiistar LAI Z1,
Samsung Galaxy J1 Ace Neo, etc.

found in typical TV boxes,
SmartWatch 3, Samsung
Galaxy Tab 2, Raspberry Pi 2, etc.

4

NEON crypto

Basic ARM insn set uses
16 32-bit registers: 512 bits.

Optional NEON extension uses
16 128-bit registers: 2048 bits.

Cortex-A7 and Cortex-A8
(and Cortex-A15 and Cortex-A17
and Qualcomm Scorpion
and Qualcomm Krait)
always have NEON insns.

Cortex-A5 and Cortex-A9
sometimes have NEON insns.

5

2012 Benchmark
“NEON”
new Cortex-A7
for various
operations

e.g. Curve25519
460200 cycles
498284 cycles

Compare
cycles on
for NIST
9 million
4.8 million
3.9 million

4

8 are in-order,
e. \Rightarrow Simpler,
energy-efficient.

million Cortex-A7
sold.

st and mid-range
obiistar Buddy,
obiistar LAI Z1,
1 Ace Neo, etc.

al TV boxes,
3, Samsung
y Pi 2, etc.

NEON crypto

Basic ARM insn set uses
16 32-bit registers: 512 bits.

Optional NEON extension uses
16 128-bit registers: 2048 bits.

Cortex-A7 and Cortex-A8
(and Cortex-A15 and Cortex-A17
and Qualcomm Scorpion
and Qualcomm Krait)
always have NEON insns.

Cortex-A5 and Cortex-A9
sometimes have NEON insns.

5

2012 Bernstein–Scott
“NEON crypto” so
new Cortex-A8 spe
for various crypto

e.g. Curve25519 E
460200 cycles on C
498284 cycles on C

Compare to OpenS
cycles on Cortex-A
for NIST P-256 EC
9 million for Open
4.8 million for Ope
3.9 million for Ope

4

NEON crypto

Basic ARM insn set uses
16 32-bit registers: 512 bits.

Optional NEON extension uses
16 128-bit registers: 2048 bits.

Cortex-A7 and Cortex-A8
(and Cortex-A15 and Cortex-A17
and Qualcomm Scorpion
and Qualcomm Krait)
always have NEON insns.

Cortex-A5 and Cortex-A9
sometimes have NEON insns.

5

2012 Bernstein–Schwabe
“NEON crypto” software:
new Cortex-A8 speed record
for various crypto primitives

e.g. Curve25519 ECDH:
460200 cycles on Cortex-A8.
498284 cycles on Cortex-A8.

Compare to OpenSSL
cycles on Cortex-A8-slow
for NIST P-256 ECDH:

9 million for OpenSSL 0.9.8
4.8 million for OpenSSL 1.0
3.9 million for OpenSSL 1.0

NEON crypto

Basic ARM insn set uses
16 32-bit registers: 512 bits.

Optional NEON extension uses
16 128-bit registers: 2048 bits.

Cortex-A7 and Cortex-A8
(and Cortex-A15 and Cortex-A17
and Qualcomm Scorpion
and Qualcomm Krait)
always have NEON insns.

Cortex-A5 and Cortex-A9
sometimes have NEON insns.

2012 Bernstein–Schwabe

“NEON crypto” software:
new Cortex-A8 speed records
for various crypto primitives.

e.g. Curve25519 ECDH:
460200 cycles on Cortex-A8-fast,
498284 cycles on Cortex-A8-slow.

Compare to OpenSSL
cycles on Cortex-A8-slow
for NIST P-256 ECDH:
9 million for OpenSSL 0.9.8k.
4.8 million for OpenSSL 1.0.1c.
3.9 million for OpenSSL 1.0.2j.

crypto

ARM insn set uses
t registers: 512 bits.

NEON extension uses
bit registers: 2048 bits.

A7 and Cortex-A8
Cortex-A15 and Cortex-A17

Qualcomm Scorpion
(Qualcomm Krait)

have NEON insns.

A5 and Cortex-A9

es have NEON insns.

5

2012 Bernstein–Schwabe

“NEON crypto” software:
new Cortex-A8 speed records
for various crypto primitives.

e.g. Curve25519 ECDH:
460200 cycles on Cortex-A8-fast,
498284 cycles on Cortex-A8-slow.

Compare to OpenSSL
cycles on Cortex-A8-slow
for NIST P-256 ECDH:
9 million for OpenSSL 0.9.8k.
4.8 million for OpenSSL 1.0.1c.
3.9 million for OpenSSL 1.0.2j.

6

NEON in

4x $a = b$
is a vect

$a[0]$

$a[1]$

$a[2]$

$a[3]$

5

2012 Bernstein–Schwabe

“NEON crypto” software:

new Cortex-A8 speed records
for various crypto primitives.

e.g. Curve25519 ECDH:

460200 cycles on Cortex-A8-fast,
498284 cycles on Cortex-A8-slow.

Compare to OpenSSL

cycles on Cortex-A8-slow

for NIST P-256 ECDH:

9 million for OpenSSL 0.9.8k.

4.8 million for OpenSSL 1.0.1c.

3.9 million for OpenSSL 1.0.2j.

6

NEON instructions

$4 \times a = b + c$

is a vector of 4 32

$a[0] = b[0] +$

$a[1] = b[1] +$

$a[2] = b[2] +$

$a[3] = b[3] +$

5

2012 Bernstein–Schwabe

“NEON crypto” software:

new Cortex-A8 speed records
for various crypto primitives.

e.g. Curve25519 ECDH:

460200 cycles on Cortex-A8-fast,
498284 cycles on Cortex-A8-slow.

Compare to OpenSSL

cycles on Cortex-A8-slow

for NIST P-256 ECDH:

9 million for OpenSSL 0.9.8k.

4.8 million for OpenSSL 1.0.1c.

3.9 million for OpenSSL 1.0.2j.

6

NEON instructions

$4x\ a = b + c$

is a vector of 4 32-bit additi

$a[0] = b[0] + c[0];$

$a[1] = b[1] + c[1];$

$a[2] = b[2] + c[2];$

$a[3] = b[3] + c[3].$

2012 Bernstein–Schwabe

“NEON crypto” software:

new Cortex-A8 speed records
for various crypto primitives.

e.g. Curve25519 ECDH:

460200 cycles on Cortex-A8-fast,
498284 cycles on Cortex-A8-slow.

Compare to OpenSSL

cycles on Cortex-A8-slow

for NIST P-256 ECDH:

9 million for OpenSSL 0.9.8k.

4.8 million for OpenSSL 1.0.1c.

3.9 million for OpenSSL 1.0.2j.

NEON instructions

$4x\ a = b + c$

is a vector of 4 32-bit additions:

$$a[0] = b[0] + c[0];$$

$$a[1] = b[1] + c[1];$$

$$a[2] = b[2] + c[2];$$

$$a[3] = b[3] + c[3].$$

2012 Bernstein–Schwabe

“NEON crypto” software:

new Cortex-A8 speed records
for various crypto primitives.

e.g. Curve25519 ECDH:

460200 cycles on Cortex-A8-fast,
498284 cycles on Cortex-A8-slow.

Compare to OpenSSL

cycles on Cortex-A8-slow

for NIST P-256 ECDH:

9 million for OpenSSL 0.9.8k.

4.8 million for OpenSSL 1.0.1c.

3.9 million for OpenSSL 1.0.2j.

NEON instructions

$4 \times a = b + c$

is a vector of 4 32-bit additions:

$$a[0] = b[0] + c[0];$$

$$a[1] = b[1] + c[1];$$

$$a[2] = b[2] + c[2];$$

$$a[3] = b[3] + c[3].$$

Cortex-A8 NEON arithmetic unit
can do this every cycle.

2012 Bernstein–Schwabe

“NEON crypto” software:

new Cortex-A8 speed records
for various crypto primitives.

e.g. Curve25519 ECDH:

460200 cycles on Cortex-A8-fast,
498284 cycles on Cortex-A8-slow.

Compare to OpenSSL

cycles on Cortex-A8-slow

for NIST P-256 ECDH:

9 million for OpenSSL 0.9.8k.

4.8 million for OpenSSL 1.0.1c.

3.9 million for OpenSSL 1.0.2j.

NEON instructions

$4x\ a = b + c$

is a vector of 4 32-bit additions:

$$a[0] = b[0] + c[0];$$

$$a[1] = b[1] + c[1];$$

$$a[2] = b[2] + c[2];$$

$$a[3] = b[3] + c[3].$$

Cortex-A8 NEON arithmetic unit
can do this every cycle.

Stage N2: reads b and c.

Stage N3: performs addition.

Stage N4: a is ready.



rnstein–Schwabe

crypto” software:

Cortex-A8 speed records

us crypto primitives.

ve25519 ECDH:

cycles on Cortex-A8-fast,

cycles on Cortex-A8-slow.

e to OpenSSL

n Cortex-A8-slow

P-256 ECDH:

n for OpenSSL 0.9.8k.

on for OpenSSL 1.0.1c.

on for OpenSSL 1.0.2j.

NEON instructions

$4x\ a = b + c$

is a vector of 4 32-bit additions:

$a[0] = b[0] + c[0];$

$a[1] = b[1] + c[1];$

$a[2] = b[2] + c[2];$

$a[3] = b[3] + c[3].$

Cortex-A8 NEON arithmetic unit
can do this every cycle.

Stage N2: reads b and c.

Stage N3: performs addition.

Stage N4: a is ready.

ADD $\xrightarrow{2\ \text{cycles}}$ ADD $\xrightarrow{2\ \text{cycles}}$ ADD

$4x\ a = b$

is a vect

$a[0]$

$a[1]$

$a[2]$

$a[3]$

Stage N

Stage N

Stage N

Stage N

ADD $\xrightarrow{2}$

Also log

NEON instructions

$$4x \ a = b + c$$

is a vector of 4 32-bit additions:

$$a[0] = b[0] + c[0];$$

$$a[1] = b[1] + c[1];$$

$$a[2] = b[2] + c[2];$$

$$a[3] = b[3] + c[3].$$

Cortex-A8 NEON arithmetic unit
can do this every cycle.

Stage N2: reads b and c.

Stage N3: performs addition.

Stage N4: a is ready.



$$4x \ a = b - c$$

is a vector of 4 32

$$a[0] = b[0] -$$

$$a[1] = b[1] -$$

$$a[2] = b[2] -$$


$$a[3] = b[3] -$$

Stage N1: reads c

Stage N2: reads b

Stage N3: perform

Stage N4: a is rea



Also logic insns, sh

NEON instructions

$$4x \ a = b + c$$

is a vector of 4 32-bit additions:

$$a[0] = b[0] + c[0];$$

$$a[1] = b[1] + c[1];$$

$$a[2] = b[2] + c[2];$$

$$a[3] = b[3] + c[3].$$

Cortex-A8 NEON arithmetic unit can do this every cycle.

Stage N2: reads b and c.

Stage N3: performs addition.

Stage N4: a is ready.



$$4x \ a = b - c$$

is a vector of 4 32-bit subtra

$$a[0] = b[0] - c[0];$$

$$a[1] = b[1] - c[1];$$

$$a[2] = b[2] - c[2];$$


$$a[3] = b[3] - c[3].$$

Stage N1: reads c.

Stage N2: reads b, negates

Stage N3: performs addition

Stage N4: a is ready.



Also logic insns, shifts, etc.

NEON instructions

$$4x \ a = b + c$$

is a vector of 4 32-bit additions:

$$a[0] = b[0] + c[0];$$

$$a[1] = b[1] + c[1];$$

$$a[2] = b[2] + c[2];$$

$$a[3] = b[3] + c[3].$$

Cortex-A8 NEON arithmetic unit can do this every cycle.

Stage N2: reads b and c.

Stage N3: performs addition.

Stage N4: a is ready.



$$4x \ a = b - c$$

is a vector of 4 32-bit subtractions:

$$a[0] = b[0] - c[0];$$

$$a[1] = b[1] - c[1];$$

$$a[2] = b[2] - c[2];$$


$$a[3] = b[3] - c[3].$$

Stage N1: reads c.

Stage N2: reads b, negates c.

Stage N3: performs addition.

Stage N4: a is ready.



Also logic insns, shifts, etc.

Instructions

$b + c$

or of 4 32-bit additions:

$$= b[0] + c[0];$$

$$= b[1] + c[1];$$

$$= b[2] + c[2];$$

$$= b[3] + c[3].$$


A8 NEON arithmetic unit

this every cycle.

2: reads b and c.

3: performs addition.

4: a is ready.



7

$$4x \ a = b - c$$

is a vector of 4 32-bit subtractions:

$$a[0] = b[0] - c[0];$$

$$a[1] = b[1] - c[1];$$

$$a[2] = b[2] - c[2];$$


$$a[3] = b[3] - c[3].$$

Stage N1: reads c.

Stage N2: reads b, negates c.

Stage N3: performs addition.

Stage N4: a is ready.



Also logic insns, shifts, etc.

8

Multiplic

$c[0,1]$

$c[2,3]$

Two cyc

Multiply

$c[0,1]$

$c[2,3]$

Also two

Stage N

Stage N

Stage N

:

:

Stage N

$$4x \ a = b - c$$

is a vector of 4 32-bit subtractions:

$$a[0] = b[0] - c[0];$$

$$a[1] = b[1] - c[1];$$

$$a[2] = b[2] - c[2];$$

$$a[3] = b[3] - c[3].$$

Stage N1: reads c.

Stage N2: reads b, negates c.

Stage N3: performs addition.

Stage N4: a is ready.

$$\text{ADD} \xrightarrow{2 \text{ or } 3 \text{ cycles}} \text{SUB}$$

Also logic insns, shifts, etc.

$$\text{ADD} \xrightarrow{2 \text{ cycles}} \text{ADD}$$

Multiplication insns

$c[0,1] = a[0]$ shift

$c[2,3] = a[1]$ shift

Two cycles on Core

Multiply-accumulate

$c[0,1] += a[0]$ shift

$c[2,3] += a[1]$ shift

Also two cycles on

Stage N1: reads b

Stage N2: reads a

Stage N3: reads c

⋮

Stage N8: c is ready

$4x \ a = b - c$

is a vector of 4 32-bit subtractions:

$$a[0] = b[0] - c[0];$$

$$a[1] = b[1] - c[1];$$

$$a[2] = b[2] - c[2];$$

$$a[3] = b[3] - c[3].$$

Stage N1: reads c.

Stage N2: reads b, negates c.

Stage N3: performs addition.

Stage N4: a is ready.

ADD $\xrightarrow{\text{2 or 3 cycles}}$ SUB

Also logic insns, shifts, etc.

es
→ ADD

Multiplication insn:

$$c[0,1] = a[0] \text{ signed} * b[0,1]$$

$$c[2,3] = a[1] \text{ signed} * b[2,3]$$

Two cycles on Cortex-A8.

Multiply-accumulate insn:

$$c[0,1] += a[0] \text{ signed} * b[0,1]$$

$$c[2,3] += a[1] \text{ signed} * b[2,3]$$

Also two cycles on Cortex-A8.

Stage N1: reads b.

Stage N2: reads a.

Stage N3: reads c if accumulate

⋮

Stage N8: c is ready.

$4x\ a = b - c$

is a vector of 4 32-bit subtractions:

$$a[0] = b[0] - c[0];$$

$$a[1] = b[1] - c[1];$$

$$a[2] = b[2] - c[2];$$

$$a[3] = b[3] - c[3].$$

Stage N1: reads c.

Stage N2: reads b, negates c.

Stage N3: performs addition.

Stage N4: a is ready.

ADD $\xrightarrow{\text{2 or 3 cycles}}$ SUB

Also logic insns, shifts, etc.

Multiplication insn:

$$c[0,1] = a[0] \text{ signed} * b[0];$$

$$c[2,3] = a[1] \text{ signed} * b[1]$$

Two cycles on Cortex-A8.

Multiply-accumulate insn:

$$c[0,1] += a[0] \text{ signed} * b[0];$$

$$c[2,3] += a[1] \text{ signed} * b[1]$$

Also two cycles on Cortex-A8.

Stage N1: reads b.

Stage N2: reads a.

Stage N3: reads c if accumulate.

⋮

Stage N8: c is ready.

$b - c$

or of 4 32-bit subtractions:

$$= b[0] - c[0];$$

$$= b[1] - c[1];$$

$$= b[2] - c[2];$$

$$= b[3] - c[3].$$

1: reads c.

2: reads b, negates c.

3: performs addition.

4: a is ready.

or 3 cycles
 $\xrightarrow{\hspace{1.5cm}}$ SUB

ic insns, shifts, etc.

Multiplication insn:

$$c[0,1] = a[0] \text{ signed} * b[0];$$

$$c[2,3] = a[1] \text{ signed} * b[1]$$

Two cycles on Cortex-A8.

Multiply-accumulate insn:

$$c[0,1] += a[0] \text{ signed} * b[0];$$

$$c[2,3] += a[1] \text{ signed} * b[1]$$

Also two cycles on Cortex-A8.

Stage N1: reads b.

Stage N2: reads a.

Stage N3: reads c if accumulate.

⋮

Stage N8: c is ready.

Typical s

c[0,1]

c[2,3]

c[0,1]

c[2,3]

c[0,1]

c[2,3]

Cortex-A

Reads c

-bit subtractions:

c[0];

c[1];

c[2];

c[3].

o, negates c.

ns addition.

dy.

- SUB

hifts, etc.

Multiplication insn:

c[0,1] = a[0] signed* b[0];

c[2,3] = a[1] signed* b[1]

Two cycles on Cortex-A8.

Multiply-accumulate insn:

c[0,1] += a[0] signed* b[0];

c[2,3] += a[1] signed* b[1]

Also two cycles on Cortex-A8.

Stage N1: reads b.

Stage N2: reads a.

Stage N3: reads c if accumulate.

:

Stage N8: c is ready.

Typical sequence o

c[0,1] = a[0] si

c[2,3] = a[1] si

c[0,1] += e[2] s

c[2,3] += e[3] s

c[0,1] += g[0] s

c[2,3] += g[1] s

Cortex-A8 recogni

Reads c in N6 inst

actions:

Multiplication insn:

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

Two cycles on Cortex-A8.

Multiply-accumulate insn:

```
c[0,1] += a[0] signed* b[0];
```

```
c[2,3] += a[1] signed* b[1]
```

Also two cycles on Cortex-A8.

Stage N1: reads b.

Stage N2: reads a.

Stage N3: reads c if accumulate.

⋮

Stage N8: c is ready.

Typical sequence of three insns:

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1];
```

```
c[0,1] += e[2] signed* f[0];
```

```
c[2,3] += e[3] signed* f[1];
```

```
c[0,1] += g[0] signed* h[0];
```

```
c[2,3] += g[1] signed* h[1];
```

Cortex-A8 recognizes this pattern.

Reads c in N6 instead of N3.

Multiplication insn:

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

Two cycles on Cortex-A8.

Multiply-accumulate insn:

```
c[0,1] += a[0] signed* b[0];
```

```
c[2,3] += a[1] signed* b[1]
```

Also two cycles on Cortex-A8.

Stage N1: reads b.

Stage N2: reads a.

Stage N3: reads c if accumulate.

⋮

Stage N8: c is ready.

Typical sequence of three insns:

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

```
c[0,1] += e[2] signed* f[2];
```

```
c[2,3] += e[3] signed* f[3]
```

```
c[0,1] += g[0] signed* h[2];
```

```
c[2,3] += g[1] signed* h[3]
```

Cortex-A8 recognizes this pattern.

Reads c in N6 instead of N3.

ation insn:

```
= a[0] signed* b[0];
```

```
= a[1] signed* b[1]
```

cles on Cortex-A8.

-accumulate insn:

```
+= a[0] signed* b[0];
```

```
+= a[1] signed* b[1]
```

o cycles on Cortex-A8.

1: reads b.

2: reads a.

3: reads c if accumulate.

8: c is ready.

Typical sequence of three insns:

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

```
c[0,1] += e[2] signed* f[2];
```

```
c[2,3] += e[3] signed* f[3]
```

```
c[0,1] += g[0] signed* h[2];
```

```
c[2,3] += g[1] signed* h[3]
```

Cortex-A8 recognizes this pattern.

Reads c in N6 instead of N3.

Time	N1
1	<i>b</i>
2	
3	<i>f</i>
4	
5	<i>h</i>
6	
7	
8	
9	
10	
11	
12	

n:

igned* b[0];

igned* b[1]

Cortex-A8.

te insn:

igned* b[0];

igned* b[1]

Cortex-A8.

.

.

if accumulate.

dy.

Typical sequence of three insns:

`c[0,1] = a[0] signed* b[0];`

`c[2,3] = a[1] signed* b[1]`

`c[0,1] += e[2] signed* f[2];`

`c[2,3] += e[3] signed* f[3]`

`c[0,1] += g[0] signed* h[2];`

`c[2,3] += g[1] signed* h[3]`

Cortex-A8 recognizes this pattern.

Reads c in N6 instead of N3.

Time	N1	N2	N3	N4
1	<i>b</i>			
2		<i>a</i>		
3	<i>f</i>		×	
4		<i>e</i>		×
5	<i>h</i>		×	
6		<i>g</i>		×
7			×	
8				×
9				
10				
11				
12				

Typical sequence of three insns:

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

```
c[0,1] += e[2] signed* f[2];
```

```
c[2,3] += e[3] signed* f[3]
```

```
c[0,1] += g[0] signed* h[2];
```

```
c[2,3] += g[1] signed* h[3]
```

Cortex-A8 recognizes this pattern.

Reads c in N6 instead of N3.

Time	N1	N2	N3	N4	N5	N6	N7
1	<i>b</i>						
2		<i>a</i>					
3	<i>f</i>		×				
4		<i>e</i>		×			
5	<i>h</i>		×		×		
6		<i>g</i>		×		×	
7			×		×		
8				×		×	
9					×		+
10						×	
11							+
12							

ulate.

sequence of three insns:

```
= a[0] signed* b[0];
```

```
= a[1] signed* b[1]
```

```
+= e[2] signed* f[2];
```

```
+= e[3] signed* f[3]
```

```
+= g[0] signed* h[2];
```

```
+= g[1] signed* h[3]
```

A8 recognizes this pattern.

in N6 instead of N3.

Time	N1	N2	N3	N4	N5	N6	N7	N8
1	<i>b</i>							
2		<i>a</i>						
3	<i>f</i>		×					
4		<i>e</i>		×				
5	<i>h</i>		×		×			
6		<i>g</i>		×		×		
7			×		×			
8				×		×		<i>c</i>
9					×		+	
10						×		<i>c</i>
11							+	
12								<i>c</i>

NEON a

and perm

$r = s[1]$

Cortex-A

NEON is

that run

NEON a

Arithme

most im

can ofte

to hide l

Cortex-A

handling

10

of three insns:

```
signed* b[0];
```

```
signed* b[1]
```

```
signed* f[2];
```

```
signed* f[3]
```

```
signed* h[2];
```

```
signed* h[3]
```

izes this pattern.

stead of N3.

Time	N1	N2	N3	N4	N5	N6	N7	N8
1	<i>b</i>							
2		<i>a</i>						
3	<i>f</i>		×					
4		<i>e</i>		×				
5	<i>h</i>		×		×			
6		<i>g</i>		×		×		
7			×		×			
8				×		×		<i>c</i>
9					×		+	
10						×		<i>c</i>
11							+	
12								<i>c</i>

11

NEON also has load

and permutation i

$r = s[1] \ t[2] \ r[0]$

Cortex-A8 has a se

NEON load/store

that runs in paralle

NEON arithmetic

Arithmetic is typic

most important bo

can often schedule

to hide loads/store

Cortex-A7 is differ

handling all NEON

insns:

[0];

[1]

[2];

[3]

[2];

[3]

pattern.

3.

Time	N1	N2	N3	N4	N5	N6	N7	N8
1	<i>b</i>							
2		<i>a</i>						
3	<i>f</i>		×					
4		<i>e</i>		×				
5	<i>h</i>		×		×			
6		<i>g</i>		×		×		
7			×		×			
8				×		×		<i>c</i>
9					×		+	
10						×		<i>c</i>
11							+	
12								<i>c</i>

NEON also has load/store instructions and permutation instructions: e.g.,
 $r = s[1] \ t[2] \ r[2,3]$

Cortex-A8 has a separate NEON load/store unit that runs in parallel with NEON arithmetic unit.

Arithmetic is typically the most important bottleneck: can often schedule instructions to hide loads/stores/perms.

Cortex-A7 is different: one unit handling all NEON instructions.

Time	N1	N2	N3	N4	N5	N6	N7	N8
1	<i>b</i>							
2		<i>a</i>						
3	<i>f</i>		×					
4		<i>e</i>		×				
5	<i>h</i>		×		×			
6		<i>g</i>		×		×		
7			×		×			
8				×		×		<i>c</i>
9					×		+	
10						×		<i>c</i>
11							+	
12								<i>c</i>

NEON also has load/store insns and permutation insns: e.g.,
 $r = s[1] \ t[2] \ r[2,3]$

Cortex-A8 has a separate NEON load/store unit that runs in parallel with NEON arithmetic unit.

Arithmetic is typically most important bottleneck: can often schedule insns to hide loads/stores/perms.

Cortex-A7 is different: one unit handling all NEON insns.

	N2	N3	N4	N5	N6	N7	N8
<i>a</i>							
<i>e</i>	×		×				
<i>g</i>	×		×	×	×		
			×		×		<i>c</i>
				×		+	<i>c</i>
					×		<i>c</i>
						+	<i>c</i>

NEON also has load/store insns and permutation insns: e.g.,
 $r = s[1] \ t[2] \ r[2,3]$

Cortex-A8 has a separate NEON load/store unit that runs in parallel with NEON arithmetic unit.

Arithmetic is typically most important bottleneck: can often schedule insns to hide loads/stores/perms.

Cortex-A7 is different: one unit handling all NEON insns.

Curve25

Radix 2^2
 (f_0, f_1, f_2)

to repres

$f = f_0 +$

$2^{102} f_4 +$

$2^{204} f_8 +$

Unscaled

f is valu

$f_0 t^0 + 2$

$f_4 t^4 + 2$

$f_8 t^8 + 2$

N4	N5	N6	N7	N8
×				
	×			
×				
	×			c
×		+		
	×			c
		+		
				c

11

NEON also has load/store insns and permutation insns: e.g.,
 $r = s[1] \ t[2] \ r[2,3]$

Cortex-A8 has a separate NEON load/store unit that runs in parallel with NEON arithmetic unit.

Arithmetic is typically most important bottleneck: can often schedule insns to hide loads/stores/perms.

Cortex-A7 is different: one unit handling all NEON insns.

12

Curve25519 on NEON

Radix $2^{25.5}$: Use s ($f_0, f_1, f_2, f_3, f_4, f_5, \dots$) to represent the int
 $f = f_0 + 2^{26} f_1 + 2^{52} f_2 + 2^{78} f_3 + 2^{104} f_4 + 2^{130} f_5 + 2^{156} f_6 + 2^{182} f_7 + 2^{208} f_8 + 2^{234} f_9 \dots$

Unscaled polynomial
 f is value at $2^{25.5}$
 $f_0 t^0 + 2^{0.5} f_1 t^1 + \dots + f_4 t^4 + 2^{0.5} f_5 t^5 + \dots + f_8 t^8 + 2^{0.5} f_9 t^9$.

NEON also has load/store insns and permutation insns: e.g.,
 $r = s[1] \ t[2] \ r[2,3]$

Cortex-A8 has a separate NEON load/store unit that runs in parallel with NEON arithmetic unit.

Arithmetic is typically most important bottleneck: can often schedule insns to hide loads/stores/perms.

Cortex-A7 is different: one unit handling all NEON insns.

Curve25519 on NEON

Radix $2^{25.5}$: Use small integers $(f_0, f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8, f_9)$ to represent the integer
 $f = f_0 + 2^{26} f_1 + 2^{51} f_2 + 2^{76} f_3 + 2^{102} f_4 + 2^{128} f_5 + 2^{153} f_6 + 2^{179} f_7 + 2^{204} f_8 + 2^{230} f_9$ modulo 2^{255}

Unscaled polynomial view:
 f is value at $2^{25.5}$ of the polynomial
 $f_0 t^0 + 2^{0.5} f_1 t^1 + f_2 t^2 + 2^{0.5} f_3 t^3 + f_4 t^4 + 2^{0.5} f_5 t^5 + f_6 t^6 + 2^{0.5} f_7 t^7 + f_8 t^8 + 2^{0.5} f_9 t^9$.

NEON also has load/store insns
and permutation insns: e.g.,
 $r = s[1] \ t[2] \ r[2,3]$

Cortex-A8 has a separate
NEON load/store unit
that runs in parallel with
NEON arithmetic unit.

Arithmetic is typically
most important bottleneck:
can often schedule insns
to hide loads/stores/perms.

Cortex-A7 is different: one unit
handling all NEON insns.

Curve25519 on NEON

Radix $2^{25.5}$: Use small integers
($f_0, f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8, f_9$)

to represent the integer

$$f = f_0 + 2^{26} f_1 + 2^{51} f_2 + 2^{77} f_3 + 2^{102} f_4 + 2^{128} f_5 + 2^{153} f_6 + 2^{179} f_7 + 2^{204} f_8 + 2^{230} f_9 \text{ modulo } 2^{255} - 19.$$

Unscaled polynomial view:

$$f \text{ is value at } 2^{25.5} \text{ of the poly } f_0 t^0 + 2^{0.5} f_1 t^1 + f_2 t^2 + 2^{0.5} f_3 t^3 + f_4 t^4 + 2^{0.5} f_5 t^5 + f_6 t^6 + 2^{0.5} f_7 t^7 + f_8 t^8 + 2^{0.5} f_9 t^9.$$

also has load/store insns
 permutation insns: e.g.,
 t[2] r[2,3]
 A8 has a separate
 load/store unit
 s in parallel with
 arithmetic unit.
 ic is typically
 important bottleneck:
 n schedule insns
 oads/stores/perms.
 A7 is different: one unit
 ; all NEON insns.

Curve25519 on NEON

Radix $2^{25.5}$: Use small integers
 $(f_0, f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8, f_9)$
 to represent the integer
 $f = f_0 + 2^{26} f_1 + 2^{51} f_2 + 2^{77} f_3 +$
 $2^{102} f_4 + 2^{128} f_5 + 2^{153} f_6 + 2^{179} f_7 +$
 $2^{204} f_8 + 2^{230} f_9$ modulo $2^{255} - 19$.

Unscaled polynomial view:

f is value at $2^{25.5}$ of the poly
 $f_0 t^0 + 2^{0.5} f_1 t^1 + f_2 t^2 + 2^{0.5} f_3 t^3 +$
 $f_4 t^4 + 2^{0.5} f_5 t^5 + f_6 t^6 + 2^{0.5} f_7 t^7 +$
 $f_8 t^8 + 2^{0.5} f_9 t^9$.

$$h \equiv fg$$

$$h_0 = f_0 g_0$$

$$h_1 = f_0 g_1$$

$$h_2 = f_0 g_2$$

$$h_3 = f_0 g_3$$

$$h_4 = f_0 g_4$$

$$h_5 = f_0 g_5$$

$$h_6 = f_0 g_6$$

$$h_7 = f_0 g_7$$

$$h_8 = f_0 g_8$$

$$h_9 = f_0 g_9$$

Proof: r

Curve25519 on NEON

Radix $2^{25.5}$: Use small integers
 $(f_0, f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8, f_9)$

to represent the integer

$$f = f_0 + 2^{26} f_1 + 2^{51} f_2 + 2^{77} f_3 + 2^{102} f_4 + 2^{128} f_5 + 2^{153} f_6 + 2^{179} f_7 + 2^{204} f_8 + 2^{230} f_9 \text{ modulo } 2^{255} - 19.$$

Unscaled polynomial view:

f is value at $2^{25.5}$ of the poly

$$f_0 t^0 + 2^{0.5} f_1 t^1 + f_2 t^2 + 2^{0.5} f_3 t^3 + f_4 t^4 + 2^{0.5} f_5 t^5 + f_6 t^6 + 2^{0.5} f_7 t^7 + f_8 t^8 + 2^{0.5} f_9 t^9.$$

$$h \equiv fg \pmod{2^{255}}$$

$$h_0 = f_0 g_0 + 38 f_1 g_9 + 19 f_2 g_8 + \dots$$

$$h_1 = f_0 g_1 + f_1 g_0 + 19 f_2 g_9 + \dots$$

$$h_2 = f_0 g_2 + 2 f_1 g_1 + 19 f_2 g_8 + \dots$$

$$h_3 = f_0 g_3 + f_1 g_2 + 19 f_2 g_7 + \dots$$

$$h_4 = f_0 g_4 + 2 f_1 g_3 + 19 f_2 g_6 + \dots$$

$$h_5 = f_0 g_5 + f_1 g_4 + 19 f_2 g_5 + \dots$$

$$h_6 = f_0 g_6 + 2 f_1 g_5 + 19 f_2 g_4 + \dots$$

$$h_7 = f_0 g_7 + f_1 g_6 + 19 f_2 g_3 + \dots$$

$$h_8 = f_0 g_8 + 2 f_1 g_7 + 19 f_2 g_2 + \dots$$

$$h_9 = f_0 g_9 + f_1 g_8 + 19 f_2 g_1 + \dots$$

Proof: multiply po

Curve25519 on NEON

Radix $2^{25.5}$: Use small integers
 $(f_0, f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8, f_9)$

to represent the integer

$$f = f_0 + 2^{26} f_1 + 2^{51} f_2 + 2^{77} f_3 + 2^{102} f_4 + 2^{128} f_5 + 2^{153} f_6 + 2^{179} f_7 + 2^{204} f_8 + 2^{230} f_9 \text{ modulo } 2^{255} - 19.$$

Unscaled polynomial view:

f is value at $2^{25.5}$ of the poly

$$f_0 t^0 + 2^{0.5} f_1 t^1 + f_2 t^2 + 2^{0.5} f_3 t^3 + f_4 t^4 + 2^{0.5} f_5 t^5 + f_6 t^6 + 2^{0.5} f_7 t^7 + f_8 t^8 + 2^{0.5} f_9 t^9.$$

$$h \equiv fg \pmod{2^{255} - 19} \text{ w}$$

$$h_0 = f_0 g_0 + 38 f_1 g_9 + 19 f_2 g_8 + 38 f_3 g_7$$

$$h_1 = f_0 g_1 + f_1 g_0 + 19 f_2 g_9 + 19 f_3 g_8$$

$$h_2 = f_0 g_2 + 2 f_1 g_1 + f_2 g_0 + 38 f_3 g_9$$

$$h_3 = f_0 g_3 + f_1 g_2 + f_2 g_1 + f_3 g_0$$

$$h_4 = f_0 g_4 + 2 f_1 g_3 + f_2 g_2 + 2 f_3 g_1$$

$$h_5 = f_0 g_5 + f_1 g_4 + f_2 g_3 + f_3 g_2$$

$$h_6 = f_0 g_6 + 2 f_1 g_5 + f_2 g_4 + 2 f_3 g_3$$

$$h_7 = f_0 g_7 + f_1 g_6 + f_2 g_5 + f_3 g_4$$

$$h_8 = f_0 g_8 + 2 f_1 g_7 + f_2 g_6 + 2 f_3 g_5$$

$$h_9 = f_0 g_9 + f_1 g_8 + f_2 g_7 + f_3 g_6$$

Proof: multiply polys mod t

Curve25519 on NEON

Radix $2^{25.5}$: Use small integers

$(f_0, f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8, f_9)$

to represent the integer

$$f = f_0 + 2^{26}f_1 + 2^{51}f_2 + 2^{77}f_3 + 2^{102}f_4 + 2^{128}f_5 + 2^{153}f_6 + 2^{179}f_7 + 2^{204}f_8 + 2^{230}f_9 \text{ modulo } 2^{255} - 19.$$

Unscaled polynomial view:

f is value at $2^{25.5}$ of the poly

$$f_0 t^0 + 2^{0.5} f_1 t^1 + f_2 t^2 + 2^{0.5} f_3 t^3 + f_4 t^4 + 2^{0.5} f_5 t^5 + f_6 t^6 + 2^{0.5} f_7 t^7 + f_8 t^8 + 2^{0.5} f_9 t^9.$$

$h \equiv fg \pmod{2^{255} - 19}$ where

$$h_0 = f_0 g_0 + 38f_1 g_9 + 19f_2 g_8 + 38f_3 g_7 + 19f_4 g_6 +$$

$$h_1 = f_0 g_1 + f_1 g_0 + 19f_2 g_9 + 19f_3 g_8 + 19f_4 g_7 +$$

$$h_2 = f_0 g_2 + 2f_1 g_1 + f_2 g_0 + 38f_3 g_9 + 19f_4 g_8 +$$

$$h_3 = f_0 g_3 + f_1 g_2 + f_2 g_1 + f_3 g_0 + 19f_4 g_9 +$$

$$h_4 = f_0 g_4 + 2f_1 g_3 + f_2 g_2 + 2f_3 g_1 + f_4 g_0 +$$

$$h_5 = f_0 g_5 + f_1 g_4 + f_2 g_3 + f_3 g_2 + f_4 g_1 +$$

$$h_6 = f_0 g_6 + 2f_1 g_5 + f_2 g_4 + 2f_3 g_3 + f_4 g_2 +$$

$$h_7 = f_0 g_7 + f_1 g_6 + f_2 g_5 + f_3 g_4 + f_4 g_3 +$$

$$h_8 = f_0 g_8 + 2f_1 g_7 + f_2 g_6 + 2f_3 g_5 + f_4 g_4 +$$

$$h_9 = f_0 g_9 + f_1 g_8 + f_2 g_7 + f_3 g_6 + f_4 g_5 +$$

Proof: multiply polys mod $t^{10} - 19$.

25.5: Use small integers

$(f_0, f_1, f_2, f_3, f_4, f_5, f_6, f_7, f_8, f_9)$

represent the integer

$$-2^{26}f_1 + 2^{51}f_2 + 2^{77}f_3 + 2^{128}f_5 + 2^{153}f_6 + 2^{179}f_7 + 2^{230}f_9 \text{ modulo } 2^{255} - 19.$$

and polynomial view:

see at $2^{25.5}$ of the poly

$$2^{0.5}f_1t^1 + f_2t^2 + 2^{0.5}f_3t^3 + 2^{0.5}f_5t^5 + f_6t^6 + 2^{0.5}f_7t^7 + 2^{0.5}f_9t^9.$$

$h \equiv fg \pmod{2^{255} - 19}$ where

$$h_0 = f_0g_0 + 38f_1g_9 + 19f_2g_8 + 38f_3g_7 + 19f_4g_6 + 38f_5g_5 + 19f_6g_4 + 38f_7g_3 + 19f_8g_2 + 38f_9g_1$$

$$h_1 = f_0g_1 + f_1g_0 + 19f_2g_9 + 19f_3g_8 + 19f_4g_7 + 19f_5g_6 + 19f_6g_5 + 19f_7g_4 + 19f_8g_3 + 19f_9g_2$$

$$h_2 = f_0g_2 + 2f_1g_1 + f_2g_0 + 38f_3g_9 + 19f_4g_8 + 38f_5g_7 + 19f_6g_6 + 19f_7g_5 + 19f_8g_4 + 19f_9g_3$$

$$h_3 = f_0g_3 + f_1g_2 + f_2g_1 + f_3g_0 + 19f_4g_9 + 19f_5g_8 + 19f_6g_7 + 19f_7g_6 + 19f_8g_5 + 19f_9g_4$$

$$h_4 = f_0g_4 + 2f_1g_3 + f_2g_2 + 2f_3g_1 + f_4g_0 + 38f_5g_9 + 19f_6g_8 + 19f_7g_7 + 19f_8g_6 + 19f_9g_5$$

$$h_5 = f_0g_5 + f_1g_4 + f_2g_3 + f_3g_2 + f_4g_1 + f_5g_0 + 19f_6g_9 + 19f_7g_8 + 19f_8g_7 + 19f_9g_6$$

$$h_6 = f_0g_6 + 2f_1g_5 + f_2g_4 + 2f_3g_3 + f_4g_2 + 2f_5g_1 + f_6g_0 + 19f_7g_9 + 19f_8g_8 + 19f_9g_7$$

$$h_7 = f_0g_7 + f_1g_6 + f_2g_5 + f_3g_4 + f_4g_3 + f_5g_2 + f_6g_1 + 19f_7g_9 + 19f_8g_8 + 19f_9g_7$$

$$h_8 = f_0g_8 + 2f_1g_7 + f_2g_6 + 2f_3g_5 + f_4g_4 + 2f_5g_3 + f_6g_2 + 19f_7g_9 + 19f_8g_8 + 19f_9g_7$$

$$h_9 = f_0g_9 + f_1g_8 + f_2g_7 + f_3g_6 + f_4g_5 + f_5g_4 + f_6g_3 + 19f_7g_9 + 19f_8g_8 + 19f_9g_7$$

Proof: multiply polys mod $t^{10} - 19$.

small integers

f_6, f_7, f_8, f_9)

integer

$2^{51}f_2 + 2^{77}f_3 +$

$2^{153}f_6 + 2^{179}f_7 +$

modulo $2^{255} - 19$.

ial view:

of the poly

$f_2t^2 + 2^{0.5}f_3t^3 +$

$f_6t^6 + 2^{0.5}f_7t^7 +$

$h \equiv fg \pmod{2^{255} - 19}$ where

$$h_0 = f_0g_0 + 38f_1g_9 + 19f_2g_8 + 38f_3g_7 + 19f_4g_6 + 38f_5g_5 + 19f_6g_4 + 38f_7g_3$$

$$h_1 = f_0g_1 + f_1g_0 + 19f_2g_9 + 19f_3g_8 + 19f_4g_7 + 19f_5g_6 + 19f_6g_5 + 19f_7g_4$$

$$h_2 = f_0g_2 + 2f_1g_1 + f_2g_0 + 38f_3g_9 + 19f_4g_8 + 38f_5g_7 + 19f_6g_6 + 38f_7g_5$$

$$h_3 = f_0g_3 + f_1g_2 + f_2g_1 + f_3g_0 + 19f_4g_9 + 19f_5g_8 + 19f_6g_7 + 19f_7g_6$$

$$h_4 = f_0g_4 + 2f_1g_3 + f_2g_2 + 2f_3g_1 + f_4g_0 + 38f_5g_9 + 19f_6g_8 + 38f_7g_7$$

$$h_5 = f_0g_5 + f_1g_4 + f_2g_3 + f_3g_2 + f_4g_1 + f_5g_0 + 19f_6g_9 + 19f_7g_8$$

$$h_6 = f_0g_6 + 2f_1g_5 + f_2g_4 + 2f_3g_3 + f_4g_2 + 2f_5g_1 + f_6g_0 + 38f_7g_9$$

$$h_7 = f_0g_7 + f_1g_6 + f_2g_5 + f_3g_4 + f_4g_3 + f_5g_2 + f_6g_1 + f_7g_0$$

$$h_8 = f_0g_8 + 2f_1g_7 + f_2g_6 + 2f_3g_5 + f_4g_4 + 2f_5g_3 + f_6g_2 + 2f_7g_1$$

$$h_9 = f_0g_9 + f_1g_8 + f_2g_7 + f_3g_6 + f_4g_5 + f_5g_4 + f_6g_3 + f_7g_2$$

Proof: multiply polys mod $t^{10} - 19$.

$h \equiv fg \pmod{2^{255} - 19}$ where

$$h_0 = f_0g_0 + 38f_1g_9 + 19f_2g_8 + 38f_3g_7 + 19f_4g_6 + 38f_5g_5 + 19f_6g_4 + 38f_7g_3 + 19f_8g_2 + 38f_9g_1$$

$$h_1 = f_0g_1 + f_1g_0 + 19f_2g_9 + 19f_3g_8 + 19f_4g_7 + 19f_5g_6 + 19f_6g_5 + 19f_7g_4 + 19f_8g_3 + 19f_9g_2$$

$$h_2 = f_0g_2 + 2f_1g_1 + f_2g_0 + 38f_3g_9 + 19f_4g_8 + 38f_5g_7 + 19f_6g_6 + 38f_7g_5 + 19f_8g_4 + 38f_9g_3$$

$$h_3 = f_0g_3 + f_1g_2 + f_2g_1 + f_3g_0 + 19f_4g_9 + 19f_5g_8 + 19f_6g_7 + 19f_7g_6 + 19f_8g_5 + 19f_9g_4$$

$$h_4 = f_0g_4 + 2f_1g_3 + f_2g_2 + 2f_3g_1 + f_4g_0 + 38f_5g_9 + 19f_6g_8 + 38f_7g_7 + 19f_8g_6 + 38f_9g_5$$

$$h_5 = f_0g_5 + f_1g_4 + f_2g_3 + f_3g_2 + f_4g_1 + f_5g_0 + 19f_6g_9 + 19f_7g_8 + 19f_8g_7 + 19f_9g_6$$

$$h_6 = f_0g_6 + 2f_1g_5 + f_2g_4 + 2f_3g_3 + f_4g_2 + 2f_5g_1 + f_6g_0 + 38f_7g_9 + 19f_8g_8 + 38f_9g_7$$

$$h_7 = f_0g_7 + f_1g_6 + f_2g_5 + f_3g_4 + f_4g_3 + f_5g_2 + f_6g_1 + f_7g_0 + 19f_8g_9 + 19f_9g_8$$

$$h_8 = f_0g_8 + 2f_1g_7 + f_2g_6 + 2f_3g_5 + f_4g_4 + 2f_5g_3 + f_6g_2 + 2f_7g_1 + f_8g_0 + 38f_9g_9$$

$$h_9 = f_0g_9 + f_1g_8 + f_2g_7 + f_3g_6 + f_4g_5 + f_5g_4 + f_6g_3 + f_7g_2 + f_8g_1 + f_9g_0$$

Proof: multiply polys mod $t^{10} - 19$.

$h \equiv fg \pmod{2^{255} - 19}$ where

$$\begin{aligned}
 h_0 &= f_0g_0 + 38f_1g_9 + 19f_2g_8 + 38f_3g_7 + 19f_4g_6 + 38f_5g_5 + 19f_6g_4 + 38f_7g_3 + 19f_8g_2 + 38f_9g_1, \\
 h_1 &= f_0g_1 + f_1g_0 + 19f_2g_9 + 19f_3g_8 + 19f_4g_7 + 19f_5g_6 + 19f_6g_5 + 19f_7g_4 + 19f_8g_3 + 19f_9g_2, \\
 h_2 &= f_0g_2 + 2f_1g_1 + f_2g_0 + 38f_3g_9 + 19f_4g_8 + 38f_5g_7 + 19f_6g_6 + 38f_7g_5 + 19f_8g_4 + 38f_9g_3, \\
 h_3 &= f_0g_3 + f_1g_2 + f_2g_1 + f_3g_0 + 19f_4g_9 + 19f_5g_8 + 19f_6g_7 + 19f_7g_6 + 19f_8g_5 + 19f_9g_4, \\
 h_4 &= f_0g_4 + 2f_1g_3 + f_2g_2 + 2f_3g_1 + f_4g_0 + 38f_5g_9 + 19f_6g_8 + 38f_7g_7 + 19f_8g_6 + 38f_9g_5, \\
 h_5 &= f_0g_5 + f_1g_4 + f_2g_3 + f_3g_2 + f_4g_1 + f_5g_0 + 19f_6g_9 + 19f_7g_8 + 19f_8g_7 + 19f_9g_6, \\
 h_6 &= f_0g_6 + 2f_1g_5 + f_2g_4 + 2f_3g_3 + f_4g_2 + 2f_5g_1 + f_6g_0 + 38f_7g_9 + 19f_8g_8 + 38f_9g_7, \\
 h_7 &= f_0g_7 + f_1g_6 + f_2g_5 + f_3g_4 + f_4g_3 + f_5g_2 + f_6g_1 + f_7g_0 + 19f_8g_9 + 19f_9g_8, \\
 h_8 &= f_0g_8 + 2f_1g_7 + f_2g_6 + 2f_3g_5 + f_4g_4 + 2f_5g_3 + f_6g_2 + 2f_7g_1 + f_8g_0 + 38f_9g_9, \\
 h_9 &= f_0g_9 + f_1g_8 + f_2g_7 + f_3g_6 + f_4g_5 + f_5g_4 + f_6g_3 + f_7g_2 + f_8g_1 + f_9g_0.
 \end{aligned}$$

Proof: multiply polys mod $t^{10} - 19$.

(mod $2^{255} - 19$) where

$$\begin{aligned}
 &+ 38f_1g_9 + 19f_2g_8 + 38f_3g_7 + 19f_4g_6 + 38f_5g_5 + 19f_6g_4 + 38f_7g_3 + 19f_8g_2 + 38f_9g_1, \\
 &+ f_1g_0 + 19f_2g_9 + 19f_3g_8 + 19f_4g_7 + 19f_5g_6 + 19f_6g_5 + 19f_7g_4 + 19f_8g_3 + 19f_9g_2, \\
 &+ 2f_1g_1 + f_2g_0 + 38f_3g_9 + 19f_4g_8 + 38f_5g_7 + 19f_6g_6 + 38f_7g_5 + 19f_8g_4 + 38f_9g_3, \\
 &+ f_1g_2 + f_2g_1 + f_3g_0 + 19f_4g_9 + 19f_5g_8 + 19f_6g_7 + 19f_7g_6 + 19f_8g_5 + 19f_9g_4, \\
 &+ 2f_1g_3 + f_2g_2 + 2f_3g_1 + f_4g_0 + 38f_5g_9 + 19f_6g_8 + 38f_7g_7 + 19f_8g_6 + 38f_9g_5, \\
 &+ f_1g_4 + f_2g_3 + f_3g_2 + f_4g_1 + f_5g_0 + 19f_6g_9 + 19f_7g_8 + 19f_8g_7 + 19f_9g_6, \\
 &+ 2f_1g_5 + f_2g_4 + 2f_3g_3 + f_4g_2 + 2f_5g_1 + f_6g_0 + 38f_7g_9 + 19f_8g_8 + 38f_9g_7, \\
 &+ f_1g_6 + f_2g_5 + f_3g_4 + f_4g_3 + f_5g_2 + f_6g_1 + f_7g_0 + 19f_8g_9 + 19f_9g_8, \\
 &+ 2f_1g_7 + f_2g_6 + 2f_3g_5 + f_4g_4 + 2f_5g_3 + f_6g_2 + 2f_7g_1 + f_8g_0 + 38f_9g_9, \\
 &+ f_1g_8 + f_2g_7 + f_3g_6 + f_4g_5 + f_5g_4 + f_6g_3 + f_7g_2 + f_8g_1 + f_9g_0.
 \end{aligned}$$

multiply polys mod $t^{10} - 19$.

Each h_i is a product of $2f_1, 2f_2, \dots, 2f_9, 19g_1, 19g_2, \dots, 19g_9$. Each h_i is under reduced sizes of (Analyze bugs can See 2011 Barbosa several h_0, h_1, \dots for subse

$t^{10} - 19$) where

$$\begin{aligned}
 & 19f_2g_8 + 38f_3g_7 + 19f_4g_6 + 38f_5g_5 + 19f_6g_4 + 38f_7g_3 + 19f_8g_2 + 38f_9g_1, \\
 & 19f_2g_9 + 19f_3g_8 + 19f_4g_7 + 19f_5g_6 + 19f_6g_5 + 19f_7g_4 + 19f_8g_3 + 19f_9g_2, \\
 & f_2g_0 + 38f_3g_9 + 19f_4g_8 + 38f_5g_7 + 19f_6g_6 + 38f_7g_5 + 19f_8g_4 + 38f_9g_3, \\
 & f_2g_1 + f_3g_0 + 19f_4g_9 + 19f_5g_8 + 19f_6g_7 + 19f_7g_6 + 19f_8g_5 + 19f_9g_4, \\
 & f_2g_2 + 2f_3g_1 + f_4g_0 + 38f_5g_9 + 19f_6g_8 + 38f_7g_7 + 19f_8g_6 + 38f_9g_5, \\
 & f_2g_3 + f_3g_2 + f_4g_1 + f_5g_0 + 19f_6g_9 + 19f_7g_8 + 19f_8g_7 + 19f_9g_6, \\
 & f_2g_4 + 2f_3g_3 + f_4g_2 + 2f_5g_1 + f_6g_0 + 38f_7g_9 + 19f_8g_8 + 38f_9g_7, \\
 & f_2g_5 + f_3g_4 + f_4g_3 + f_5g_2 + f_6g_1 + f_7g_0 + 19f_8g_9 + 19f_9g_8, \\
 & f_2g_6 + 2f_3g_5 + f_4g_4 + 2f_5g_3 + f_6g_2 + 2f_7g_1 + f_8g_0 + 38f_9g_9, \\
 & f_2g_7 + f_3g_6 + f_4g_5 + f_5g_4 + f_6g_3 + f_7g_2 + f_8g_1 + f_9g_0.
 \end{aligned}$$

polys mod $t^{10} - 19$.

Each h_i is a sum of products after products of $2f_1, 2f_3, 2f_5, 2f_7, 19g_1, 19g_2, \dots, 19g_9$.

Each h_i fits into 6 under reasonable sizes of f_1, g_1, \dots, g_9 .

(Analyze this very bugs can slip past See 2011 Brumley

Barbosa–Vercaute several recent Ope

h_0, h_1, \dots are too for subsequent mu

where

$$\begin{aligned}
 h_7 &+ 19f_4g_6 + 38f_5g_5 + 19f_6g_4 + 38f_7g_3 + 19f_8g_2 + 38f_9g_1, \\
 h_8 &+ 19f_4g_7 + 19f_5g_6 + 19f_6g_5 + 19f_7g_4 + 19f_8g_3 + 19f_9g_2, \\
 h_9 &+ 19f_4g_8 + 38f_5g_7 + 19f_6g_6 + 38f_7g_5 + 19f_8g_4 + 38f_9g_3, \\
 h_{10} &+ 19f_4g_9 + 19f_5g_8 + 19f_6g_7 + 19f_7g_6 + 19f_8g_5 + 19f_9g_4, \\
 h_{11} &+ f_4g_0 + 38f_5g_9 + 19f_6g_8 + 38f_7g_7 + 19f_8g_6 + 38f_9g_5, \\
 h_{12} &+ f_4g_1 + f_5g_0 + 19f_6g_9 + 19f_7g_8 + 19f_8g_7 + 19f_9g_6, \\
 h_{13} &+ f_4g_2 + 2f_5g_1 + f_6g_0 + 38f_7g_9 + 19f_8g_8 + 38f_9g_7, \\
 h_{14} &+ f_4g_3 + f_5g_2 + f_6g_1 + f_7g_0 + 19f_8g_9 + 19f_9g_8, \\
 h_{15} &+ f_4g_4 + 2f_5g_3 + f_6g_2 + 2f_7g_1 + f_8g_0 + 38f_9g_9, \\
 h_{16} &+ f_4g_5 + f_5g_4 + f_6g_3 + f_7g_2 + f_8g_1 + f_9g_0.
 \end{aligned}$$

— 19.

Each h_i is a sum of ten products after precomputation of $2f_1, 2f_3, 2f_5, 2f_7, 2f_9, 19g_1, 19g_2, \dots, 19g_9$.

Each h_i fits into 64 bits under reasonable limits on sizes of $f_1, g_1, \dots, f_9, g_9$.

(Analyze this very carefully: bugs can slip past most tests. See 2011 Brumley–Page–Barbosa–Vercauteren and several recent OpenSSL bugs. h_0, h_1, \dots are too large for subsequent multiplication.)

$$\begin{aligned}
&38f_5g_5 + 19f_6g_4 + 38f_7g_3 + 19f_8g_2 + 38f_9g_1, \\
&19f_5g_6 + 19f_6g_5 + 19f_7g_4 + 19f_8g_3 + 19f_9g_2, \\
&38f_5g_7 + 19f_6g_6 + 38f_7g_5 + 19f_8g_4 + 38f_9g_3, \\
&19f_5g_8 + 19f_6g_7 + 19f_7g_6 + 19f_8g_5 + 19f_9g_4, \\
&38f_5g_9 + 19f_6g_8 + 38f_7g_7 + 19f_8g_6 + 38f_9g_5, \\
&f_5g_0 + 19f_6g_9 + 19f_7g_8 + 19f_8g_7 + 19f_9g_6, \\
&2f_5g_1 + f_6g_0 + 38f_7g_9 + 19f_8g_8 + 38f_9g_7, \\
&f_5g_2 + f_6g_1 + f_7g_0 + 19f_8g_9 + 19f_9g_8, \\
&2f_5g_3 + f_6g_2 + 2f_7g_1 + f_8g_0 + 38f_9g_9, \\
&f_5g_4 + f_6g_3 + f_7g_2 + f_8g_1 + f_9g_0.
\end{aligned}$$

Each h_i is a sum of ten products after precomputation of $2f_1, 2f_3, 2f_5, 2f_7, 2f_9, 19g_1, 19g_2, \dots, 19g_9$.

Each h_i fits into 64 bits under reasonable limits on sizes of $f_1, g_1, \dots, f_9, g_9$.

(Analyze this very carefully: bugs can slip past most tests! See 2011 Brumley–Page–Barbosa–Vercauteren and several recent OpenSSL bugs.)

h_0, h_1, \dots are too large for subsequent multiplication.

$$\begin{aligned}
& f_6 g_4 + 38 f_7 g_3 + 19 f_8 g_2 + 38 f_9 g_1, \\
& f_6 g_5 + 19 f_7 g_4 + 19 f_8 g_3 + 19 f_9 g_2, \\
& f_6 g_6 + 38 f_7 g_5 + 19 f_8 g_4 + 38 f_9 g_3, \\
& f_6 g_7 + 19 f_7 g_6 + 19 f_8 g_5 + 19 f_9 g_4, \\
& f_6 g_8 + 38 f_7 g_7 + 19 f_8 g_6 + 38 f_9 g_5, \\
& f_6 g_9 + 19 f_7 g_8 + 19 f_8 g_7 + 19 f_9 g_6, \\
& f_6 g_0 + 38 f_7 g_9 + 19 f_8 g_8 + 38 f_9 g_7, \\
& f_6 g_1 + f_7 g_0 + 19 f_8 g_9 + 19 f_9 g_8, \\
& f_6 g_2 + 2 f_7 g_1 + f_8 g_0 + 38 f_9 g_9, \\
& f_6 g_3 + f_7 g_2 + f_8 g_1 + f_9 g_0.
\end{aligned}$$

Each h_i is a sum of ten products after precomputation of $2f_1, 2f_3, 2f_5, 2f_7, 2f_9, 19g_1, 19g_2, \dots, 19g_9$.

Each h_i fits into 64 bits under reasonable limits on sizes of $f_1, g_1, \dots, f_9, g_9$.

(Analyze this very carefully: bugs can slip past most tests! See 2011 Brumley–Page–Barbosa–Vercauteren and several recent OpenSSL bugs.)

h_0, h_1, \dots are too large for subsequent multiplication.

Carry h_0 replace ($(h_0 \bmod$ This ma

Similarly Eventua

We actu Slightly (given d but more

Some th

- Mix si
- Interle

$$\begin{aligned}
&+ 19f_8g_2 + 38f_9g_1, \\
&+ 19f_8g_3 + 19f_9g_2, \\
&+ 19f_8g_4 + 38f_9g_3, \\
&+ 19f_8g_5 + 19f_9g_4, \\
&+ 19f_8g_6 + 38f_9g_5, \\
&+ 19f_8g_7 + 19f_9g_6, \\
&+ 19f_8g_8 + 38f_9g_7, \\
&+ 19f_8g_9 + 19f_9g_8, \\
&+ f_8g_0 + 38f_9g_9, \\
&+ f_8g_1 + f_9g_0.
\end{aligned}$$

Each h_i is a sum of ten products after precomputation of $2f_1, 2f_3, 2f_5, 2f_7, 2f_9, 19g_1, 19g_2, \dots, 19g_9$.

Each h_i fits into 64 bits under reasonable limits on sizes of $f_1, g_1, \dots, f_9, g_9$.

(Analyze this very carefully: bugs can slip past most tests! See 2011 Brumley–Page–Barbosa–Vercauteren and several recent OpenSSL bugs.)

h_0, h_1, \dots are too large for subsequent multiplication.

Carry $h_0 \rightarrow h_1$: i.e. replace (h_0, h_1) with $(h_0 \bmod 2^{26}, h_1 + \lfloor h_0 / 2^{26} \rfloor)$. This makes h_0 small.

Similarly for other carries. Eventually all h_i are small.

We actually use signed integers. Slightly more expensive (given details of implementation) but more room for error.

Some things we have to do:

- Mix signed, unsigned integers
- Interleave reduction and multiplication

Each h_i is a sum of ten products after precomputation of $2f_1, 2f_3, 2f_5, 2f_7, 2f_9, 19g_1, 19g_2, \dots, 19g_9$.

Each h_i fits into 64 bits under reasonable limits on sizes of $f_1, g_1, \dots, f_9, g_9$.

(Analyze this very carefully: bugs can slip past most tests! See 2011 Brumley–Page–Barbosa–Vercauteren and several recent OpenSSL bugs.)

h_0, h_1, \dots are too large for subsequent multiplication.

Carry $h_0 \rightarrow h_1$: i.e., replace (h_0, h_1) with $(h_0 \bmod 2^{26}, h_1 + \lfloor h_0/2^{26} \rfloor)$. This makes h_0 small.

Similarly for other h_i .

Eventually all h_i are small enough.

We actually use signed coefficients. Slightly more expensive carry propagation (given details of insn set) but more room for $ab + c^2$ etc.

Some things we haven't tried:

- Mix signed, unsigned carries
- Interleave reduction, carry

Each h_i is a sum of ten products after precomputation of $2f_1, 2f_3, 2f_5, 2f_7, 2f_9, 19g_1, 19g_2, \dots, 19g_9$.

Each h_i fits into 64 bits under reasonable limits on sizes of $f_1, g_1, \dots, f_9, g_9$.

(Analyze this very carefully: bugs can slip past most tests! See 2011 Brumley–Page–Barbosa–Vercauteren and several recent OpenSSL bugs.)

h_0, h_1, \dots are too large for subsequent multiplication.

Carry $h_0 \rightarrow h_1$: i.e., replace (h_0, h_1) with $(h_0 \bmod 2^{26}, h_1 + \lfloor h_0/2^{26} \rfloor)$. This makes h_0 small.

Similarly for other h_i .

Eventually all h_i are small enough.

We actually use signed coeffs.

Slightly more expensive carries (given details of insn set)

but more room for $ab + c^2$ etc.

Some things we haven't tried yet:

- Mix signed, unsigned carries.
- Interleave reduction, carrying.

is a sum of ten
 s after precomputation
 $f_3, 2f_5, 2f_7, 2f_9,$
 $g_2, \dots, 19g_9.$

fits into 64 bits
 reasonable limits on
 $f_1, g_1, \dots, f_9, g_9.$

e this very carefully:
 n slip past most tests!
 1 Brumley–Page–
 –Vercauteren and
 recent OpenSSL bugs.)
 . are too large
 equent multiplication.

Carry $h_0 \rightarrow h_1$: i.e.,
 replace (h_0, h_1) with
 $(h_0 \bmod 2^{26}, h_1 + \lfloor h_0/2^{26} \rfloor).$
 This makes h_0 small.

Similarly for other h_i .

Eventually all h_i are small enough.

We actually use signed coeffs.

Slightly more expensive carries

(given details of insn set)

but more room for $ab + c^2$ etc.

Some things we haven't tried yet:

- Mix signed, unsigned carries.
- Interleave reduction, carrying.

Minor ch
 Result o
 used unt
 Find an
 for the C
 while the
 Sometim
 higher-le
 Example
 $h_2 \rightarrow h_3$
 $h_7 \rightarrow h_8$
 have lon

of ten
 computation
 $2f_9$,
 g_9 .

4 bits
 limits on
 f_9, g_9 .

carefully:
 most tests!

—Page—
 ren and
 enSSL bugs.)

large
 multiplication.

Carry $h_0 \rightarrow h_1$: i.e.,
 replace (h_0, h_1) with
 $(h_0 \bmod 2^{26}, h_1 + \lfloor h_0/2^{26} \rfloor)$.

This makes h_0 small.

Similarly for other h_i .

Eventually all h_i are small enough.

We actually use signed coeffs.

Slightly more expensive carries
 (given details of insn set)
 but more room for $ab + c^2$ etc.

Some things we haven't tried yet:

- Mix signed, unsigned carries.
- Interleave reduction, carrying.

Minor challenge: p
 Result of each instr
 used until a few cy
 Find an independe
 for the CPU to sta
 while the first instr
 Sometimes helps t
 higher-level compu
 Example: carries h
 $h_2 \rightarrow h_3 \rightarrow h_4 \rightarrow$
 $h_7 \rightarrow h_8 \rightarrow h_9 \rightarrow$
 have long chain of

Carry $h_0 \rightarrow h_1$: i.e.,
 replace (h_0, h_1) with
 $(h_0 \bmod 2^{26}, h_1 + \lfloor h_0/2^{26} \rfloor)$.

This makes h_0 small.

Similarly for other h_i .

Eventually all h_i are small enough.

We actually use signed coeffs.

Slightly more expensive carries

(given details of insn set)

but more room for $ab + c^2$ etc.

Some things we haven't tried yet:

- Mix signed, unsigned carries.
- Interleave reduction, carrying.

Minor challenge: pipelining.

Result of each insn cannot be
 used until a few cycles later.

Find an independent insn
 for the CPU to start working
 while the first insn is in prog.

Sometimes helps to adjust
 higher-level computations.

Example: carries $h_0 \rightarrow h_1 \rightarrow$
 $h_2 \rightarrow h_3 \rightarrow h_4 \rightarrow h_5 \rightarrow h_6 \rightarrow$
 $h_7 \rightarrow h_8 \rightarrow h_9 \rightarrow h_0 \rightarrow h_1$
 have long chain of dependencies.

Carry $h_0 \rightarrow h_1$: i.e.,
 replace (h_0, h_1) with
 $(h_0 \bmod 2^{26}, h_1 + \lfloor h_0/2^{26} \rfloor)$.

This makes h_0 small.

Similarly for other h_i .

Eventually all h_i are small enough.

We actually use signed coeffs.

Slightly more expensive carries

(given details of insn set)

but more room for $ab + c^2$ etc.

Some things we haven't tried yet:

- Mix signed, unsigned carries.
- Interleave reduction, carrying.

Minor challenge: pipelining.

Result of each insn cannot be
 used until a few cycles later.

Find an independent insn
 for the CPU to start working on
 while the first insn is in progress.

Sometimes helps to adjust
 higher-level computations.

Example: carries $h_0 \rightarrow h_1 \rightarrow$
 $h_2 \rightarrow h_3 \rightarrow h_4 \rightarrow h_5 \rightarrow h_6 \rightarrow$

$h_7 \rightarrow h_8 \rightarrow h_9 \rightarrow h_0 \rightarrow h_1$

have long chain of dependencies.

$h_0 \rightarrow h_1$: i.e.,
 (h_0, h_1) with
 $2^{26}, h_1 + \lfloor h_0/2^{26} \rfloor$).
 makes h_0 small.
 for other h_i .
 ally all h_i are small enough.
 ally use signed coeffs.
 more expensive carries
 (details of insn set)
 e room for $ab + c^2$ etc.
 ings we haven't tried yet:
 gned, unsigned carries.
 ave reduction, carrying.

Minor challenge: pipelining.
 Result of each insn cannot be
 used until a few cycles later.
 Find an independent insn
 for the CPU to start working on
 while the first insn is in progress.
 Sometimes helps to adjust
 higher-level computations.
 Example: carries $h_0 \rightarrow h_1 \rightarrow$
 $h_2 \rightarrow h_3 \rightarrow h_4 \rightarrow h_5 \rightarrow h_6 \rightarrow$
 $h_7 \rightarrow h_8 \rightarrow h_9 \rightarrow h_0 \rightarrow h_1$
 have long chain of dependencies.

Alternat
 $h_0 \rightarrow h_1$
 $h_1 \rightarrow h_2$
 $h_2 \rightarrow h_3$
 $h_3 \rightarrow h_4$
 $h_4 \rightarrow h_5$
 $h_5 \rightarrow h_6$
 12 carries
 but later
 Now mu
 to find i
 for CPU

e.,
 th
 $\lfloor h_0/2^{26} \rfloor$).
 all.
 h_i .
 re small enough.
 gned coeffs.
 nsive carries
 (insn set)
 $ab + c^2$ etc.
 haven't tried yet:
 gned carries.
 tion, carrying.

Minor challenge: pipelining.
 Result of each insn cannot be
 used until a few cycles later.
 Find an independent insn
 for the CPU to start working on
 while the first insn is in progress.
 Sometimes helps to adjust
 higher-level computations.
 Example: carries $h_0 \rightarrow h_1 \rightarrow$
 $h_2 \rightarrow h_3 \rightarrow h_4 \rightarrow h_5 \rightarrow h_6 \rightarrow$
 $h_7 \rightarrow h_8 \rightarrow h_9 \rightarrow h_0 \rightarrow h_1$
 have long chain of dependencies.

Alternative: carry
 $h_0 \rightarrow h_1$ and $h_5 \rightarrow$
 $h_1 \rightarrow h_2$ and $h_6 \rightarrow$
 $h_2 \rightarrow h_3$ and $h_7 \rightarrow$
 $h_3 \rightarrow h_4$ and $h_8 \rightarrow$
 $h_4 \rightarrow h_5$ and $h_9 \rightarrow$
 $h_5 \rightarrow h_6$ and $h_0 \rightarrow$
 12 carries instead
 but latency is muc
 Now much easier
 to find independen
 for CPU to handle

Minor challenge: pipelining.

Result of each insn cannot be used until a few cycles later.

Find an independent insn for the CPU to start working on while the first insn is in progress.

Sometimes helps to adjust higher-level computations.

Example: carries $h_0 \rightarrow h_1 \rightarrow h_2 \rightarrow h_3 \rightarrow h_4 \rightarrow h_5 \rightarrow h_6 \rightarrow h_7 \rightarrow h_8 \rightarrow h_9 \rightarrow h_0 \rightarrow h_1$ have long chain of dependencies.

Alternative: carry

$h_0 \rightarrow h_1$ and $h_5 \rightarrow h_6$;

$h_1 \rightarrow h_2$ and $h_6 \rightarrow h_7$;

$h_2 \rightarrow h_3$ and $h_7 \rightarrow h_8$;

$h_3 \rightarrow h_4$ and $h_8 \rightarrow h_9$;

$h_4 \rightarrow h_5$ and $h_9 \rightarrow h_0$;

$h_5 \rightarrow h_6$ and $h_0 \rightarrow h_1$.

12 carries instead of 11, but latency is much smaller.

Now much easier to find independent insns for CPU to handle in parallel.

Minor challenge: pipelining.

Result of each insn cannot be used until a few cycles later.

Find an independent insn for the CPU to start working on while the first insn is in progress.

Sometimes helps to adjust higher-level computations.

Example: carries $h_0 \rightarrow h_1 \rightarrow h_2 \rightarrow h_3 \rightarrow h_4 \rightarrow h_5 \rightarrow h_6 \rightarrow h_7 \rightarrow h_8 \rightarrow h_9 \rightarrow h_0 \rightarrow h_1$ have long chain of dependencies.

Alternative: carry

$h_0 \rightarrow h_1$ and $h_5 \rightarrow h_6$;

$h_1 \rightarrow h_2$ and $h_6 \rightarrow h_7$;

$h_2 \rightarrow h_3$ and $h_7 \rightarrow h_8$;

$h_3 \rightarrow h_4$ and $h_8 \rightarrow h_9$;

$h_4 \rightarrow h_5$ and $h_9 \rightarrow h_0$;

$h_5 \rightarrow h_6$ and $h_0 \rightarrow h_1$.

12 carries instead of 11, but latency is much smaller.

Now much easier to find independent insns for CPU to handle in parallel.

challenge: pipelining.

if each insn cannot be
until a few cycles later.

independent insn

CPU to start working on
the first insn is in progress.

nes helps to adjust
level computations.

e: carries $h_0 \rightarrow h_1 \rightarrow$
 $\rightarrow h_4 \rightarrow h_5 \rightarrow h_6 \rightarrow$
 $\rightarrow h_9 \rightarrow h_0 \rightarrow h_1$

g chain of dependencies.

Alternative: carry

$h_0 \rightarrow h_1$ and $h_5 \rightarrow h_6$;

$h_1 \rightarrow h_2$ and $h_6 \rightarrow h_7$;

$h_2 \rightarrow h_3$ and $h_7 \rightarrow h_8$;

$h_3 \rightarrow h_4$ and $h_8 \rightarrow h_9$;

$h_4 \rightarrow h_5$ and $h_9 \rightarrow h_0$;

$h_5 \rightarrow h_6$ and $h_0 \rightarrow h_1$.

12 carries instead of 11,
but latency is much smaller.

Now much easier

to find independent insns
for CPU to handle in parallel.

Major ch

e.g. 4x a
does 4 a
but need
of inputs

On Cort

occasionally

run in pa

but frequ

would be

On Cort

every op

pipelining.
 cannot be
 cycles later.
 ent insn
 art working on
 is in progress.
 to adjust
 utations.
 $h_0 \rightarrow h_1 \rightarrow$
 $h_5 \rightarrow h_6 \rightarrow$
 $h_0 \rightarrow h_1$
 F dependencies.

Alternative: carry

$h_0 \rightarrow h_1$ and $h_5 \rightarrow h_6$;
 $h_1 \rightarrow h_2$ and $h_6 \rightarrow h_7$;
 $h_2 \rightarrow h_3$ and $h_7 \rightarrow h_8$;
 $h_3 \rightarrow h_4$ and $h_8 \rightarrow h_9$;
 $h_4 \rightarrow h_5$ and $h_9 \rightarrow h_0$;
 $h_5 \rightarrow h_6$ and $h_0 \rightarrow h_1$.

12 carries instead of 11,
 but latency is much smaller.

Now much easier
 to find independent insns
 for CPU to handle in parallel.

Major challenge: v
 e.g. 4x $a = b + c$
 does 4 additions a
 but needs particula
 of inputs and outp
 On Cortex-A8,
occasional permut
 run in parallel with
 but frequent perm
 would be a bottler
 On Cortex-A7,
 every operation co

Alternative: carry

$h_0 \rightarrow h_1$ and $h_5 \rightarrow h_6$;

$h_1 \rightarrow h_2$ and $h_6 \rightarrow h_7$;

$h_2 \rightarrow h_3$ and $h_7 \rightarrow h_8$;

$h_3 \rightarrow h_4$ and $h_8 \rightarrow h_9$;

$h_4 \rightarrow h_5$ and $h_9 \rightarrow h_0$;

$h_5 \rightarrow h_6$ and $h_0 \rightarrow h_1$.

12 carries instead of 11,
but latency is much smaller.

Now much easier
to find independent insns
for CPU to handle in parallel.

Major challenge: vectorization

e.g. $4x\ a = b + c$

does 4 additions at once,
but needs particular arrangement
of inputs and outputs.

On Cortex-A8,
occasional permutations
run in parallel with arithmetic
but frequent permutations
would be a bottleneck.

On Cortex-A7,
every operation costs cycles.

Alternative: carry

$h_0 \rightarrow h_1$ and $h_5 \rightarrow h_6$;

$h_1 \rightarrow h_2$ and $h_6 \rightarrow h_7$;

$h_2 \rightarrow h_3$ and $h_7 \rightarrow h_8$;

$h_3 \rightarrow h_4$ and $h_8 \rightarrow h_9$;

$h_4 \rightarrow h_5$ and $h_9 \rightarrow h_0$;

$h_5 \rightarrow h_6$ and $h_0 \rightarrow h_1$.

12 carries instead of 11,
but latency is much smaller.

Now much easier
to find independent insns
for CPU to handle in parallel.

Major challenge: vectorization.

e.g. $4x\ a = b + c$

does 4 additions at once,
but needs particular arrangement
of inputs and outputs.

On Cortex-A8,
occasional permutations
run in parallel with arithmetic,
but frequent permutations
would be a bottleneck.

On Cortex-A7,
every operation costs cycles.

ive: carry

and $h_5 \rightarrow h_6$;

and $h_6 \rightarrow h_7$;

and $h_7 \rightarrow h_8$;

and $h_8 \rightarrow h_9$;

and $h_9 \rightarrow h_0$;

and $h_0 \rightarrow h_1$.

es instead of 11,
ncy is much smaller.

ch easier

ndependent insns

to handle in parallel.

Major challenge: vectorization.

e.g. $4x \ a = b + c$

does 4 additions at once,
but needs particular arrangement
of inputs and outputs.

On Cortex-A8,

occasional permutations

run in parallel with arithmetic,

but frequent permutations

would be a bottleneck.

On Cortex-A7,

every operation costs cycles.

Often hi

do a pai

$h = fg$;

Vectoriz

Merge f_0

and f'_0, f

into vect

Similarly

Then co

Comput

into NEO

$c[0, 1]$

$c[2, 3]$

Major challenge: vectorization.

e.g. $4x \ a = b + c$

does 4 additions at once,

but needs particular arrangement
of inputs and outputs.

On Cortex-A8,

occasional permutations

run in parallel with arithmetic,

but frequent permutations

would be a bottleneck.

On Cortex-A7,

every operation costs cycles.

Often higher-level

do a pair of mults

$h = fg; h' = f'g'$

Vectorize across the

Merge f_0, f_1, \dots, f_9

and f'_0, f'_1, \dots, f'_9

into vectors (f_i, f'_i)

Similarly (g_i, g'_i) .

Then compute $(h_i$

Computation fits neatly

into NEON insns:

$c[0,1] = a[0]$ si

$c[2,3] = a[1]$ si

Major challenge: vectorization.

e.g. $4x \ a = b + c$

does 4 additions at once,

but needs particular arrangement of inputs and outputs.

On Cortex-A8,

occasional permutations

run in parallel with arithmetic,

but frequent permutations

would be a bottleneck.

On Cortex-A7,

every operation costs cycles.

Often higher-level operations

do a pair of mults in parallel

$$h = fg; h' = f'g'.$$

Vectorize across those mults

Merge f_0, f_1, \dots, f_9

and f'_0, f'_1, \dots, f'_9

into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally

into NEON insns: e.g.,

$c[0,1] = a[0] \text{ signed} * b[0,1]$

$c[2,3] = a[1] \text{ signed} * b[2,3]$

Major challenge: vectorization.

e.g. $4x \ a = b + c$

does 4 additions at once,

but needs particular arrangement of inputs and outputs.

On Cortex-A8,

occasional permutations

run in parallel with arithmetic,

but frequent permutations

would be a bottleneck.

On Cortex-A7,

every operation costs cycles.

Often higher-level operations do a pair of mults in parallel:

$$h = fg; \ h' = f'g'.$$

Vectorize across those mults.

Merge f_0, f_1, \dots, f_9

and f'_0, f'_1, \dots, f'_9

into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally

into NEON insns: e.g.,

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

challenge: vectorization.

$$a = b + c$$

additions at once,

needs particular arrangement

of inputs and outputs.

ex-A8,

data permutations

parallel with arithmetic,

data permutations

are a bottleneck.

ex-A7,

operation costs cycles.

Often higher-level operations
do a pair of mults in parallel:
 $h = fg; h' = f'g'$.

Vectorize across those mults.

Merge f_0, f_1, \dots, f_9

and f'_0, f'_1, \dots, f'_9

into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally

into NEON insns: e.g.,

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

Example

$$C = X_1$$

inside po

for Edwa

vectorization.

at once,

an arrangement

outputs.

operations

in arithmetic,

computations

check.

costs cycles.

Often higher-level operations
do a pair of mults in parallel:
 $h = fg; h' = f'g'$.

Vectorize across those mults.

Merge f_0, f_1, \dots, f_9

and f'_0, f'_1, \dots, f'_9

into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally

into NEON insns: e.g.,

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

Example: Recall

$C = X_1 \cdot X_2; D =$

inside point-addition

for Edwards curves

on.

Often higher-level operations
do a pair of mults in parallel:
 $h = fg; h' = f'g'$.

ment

Vectorize across those mults.

Merge f_0, f_1, \dots, f_9

and f'_0, f'_1, \dots, f'_9

into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally

into NEON insns: e.g.,

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formula

for Edwards curves.

ic,

Often higher-level operations
do a pair of mults in parallel:
 $h = fg; h' = f'g'$.

Vectorize across those mults.

Merge f_0, f_1, \dots, f_9

and f'_0, f'_1, \dots, f'_9

into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally

into NEON insns: e.g.,

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Often higher-level operations
do a pair of mults in parallel:
 $h = fg; h' = f'g'$.

Vectorize across those mults.

Merge f_0, f_1, \dots, f_9
and f'_0, f'_1, \dots, f'_9
into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally
into NEON insns: e.g.,

```
c[0,1] = a[0] signed* b[0];
```

```
c[2,3] = a[1] signed* b[1]
```

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Example: Can compute

$2P, 3P, 4P, 5P, 6P, 7P$ as

$$2P = P + P;$$

$$3P = 2P + P \text{ and } 4P = 2P + 2P;$$

$$5P = 4P + P \text{ and } 6P = 3P + 3P$$

$$\text{and } 7P = 4P + 3P.$$

Often higher-level operations
do a pair of mults in parallel:
 $h = fg; h' = f'g'$.

Vectorize across those mults.

Merge f_0, f_1, \dots, f_9
and f'_0, f'_1, \dots, f'_9
into vectors (f_i, f'_i) .

Similarly (g_i, g'_i) .

Then compute (h_i, h'_i) .

Computation fits naturally
into NEON insns: e.g.,

`c[0,1] = a[0] signed* b[0];`

`c[2,3] = a[1] signed* b[1]`

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Example: Can compute

$2P, 3P, 4P, 5P, 6P, 7P$ as

$$2P = P + P;$$

$$3P = 2P + P \text{ and } 4P = 2P + 2P;$$

$$5P = 4P + P \text{ and } 6P = 3P + 3P$$

$$\text{and } 7P = 4P + 3P.$$

Example: Typical algorithms

for fixed-base scalarmult

have many parallel point adds.

gher-level operations
r of mults in parallel:

$$h' = f'g'.$$

e across those mults.

$$f_0, f_1, \dots, f_9$$

$$f'_1, \dots, f'_9$$

tors (f_i, f'_i) .

(g_i, g'_i) .

mpute (h_i, h'_i) .

ation fits naturally

ON insns: e.g.,

`= a[0] signed* b[0];`

`= a[1] signed* b[1]`

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Example: Can compute

$2P, 3P, 4P, 5P, 6P, 7P$ as

$$2P = P + P;$$

$$3P = 2P + P \text{ and } 4P = 2P + 2P;$$

$$5P = 4P + P \text{ and } 6P = 3P + 3P$$

$$\text{and } 7P = 4P + 3P.$$

Example: Typical algorithms

for fixed-base scalarmult

have many parallel point adds.

Example

with a b

can vect

operations
in parallel:

those mults.

g

).

, h'_i).

naturally

e.g.,

igned* b[0];

igned* b[1]

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Example: Can compute

$2P, 3P, 4P, 5P, 6P, 7P$ as

$$2P = P + P;$$

$$3P = 2P + P \text{ and } 4P = 2P + 2P;$$

$$5P = 4P + P \text{ and } 6P = 3P + 3P$$

$$\text{and } 7P = 4P + 3P.$$

Example: Typical algorithms

for fixed-base scalarmult

have many parallel point adds.

Example: A busy s
with a backlog of
can vectorize across

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Example: Can compute

$2P, 3P, 4P, 5P, 6P, 7P$ as

$$2P = P + P;$$

$$3P = 2P + P \text{ and } 4P = 2P + 2P;$$

$$5P = 4P + P \text{ and } 6P = 3P + 3P$$

$$\text{and } 7P = 4P + 3P.$$

Example: Typical algorithms

for fixed-base scalarmult

have many parallel point adds.

Example: A busy server
with a backlog of scalarmult
can vectorize across them.

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Example: Can compute

$2P, 3P, 4P, 5P, 6P, 7P$ as

$$2P = P + P;$$

$$3P = 2P + P \text{ and } 4P = 2P + 2P;$$

$$5P = 4P + P \text{ and } 6P = 3P + 3P$$

$$\text{and } 7P = 4P + 3P.$$

Example: Typical algorithms

for fixed-base scalarmult

have many parallel point adds.

Example: A busy server
with a backlog of scalarmults
can vectorize across them.

Example: Recall

$$C = X_1 \cdot X_2; D = Y_1 \cdot Y_2$$

inside point-addition formulas
for Edwards curves.

Example: Can compute

$2P, 3P, 4P, 5P, 6P, 7P$ as

$$2P = P + P;$$

$$3P = 2P + P \text{ and } 4P = 2P + 2P;$$

$$5P = 4P + P \text{ and } 6P = 3P + 3P$$

$$\text{and } 7P = 4P + 3P.$$

Example: Typical algorithms

for fixed-base scalarmult

have many parallel point adds.

Example: A busy server
with a backlog of scalarmults
can vectorize across them.

Beware a disadvantage of
vectorizing across two mults:

256-bit f, f', g, g', h, h'

occupy at least 1536 bits,

leaving very little room

for temporary registers.

We use some loads and stores
inside vectorized `mulmul`.

Mostly invisible on Cortex-A8,

but bigger issue on Cortex-A7.

e: Recall

$$\cdot X_2; D = Y_1 \cdot Y_2$$

oint-addition formulas

ards curves.

e: Can compute

$$4P, 5P, 6P, 7P \text{ as}$$

$$+ P;$$

$$P + P \text{ and } 4P = 2P + 2P;$$

$$P + P \text{ and } 6P = 3P + 3P$$

$$= 4P + 3P.$$

e: Typical algorithms

-base scalarmult

ny parallel point adds.

Example: A busy server
with a backlog of scalarmults
can vectorize across them.

Beware a disadvantage of
vectorizing across two mults:

256-bit f, f', g, g', h, h'
occupy at least 1536 bits,
leaving very little room
for temporary registers.

We use some loads and stores
inside vectorized mulmul.

Mostly invisible on Cortex-A8,
but bigger issue on Cortex-A7.

Some fie
inside a

Example

convert

$$Z^{-1}X \in$$

Easy, co

$$11M + 2$$

$$z_2 = z_1$$

$$z_8 = z_2$$

$$z_9 = z_1$$

$$z_{11} = z_2$$

$$z_{22} = z_1$$

$$z_{_5_0} =$$

$$z_{_10_5} =$$

$Y_1 \cdot Y_2$
 on formulas

S.

compute

$7P$ as

$$4P = 2P + 2P;$$

$$6P = 3P + 3P$$

P .

algorithms

armult

l point adds.

Example: A busy server
 with a backlog of scalarmults
 can vectorize across them.

Beware a disadvantage of
 vectorizing across two mults:

256-bit f, f', g, g', h, h'

occupy at least 1536 bits,

leaving very little room

for temporary registers.

We use some loads and stores
 inside vectorized mulmul.

Mostly invisible on Cortex-A8,
 but bigger issue on Cortex-A7.

Some field ops are
 inside a single scalar

Example: At end of
 convert fraction (X)

$$Z^{-1}X \in \{0, 1, \dots\}$$

Easy, constant time

$$11\mathbf{M} + 254\mathbf{S} \text{ for } p$$

$$z2 = z1^{2^1}$$

$$z8 = z2^{2^2}$$

$$z9 = z1 * z8$$

$$z11 = z2 * z9$$

$$z22 = z11^{2^1}$$

$$z_{5_0} = z9 * z22$$

$$z_{10_5} = z_{5_0}^{2^2}$$

Example: A busy server with a backlog of scalarmults can vectorize across them.

Beware a disadvantage of vectorizing across two mults:

256-bit f, f', g, g', h, h' occupy at least 1536 bits, leaving very little room for temporary registers.

We use some loads and stores inside vectorized mulmul.

Mostly invisible on Cortex-A8, but bigger issue on Cortex-A7.

Some field ops are hard to p inside a single scalarmult.

Example: At end of ECDH, convert fraction $(X : Z)$ into $Z^{-1}X \in \{0, 1, \dots, p - 1\}$.

Easy, constant time: $Z^{-1} =$

11M + **254S** for $p = 2^{255} -$

$$z2 = z1^{2^1}$$

$$z8 = z2^{2^2}$$

$$z9 = z1 * z8$$

$$z11 = z2 * z9$$

$$z22 = z11^{2^1}$$

$$z_5_0 = z9 * z22$$

$$z_10_5 = z_5_0^{2^5}$$

Example: A busy server with a backlog of scalarmults can vectorize across them.

Beware a disadvantage of vectorizing across two mults:

256-bit f, f', g, g', h, h' occupy at least 1536 bits, leaving very little room for temporary registers.

We use some loads and stores inside vectorized mulmul.

Mostly invisible on Cortex-A8, but bigger issue on Cortex-A7.

Some field ops are hard to pair inside a single scalarmult.

Example: At end of ECDH, convert fraction $(X : Z)$ into $Z^{-1}X \in \{0, 1, \dots, p-1\}$.

Easy, constant time: $Z^{-1} = Z^{p-2}$.
11M + **254S** for $p = 2^{255} - 19$:

$$z2 = z1^{2^1}$$

$$z8 = z2^{2^2}$$

$$z9 = z1 * z8$$

$$z11 = z2 * z9$$

$$z22 = z11^{2^1}$$

$$z_5_0 = z9 * z22$$

$$z_10_5 = z_5_0^{2^5}$$

e: A busy server
 backlog of scalarmults
 orize across them.
 a disadvantage of
 ng across two mults:
 f, f', g, g', h, h'
 at least 1536 bits,
 very little room
 orary registers.
 some loads and stores
 ectorized mulmul.
 nvisible on Cortex-A8,
 er issue on Cortex-A7.

Some field ops are hard to pair inside a single scalarmult.

Example: At end of ECDH, convert fraction $(X : Z)$ into $Z^{-1}X \in \{0, 1, \dots, p - 1\}$.

Easy, constant time: $Z^{-1} = Z^{p-2}$.
11M + **254S** for $p = 2^{255} - 19$:

$$z2 = z1^{2^1}$$

$$z8 = z2^{2^2}$$

$$z9 = z1 * z8$$

$$z11 = z2 * z9$$

$$z22 = z11^{2^1}$$

$$z_5_0 = z9 * z22$$

$$z_10_5 = z_5_0^{2^5}$$

$z_10_0 =$
 $z_20_10 =$
 $z_20_0 =$
 $z_40_20 =$
 $z_40_0 =$
 $z_50_10 =$
 $z_50_0 =$
 $z_100_50 =$
 $z_100_0 =$
 $z_200_10 =$
 $z_200_0 =$
 $z_250_50 =$
 $z_250_0 =$
 $z_255_5 =$
 $z_255_2 =$

server
 scalarmults
 ss them.
 ntage of
 two mults:
 h, h'
 36 bits,
 room
 sters.
 s and stores
 mulmul.
 n Cortex-A8,
 n Cortex-A7.

Some field ops are hard to pair
 inside a single scalarmult.

Example: At end of ECDH,
 convert fraction $(X : Z)$ into
 $Z^{-1}X \in \{0, 1, \dots, p - 1\}$.

Easy, constant time: $Z^{-1} = Z^{p-2}$.
11M + **254S** for $p = 2^{255} - 19$:

$$z2 = z1^{2^1}$$

$$z8 = z2^{2^2}$$

$$z9 = z1 * z8$$

$$z11 = z2 * z9$$

$$z22 = z11^{2^1}$$

$$z_5_0 = z9 * z22$$

$$z_10_5 = z_5_0^{2^5}$$

$z_10_0 = z_10_5 * z_5_0$
 $z_20_10 = z_10_0^{2^1}$
 $z_20_0 = z_20_10^{2^1}$
 $z_40_20 = z_20_0^{2^2}$
 $z_40_0 = z_40_20^{2^1}$
 $z_50_10 = z_40_0^{2^1}$
 $z_50_0 = z_50_10^{2^1}$
 $z_100_50 = z_50_0^{2^2}$
 $z_100_0 = z_100_50^{2^1}$
 $z_200_100 = z_100_0^{2^2}$
 $z_200_0 = z_200_100^{2^1}$
 $z_250_50 = z_200_0^{2^1}$
 $z_250_0 = z_250_50^{2^1}$
 $z_255_5 = z_250_0^{2^5}$
 $z_255_21 = z_255_5^{2^1}$

Some field ops are hard to pair inside a single scalarmult.

Example: At end of ECDH, convert fraction $(X : Z)$ into $Z^{-1}X \in \{0, 1, \dots, p - 1\}$.

Easy, constant time: $Z^{-1} = Z^{p-2}$.

11M + **254S** for $p = 2^{255} - 19$:

$$z_2 = z_1^{2^1}$$

$$z_8 = z_2^{2^2}$$

$$z_9 = z_1 * z_8$$

$$z_{11} = z_2 * z_9$$

$$z_{22} = z_{11}^{2^1}$$

$$z_{5_0} = z_9 * z_{22}$$

$$z_{10_5} = z_{5_0}^{2^5}$$

$$z_{10_0} = z_{10_5} * z_{5_0}$$

$$z_{20_{10}} = z_{10_0}^{2^{10}}$$

$$z_{20_0} = z_{20_{10}} * z_{10_0}$$

$$z_{40_{20}} = z_{20_0}^{2^{20}}$$

$$z_{40_0} = z_{40_{20}} * z_{20_0}$$

$$z_{50_{10}} = z_{40_0}^{2^{10}}$$

$$z_{50_0} = z_{50_{10}} * z_{10_0}$$

$$z_{100_{50}} = z_{50_0}^{2^{50}}$$

$$z_{100_0} = z_{100_{50}} * z_{50_0}$$

$$z_{200_{100}} = z_{100_0}^{2^{100}}$$

$$z_{200_0} = z_{200_{100}} * z_{100_0}$$

$$z_{250_{50}} = z_{200_0}^{2^{50}}$$

$$z_{250_0} = z_{250_{50}} * z_{50_0}$$

$$z_{255_5} = z_{250_0}^{2^5}$$

$$z_{255_{21}} = z_{255_5} * z_{11}$$

Some field ops are hard to pair inside a single scalarmult.

Example: At end of ECDH, convert fraction $(X : Z)$ into $Z^{-1}X \in \{0, 1, \dots, p - 1\}$.

Easy, constant time: $Z^{-1} = Z^{p-2}$.

11M + **254S** for $p = 2^{255} - 19$:

$$z_2 = z_1^{2^1}$$

$$z_8 = z_2^{2^2}$$

$$z_9 = z_1 * z_8$$

$$z_{11} = z_2 * z_9$$

$$z_{22} = z_{11}^{2^1}$$

$$z_{5_0} = z_9 * z_{22}$$

$$z_{10_5} = z_{5_0}^{2^5}$$

$$z_{10_0} = z_{10_5} * z_{5_0}$$

$$z_{20_{10}} = z_{10_0}^{2^{10}}$$

$$z_{20_0} = z_{20_{10}} * z_{10_0}$$

$$z_{40_{20}} = z_{20_0}^{2^{20}}$$

$$z_{40_0} = z_{40_{20}} * z_{20_0}$$

$$z_{50_{10}} = z_{40_0}^{2^{10}}$$

$$z_{50_0} = z_{50_{10}} * z_{10_0}$$

$$z_{100_{50}} = z_{50_0}^{2^{50}}$$

$$z_{100_0} = z_{100_{50}} * z_{50_0}$$

$$z_{200_{100}} = z_{100_0}^{2^{100}}$$

$$z_{200_0} = z_{200_{100}} * z_{100_0}$$

$$z_{250_{50}} = z_{200_0}^{2^{50}}$$

$$z_{250_0} = z_{250_{50}} * z_{50_0}$$

$$z_{255_5} = z_{250_0}^{2^5}$$

$$z_{255_{21}} = z_{255_5} * z_{11}$$

field ops are hard to pair
single scalarmult.

e: At end of ECDH,
fraction $(X : Z)$ into
 $\{0, 1, \dots, p - 1\}$.

constant time: $Z^{-1} = Z^{p-2}$.

254S for $p = 2^{255} - 19$:

z^2

z^4

z^8

z^{16}

z^{32}

z^{64}

z^{128}

$$z_{10_0} = z_{10_5} * z_{5_0}$$

$$z_{20_{10}} = z_{10_0}^{2^{10}}$$

$$z_{20_0} = z_{20_{10}} * z_{10_0}$$

$$z_{40_{20}} = z_{20_0}^{2^{20}}$$

$$z_{40_0} = z_{40_{20}} * z_{20_0}$$

$$z_{50_{10}} = z_{40_0}^{2^{10}}$$

$$z_{50_0} = z_{50_{10}} * z_{10_0}$$

$$z_{100_{50}} = z_{50_0}^{2^{50}}$$

$$z_{100_0} = z_{100_{50}} * z_{50_0}$$

$$z_{200_{100}} = z_{100_0}^{2^{100}}$$

$$z_{200_0} = z_{200_{100}} * z_{100_0}$$

$$z_{250_{50}} = z_{200_0}^{2^{50}}$$

$$z_{250_0} = z_{250_{50}} * z_{50_0}$$

$$z_{255_5} = z_{250_0}^{2^5}$$

$$z_{255_{21}} = z_{255_5} * z_{11}$$

Can still
inside a

Strategy

50 mul i

$(f_0, 2f_1), (f_2,$

$(f_1, f_8), (f_3, t$

$(g_0, g_1), (g_2,$

$(g_0, 19g_1), ($

$(19g_2, 19g_3)$

$(19g_2, g_3), ($

Change

e.g., $(h_0,$

hard to pair
armult.

of ECDH,
($X : Z$) into
, $p - 1$ }.

ie: $Z^{-1} = Z^{p-2}$.

$p = 2^{255} - 19$:

```

z_10_0 = z_10_5*z_5_0
z_20_10 = z_10_0^2^10
z_20_0 = z_20_10*z_10_0
z_40_20 = z_20_0^2^20
z_40_0 = z_40_20*z_20_0
z_50_10 = z_40_0^2^10
z_50_0 = z_50_10*z_10_0
z_100_50 = z_50_0^2^50
z_100_0 = z_100_50*z_50_0
z_200_100 = z_100_0^2^100
z_200_0 = z_200_100*z_100_0
z_250_50 = z_200_0^2^50
z_250_0 = z_250_50*z_50_0
z_255_5 = z_250_0^2^5
z_255_21 = z_255_5*z11

```

5

Can still vectorize
inside a single field

Strategy in our so

50 mul insns start

$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5),$

$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7,$

$(g_0, g_1), (g_2, g_3), (g_4, g_5), ($

$(g_0, 19g_1), (g_2, 19g_3), (g_4,$

$(19g_2, 19g_3), (19g_4, 19g_5),$

$(19g_2, g_3), (19g_4, g_5), (19g$

Change carry patt

e.g., $(h_0, h_4) \rightarrow (h$

pair

$$z_{10_0} = z_{10_5} * z_{5_0}$$

$$z_{20_{10}} = z_{10_0}^{2^{10}}$$

$$z_{20_0} = z_{20_{10}} * z_{10_0}$$

$$z_{40_{20}} = z_{20_0}^{2^{20}}$$

$$z_{40_0} = z_{40_{20}} * z_{20_0}$$

$$z_{50_{10}} = z_{40_0}^{2^{10}}$$

$$z_{50_0} = z_{50_{10}} * z_{10_0}$$

 Z^{p-2}

19:

$$z_{100_{50}} = z_{50_0}^{2^{50}}$$

$$z_{100_0} = z_{100_{50}} * z_{50_0}$$

$$z_{200_{100}} = z_{100_0}^{2^{100}}$$

$$z_{200_0} = z_{200_{100}} * z_{100_0}$$

$$z_{250_{50}} = z_{200_0}^{2^{50}}$$

$$z_{250_0} = z_{250_{50}} * z_{50_0}$$

$$z_{255_5} = z_{250_0}^{2^5}$$

$$z_{255_{21}} = z_{255_5} * z_{11}$$

Can still vectorize
inside a single field op.

Strategy in our software:

50 mul insns starting from

$$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9)$$

$$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$$

$$(g_0, g_1), (g_2, g_3), (g_4, g_5), (g_6, g_7);$$

$$(g_0, 19g_1), (g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7);$$

$$(19g_2, 19g_3), (19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$$

$$(19g_2, g_3), (19g_4, g_5), (19g_6, g_7), (19g_8, g_9);$$

Change carry pattern to vec

e.g., $(h_0, h_4) \rightarrow (h_1, h_5)$.

```

z_10_0 = z_10_5*z_5_0
z_20_10 = z_10_0^2^10
z_20_0 = z_20_10*z_10_0
z_40_20 = z_20_0^2^20
z_40_0 = z_40_20*z_20_0
z_50_10 = z_40_0^2^10
z_50_0 = z_50_10*z_10_0
z_100_50 = z_50_0^2^50
z_100_0 = z_100_50*z_50_0
z_200_100 = z_100_0^2^100
z_200_0 = z_200_100*z_100_0
z_250_50 = z_200_0^2^50
z_250_0 = z_250_50*z_50_0
z_255_5 = z_250_0^2^5
z_255_21 = z_255_5*z11

```

Can still vectorize
inside a single field op.

Strategy in our software:

50 mul insns starting from

$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9);$

$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$

$(g_0, g_1), (g_2, g_3), (g_4, g_5), (g_6, g_7);$

$(g_0, 19g_1), (g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_2, 19g_3), (19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$

$(19g_2, g_3), (19g_4, g_5), (19g_6, g_7), (19g_8, g_9).$

Change carry pattern to vectorize,
e.g., $(h_0, h_4) \rightarrow (h_1, h_5).$

```

= z_10_5*z_5_0
= z_10_0^2^10
= z_20_10*z_10_0
= z_20_0^2^20
= z_40_20*z_20_0
= z_40_0^2^10
= z_50_10*z_10_0
0 = z_50_0^2^50
= z_100_50*z_50_0
00 = z_100_0^2^100
= z_200_100*z_100_0
0 = z_200_0^2^50
= z_250_50*z_50_0
= z_250_0^2^5
1 = z_255_5*z11

```

Can still vectorize
inside a single field op.

Strategy in our software:

50 mul insns starting from

$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9);$

$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$

$(g_0, g_1), (g_2, g_3), (g_4, g_5), (g_6, g_7);$

$(g_0, 19g_1), (g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_2, 19g_3), (19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$

$(19g_2, g_3), (19g_4, g_5), (19g_6, g_7), (19g_8, g_9).$

Change carry pattern to vectorize,
e.g., $(h_0, h_4) \rightarrow (h_1, h_5).$

Core arith
on mul i
Squaring

Some lo

ECDH: \approx

More de

356019 \approx

$\approx 78\%$ o

Cortex-A

Still som

z_{5_0}
 $^2^{10}$
 $*z_{10_0}$
 $^2^{20}$
 $*z_{20_0}$
 $^2^{10}$
 $*z_{10_0}$
 0^2^{50}
 $50*z_{50_0}$
 $0_0^2^{100}$
 $100*z_{100_0}$
 $_0^2^{50}$
 $50*z_{50_0}$
 0^2^5
 $_5*z_{11}$

Can still vectorize
inside a single field op.

Strategy in our software:

50 mul insns starting from

$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9);$

$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$

$(g_0, g_1), (g_2, g_3), (g_4, g_5), (g_6, g_7);$

$(g_0, 19g_1), (g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_2, 19g_3), (19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$

$(19g_2, g_3), (19g_4, g_5), (19g_6, g_7), (19g_8, g_9).$

Change carry pattern to vectorize,
e.g., $(h_0, h_4) \rightarrow (h_1, h_5).$

Core arithmetic: 1
on mul insns for e

Squarings are some

Some loss for carr

ECDH: ≈ 10 field m

More detailed ana

356019 cycles on a

$\approx 78\%$ of software

Cortex-A8-fast cyc

Still some room fo

Can still vectorize
inside a single field op.

Strategy in our software:

50 mul insns starting from

$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9);$

$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$

$(g_0, g_1), (g_2, g_3), (g_4, g_5), (g_6, g_7);$

$(g_0, 19g_1), (g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_2, 19g_3), (19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$

$(19g_2, g_3), (19g_4, g_5), (19g_6, g_7), (19g_8, g_9).$

Change carry pattern to vectorize,
e.g., $(h_0, h_4) \rightarrow (h_1, h_5).$

Core arithmetic: 100 cycles
on mul insns for each field m

Squarings are somewhat fast

Some loss for carries etc.

ECDH: ≈ 10 field muls $\cdot 255$

More detailed analysis:

356019 cycles on arithmetic,

$\approx 78\%$ of software's total

Cortex-A8-fast cycles for EC

Still some room for improve

Can still vectorize
inside a single field op.

Strategy in our software:

50 mul insns starting from

$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9);$

$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$

$(g_0, g_1), (g_2, g_3), (g_4, g_5), (g_6, g_7);$

$(g_0, 19g_1), (g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_2, 19g_3), (19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$

$(19g_2, g_3), (19g_4, g_5), (19g_6, g_7), (19g_8, g_9).$

Change carry pattern to vectorize,

e.g., $(h_0, h_4) \rightarrow (h_1, h_5).$

Core arithmetic: 100 cycles
on mul insns for each field mul.
Squarings are somewhat faster.

Some loss for carries etc.

ECDH: ≈ 10 field muls \cdot 255 bits.

More detailed analysis:

356019 cycles on arithmetic;

$\approx 78\%$ of software's total

Cortex-A8-fast cycles for ECDH.

Still some room for improvement.

Can still vectorize
inside a single field op.

Strategy in our software:

50 mul insns starting from

$(f_0, 2f_1), (f_2, 2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9);$

$(f_1, f_8), (f_3, f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$

$(g_0, g_1), (g_2, g_3), (g_4, g_5), (g_6, g_7);$

$(g_0, 19g_1), (g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_2, 19g_3), (19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$

$(19g_2, g_3), (19g_4, g_5), (19g_6, g_7), (19g_8, g_9).$

Change carry pattern to vectorize,
e.g., $(h_0, h_4) \rightarrow (h_1, h_5).$

Core arithmetic: 100 cycles
on mul insns for each field mul.
Squarings are somewhat faster.

Some loss for carries etc.

ECDH: ≈ 10 field muls \cdot 255 bits.

More detailed analysis:

356019 cycles on arithmetic;

$\approx 78\%$ of software's total

Cortex-A8-fast cycles for ECDH.

Still some room for improvement.

Each CPU is a new adventure.

e.g. Could it be better to use

Cortex-A7 FPU with radix $2^{21.25}$?

vectorize

single field op.

in our software:

insns starting from

$(2f_3), (f_4, 2f_5), (f_6, 2f_7), (f_8, 2f_9);$

$(f_0), (f_5, f_2), (f_7, f_4), (f_9, f_6);$

$(g_3), (g_4, g_5), (g_6, g_7);$

$(g_2, 19g_3), (g_4, 19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_4, 19g_5), (19g_6, 19g_7), (19g_8, 19g_9);$

$(19g_4, g_5), (19g_6, g_7), (19g_8, g_9).$

carry pattern to vectorize,

$(h_4) \rightarrow (h_1, h_5).$

Core arithmetic: 100 cycles

on mul insns for each field mul.

Squarings are somewhat faster.

Some loss for carries etc.

ECDH: ≈ 10 field muls $\cdot 255$ bits.

More detailed analysis:

356019 cycles on arithmetic;

$\approx 78\%$ of software's total

Cortex-A8-fast cycles for ECDH.

Still some room for improvement.

Each CPU is a new adventure.

e.g. Could it be better to use

Cortex-A7 FPU with radix $2^{21.25}$?

Much m

<https://>

benchma

2137 pu

hundred

39 DH p

56 signa

304 auth

d op.

ftware:

ing from

$(f_6, 2f_7), (f_8, 2f_9);$

$(f_4), (f_9, f_6);$

$(g_6, g_7);$

$(19g_5), (g_6, 19g_7), (g_8, 19g_9);$

$(19g_6, 19g_7), (19g_8, 19g_9);$

$(g_6, g_7), (19g_8, g_9).$

ern to vectorize,

$(h_1, h_5).$

Core arithmetic: 100 cycles
on mul insns for each field mul.
Squarings are somewhat faster.

Some loss for carries etc.

ECDH: ≈ 10 field muls $\cdot 255$ bits.

More detailed analysis:

356019 cycles on arithmetic;

$\approx 78\%$ of software's total

Cortex-A8-fast cycles for ECDH.

Still some room for improvement.

Each CPU is a new adventure.

e.g. Could it be better to use

Cortex-A7 FPU with radix $2^{21.25}$?

Much more work to

<https://bench.c>

benchmarks for (c

2137 public implem

hundreds of crypt

39 DH primitives,

56 signature primi

304 authenticated

Core arithmetic: 100 cycles
on mul insns for each field mul.
Squarings are somewhat faster.

Some loss for carries etc.

ECDH: ≈ 10 field muls \cdot 255 bits.

More detailed analysis:

356019 cycles on arithmetic;

$\approx 78\%$ of software's total

Cortex-A8-fast cycles for ECDH.

Still some room for improvement.

Each CPU is a new adventure.

e.g. Could it be better to use

Cortex-A7 FPU with radix $2^{21.25}$?

Much more work to do

<https://bench.cr.yp.to>:

benchmarks for (currently)

2137 public implementations

hundreds of crypto primitive

39 DH primitives,

56 signature primitives,

304 authenticated ciphers, e

Core arithmetic: 100 cycles
on mul insns for each field mul.
Squarings are somewhat faster.
Some loss for carries etc.
ECDH: ≈ 10 field muls \cdot 255 bits.
More detailed analysis:
356019 cycles on arithmetic;
 $\approx 78\%$ of software's total
Cortex-A8-fast cycles for ECDH.
Still some room for improvement.
Each CPU is a new adventure.
e.g. Could it be better to use
Cortex-A7 FPU with radix $2^{21.25}$?

Much more work to do

<https://bench.cr.yp.to>:
benchmarks for (currently)
2137 public implementations of
hundreds of crypto primitives—
39 DH primitives,
56 signature primitives,
304 authenticated ciphers, etc.

Core arithmetic: 100 cycles
 on mul insns for each field mul.
 Squarings are somewhat faster.
 Some loss for carries etc.

ECDH: ≈ 10 field muls \cdot 255 bits.

More detailed analysis:
 356019 cycles on arithmetic;
 $\approx 78\%$ of software's total
 Cortex-A8-fast cycles for ECDH.
 Still some room for improvement.

Each CPU is a new adventure.
 e.g. Could it be better to use
 Cortex-A7 FPU with radix $2^{21.25}$?

Much more work to do

<https://bench.cr.yp.to>:
 benchmarks for (currently)
 2137 public implementations of
 hundreds of crypto primitives—
 39 DH primitives,
 56 signature primitives,
 304 authenticated ciphers, etc.

Many interesting primitives
 are far slower than necessary
 on many important CPUs.

Core arithmetic: 100 cycles
 on mul insns for each field mul.
 Squarings are somewhat faster.
 Some loss for carries etc.

ECDH: ≈ 10 field muls \cdot 255 bits.

More detailed analysis:
 356019 cycles on arithmetic;
 $\approx 78\%$ of software's total
 Cortex-A8-fast cycles for ECDH.
 Still some room for improvement.

Each CPU is a new adventure.
 e.g. Could it be better to use
 Cortex-A7 FPU with radix $2^{21.25}$?

Much more work to do

<https://bench.cr.yp.to>:
 benchmarks for (currently)
 2137 public implementations of
 hundreds of crypto primitives—
 39 DH primitives,
 56 signature primitives,
 304 authenticated ciphers, etc.

Many interesting primitives
 are far slower than necessary
 on many important CPUs.

Exercise: Make them faster!