

# High-speed high-security signatures

Peter Schwabe

National Taiwan University



Joint work with Daniel J. Bernstein, Niels Duif,  
Tanja Lange, and Bo-Yin Yang

September 14, 2011

EiPSI Seminar

A look back...

臺灣大學



A look back...



“Do you *really* think you can get a Ph.D. without even *mentioning* Edwards curves in your thesis?”

# A new start: Work on Edwards signatures

臺灣大學



# Requirements



- ▶ Elliptic-curve signatures using twisted Edwards curves
- ▶ 128 bits of security

# Requirements

- ▶ Elliptic-curve signatures using twisted Edwards curves
- ▶ 128 bits of security
- ▶ Support for batch verification
- ▶ Timing-attack resistant implementation
- ▶ Foolproof session keys
- ▶ Hash-function collision resilience

# Requirements

- ▶ Elliptic-curve signatures using twisted Edwards curves
- ▶ 128 bits of security
- ▶ Support for batch verification
- ▶ Timing-attack resistant implementation
- ▶ Foolproof session keys
- ▶ Hash-function collision resilience
- ▶ The usual: make it fast

# Requirements

- ▶ Elliptic-curve signatures using twisted Edwards curves
- ▶ 128 bits of security
- ▶ Support for batch verification
- ▶ Timing-attack resistant implementation
- ▶ Foolproof session keys
- ▶ Hash-function collision resilience
- ▶ The usual: make it fast
  - ▶ Fast signing
  - ▶ Fast verification
  - ▶ Faster batch verification
  - ▶ Fast key generation



# The EdDSA signature system



## EdDSA and Ed25519 parameters

EdDSA

- ▶ Integer  $b \geq 10$

Ed25519

- ▶  $b = 256$

## EdDSA and Ed25519 parameters

## EdDSA

- ▶ Integer  $b \geq 10$
- ▶ Prime power  $q \equiv 1 \pmod{4}$
- ▶  $(b - 1)$ -bit encoding of elements of  $\mathbb{F}_q$

## Ed25519

- ▶  $b = 256$
- ▶  $q = 2^{255} - 19$  (prime)
- ▶ little-endian encoding of  $\{0, \dots, 2^{255} - 20\}$

## EdDSA and Ed25519 parameters

## EdDSA

- ▶ Integer  $b \geq 10$
- ▶ Prime power  $q \equiv 1 \pmod{4}$
- ▶  $(b - 1)$ -bit encoding of elements of  $\mathbb{F}_q$
- ▶ Hash function  $H$  with  $2b$ -bit output

## Ed25519

- ▶  $b = 256$
- ▶  $q = 2^{255} - 19$  (prime)
- ▶ little-endian encoding of  $\{0, \dots, 2^{255} - 20\}$
- ▶  $H = \text{SHA-512}$

# EdDSA and Ed25519 parameters

## EdDSA

- ▶ Integer  $b \geq 10$
- ▶ Prime power  $q \equiv 1 \pmod{4}$
- ▶  $(b-1)$ -bit encoding of elements of  $\mathbb{F}_q$
- ▶ Hash function  $H$  with  $2b$ -bit output
- ▶ Non-square  $d \in \mathbb{F}_q$
- ▶  $B \in \{(x, y) \in \mathbb{F}_q \times \mathbb{F}_q, -x^2 + y^2 = 1 + dx^2y^2\}$   
(twisted Edwards curve  $E$ )
- ▶ prime  $\ell \in (2^{b-4}, 2^{b-3})$  with  $\ell B = (0, 1)$

## Ed25519

- ▶  $b = 256$
- ▶  $q = 2^{255} - 19$  (prime)
- ▶ little-endian encoding of  $\{0, \dots, 2^{255} - 20\}$
- ▶  $H = \text{SHA-512}$
- ▶  $d = -121665/121666$
- ▶  $B = (x, 4/5)$ , with  $x$  “even”
- ▶  $\ell$  a 253-bit prime

# EdDSA and Ed25519 parameters

## EdDSA

- ▶ Integer  $b \geq 10$
- ▶ Prime power  $q \equiv 1 \pmod{4}$
- ▶  $(b - 1)$ -bit encoding of elements of  $\mathbb{F}_q$
- ▶ Hash function  $H$  with  $2b$ -bit output
- ▶ Non-square  $d \in \mathbb{F}_q$
- ▶  $B \in \{(x, y) \in \mathbb{F}_q \times \mathbb{F}_q, -x^2 + y^2 = 1 + dx^2y^2\}$   
(twisted Edwards curve  $E$ )
- ▶ prime  $\ell \in (2^{b-4}, 2^{b-3})$  with  $\ell B = (0, 1)$

## Ed25519

- ▶  $b = 256$
- ▶  $q = 2^{255} - 19$  (prime)
- ▶ little-endian encoding of  $\{0, \dots, 2^{255} - 20\}$
- ▶  $H = \text{SHA-512}$
- ▶  $d = -121665/121666$
- ▶  $B = (x, 4/5)$ , with  $x$  “even”
- ▶  $\ell$  a 253-bit prime

Ed25519 curve is birationally equivalent to the Curve25519 curve.

## EdDSA keys

- ▶ Secret key:  $b$ -bit string  $k$
- ▶ Compute  $H(k) = (h_0, \dots, h_{2b-1})$

## EdDSA keys

- ▶ Secret key:  $b$ -bit string  $k$
- ▶ Compute  $H(k) = (h_0, \dots, h_{2b-1})$
- ▶ Derive integer  $a = 2^{b-2} + \sum_{3 \leq i \leq b-3} 2^i h_i$
- ▶ Note that  $a$  is a multiple of 8



## EdDSA keys

- ▶ Secret key:  $b$ -bit string  $k$
- ▶ Compute  $H(k) = (h_0, \dots, h_{2b-1})$
- ▶ Derive integer  $a = 2^{b-2} + \sum_{3 \leq i \leq b-3} 2^i h_i$
- ▶ Note that  $a$  is a multiple of 8
- ▶ Compute  $A = aB$
- ▶ Public key: Encoding  $\underline{A}$  of  $A = (x_A, y_A)$  as  $y_A$  and one (parity) bit of  $x_A$  (needs  $b$  bits)

## EdDSA keys

- ▶ Secret key:  $b$ -bit string  $k$
- ▶ Compute  $H(k) = (h_0, \dots, h_{2b-1})$
- ▶ Derive integer  $a = 2^{b-2} + \sum_{3 \leq i \leq b-3} 2^i h_i$
- ▶ Note that  $a$  is a multiple of 8
- ▶ Compute  $A = aB$
- ▶ Public key: Encoding  $\underline{A}$  of  $A = (x_A, y_A)$  as  $y_A$  and one (parity) bit of  $x_A$  (needs  $b$  bits)
- ▶ Compute  $A$  from  $\underline{A}$ :  $x_A = \pm \sqrt{(y_A^2 - 1)/(dy_A^2 + 1)}$

# EdDSA signatures

## Signing

- ▶ Message  $M$  determines  $r = H(h_b, \dots, h_{2b-1}, M) \in \{0, \dots, 2^{2b} - 1\}$
- ▶ Define  $R = rB$
- ▶ Define  $S = (r + H(\underline{R}, \underline{A}, M)a) \bmod \ell$
- ▶ Signature:  $(\underline{R}, \underline{S})$ , with  $\underline{S}$  the  $b$ -bit little-endian encoding of  $S$
- ▶  $(\underline{R}, \underline{S})$  has  $2b$  bits (3 known to be zero)

# EdDSA signatures

## Signing

- ▶ Message  $M$  determines  $r = H(h_b, \dots, h_{2b-1}, M) \in \{0, \dots, 2^{2b} - 1\}$
- ▶ Define  $R = rB$
- ▶ Define  $S = (r + H(\underline{R}, \underline{A}, M)a) \pmod{\ell}$
- ▶ Signature:  $(\underline{R}, \underline{S})$ , with  $\underline{S}$  the  $b$ -bit little-endian encoding of  $S$
- ▶  $(\underline{R}, \underline{S})$  has  $2b$  bits (3 known to be zero)

## Verification

- ▶ Verifier parses  $A$  from  $\underline{A}$  and  $R$  from  $\underline{R}$
- ▶ Computes  $H(\underline{R}, \underline{A}, M)$
- ▶ Checks group equation

$$8SB = 8R + 8H(\underline{R}, \underline{A}, M)A$$

- ▶ Rejects if parsing fails or equation does not hold

# Security features of EdDSA



## Collision resilience

- ▶ ECDSA uses  $H(M)$
- ▶ Collisions in  $H$  allow existential forgery

# Collision resilience

- ▶ ECDSA uses  $H(M)$
- ▶ Collisions in  $H$  allow existential forgery
- ▶ Schnorr signatures and EdDSA include  $\underline{R}$  in the hash
  - ▶ Schnorr:  $H(\underline{R}, M)$
  - ▶ EdDSA:  $H(\underline{R}, \underline{A}, M)$
- ▶ Signatures are hash-function-collision resilient

# Collision resilience

- ▶ ECDSA uses  $H(M)$
- ▶ Collisions in  $H$  allow existential forgery
- ▶ Schnorr signatures and EdDSA include  $\underline{R}$  in the hash
  - ▶ Schnorr:  $H(\underline{R}, M)$
  - ▶ EdDSA:  $H(\underline{R}, \underline{A}, M)$
- ▶ Signatures are hash-function-collision resilient
- ▶ Including  $\underline{A}$  alleviates concerns about attacks against multiple keys



## Foolproof session keys

- ▶ Each message needs a different  $r$  (“session key”)
- ▶ Just knowing a few bits of  $r$  allows to recover  $a$
- ▶ Usual approach (e.g., Schnorr signatures): Choose random  $r$  for each message

## Foolproof session keys

- ▶ Each message needs a different  $r$  (“session key”)
- ▶ Just knowing a few bits of  $r$  allows to recover  $a$
- ▶ Usual approach (e.g., Schnorr signatures): Choose random  $r$  for each message
- ▶ Potential problems: Bad random-number generators, off-by-one(-byte) bugs

## Foolproof session keys

- ▶ Each message needs a different  $r$  (“session key”)
- ▶ Just knowing a few bits of  $r$  allows to recover  $a$
- ▶ Usual approach (e.g., Schnorr signatures): Choose random  $r$  for each message
- ▶ Potential problems: Bad random-number generators, off-by-one(-byte) bugs
- ▶ Even worse: No random-number generator: Sony’s PS3 security disaster

## Foolproof session keys

- ▶ Each message needs a different  $r$  (“session key”)
- ▶ Just knowing a few bits of  $r$  allows to recover  $a$
- ▶ Usual approach (e.g., Schnorr signatures): Choose random  $r$  for each message
- ▶ Potential problems: Bad random-number generators, off-by-one(-byte) bugs
- ▶ Even worse: No random-number generator: Sony’s PS3 security disaster
- ▶ EdDSA uses deterministic, pseudo-random session keys  
 $H(h_b, \dots, h_{2b-1}, M)$

## Foolproof session keys

- ▶ Each message needs a different  $r$  (“session key”)
- ▶ Just knowing a few bits of  $r$  allows to recover  $a$
- ▶ Usual approach (e.g., Schnorr signatures): Choose random  $r$  for each message
- ▶ Potential problems: Bad random-number generators, off-by-one(-byte) bugs
- ▶ Even worse: No random-number generator: Sony’s PS3 security disaster
- ▶ EdDSA uses deterministic, pseudo-random session keys  
 $H(h_b, \dots, h_{2b-1}, M)$
- ▶ Same security as Schnorr under standard PRF assumptions
- ▶ Does not consume per-message randomness
- ▶ Better for testing (deterministic output)

# Speed of Ed25519



## Fast constant-time implementation

- ▶ Recent paper by Brumley and Taveri: remote timing attack against ECDSA implementation in OpenSSL

# Fast constant-time implementation

- ▶ Recent paper by Brumley and Tuveri: remote timing attack against ECDSA implementation in OpenSSL
- ▶ Protection against timing attacks means:
  - ▶ No data flow from secret data into branch conditions
  - ▶ No data flow from secret data into load indices



# Fast constant-time implementation

- ▶ Recent paper by Brumley and Tuveri: remote timing attack against ECDSA implementation in OpenSSL
- ▶ Protection against timing attacks means:
  - ▶ No data flow from secret data into branch conditions
  - ▶ No data flow from secret data into load indices
- ▶ Choose constant-time scalar-multiplication algorithms
- ▶ Substitute table lookups by arithmetic

## Fast signing

- ▶ Main computational task: Compute  $R = rB$

# Fast signing

- ▶ Main computational task: Compute  $R = rB$
- ▶ First compute  $r \bmod \ell$ , write it as  $r_0 + 16r_1 + \cdots + 16^{63}r_{63}$ , with

$$r_i \in \{-8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7\}$$

# Fast signing

- ▶ Main computational task: Compute  $R = rB$
- ▶ First compute  $r \bmod \ell$ , write it as  $r_0 + 16r_1 + \dots + 16^{63}r_{63}$ , with

$$r_i \in \{-8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7\}$$

- ▶ Precompute  $16^i|r_i|B$  for  $i = 0, \dots, 63$  and  $|r_i| \in \{1, \dots, 8\}$ , in a lookup table at compile time

# Fast signing

- ▶ Main computational task: Compute  $R = rB$
- ▶ First compute  $r \bmod \ell$ , write it as  $r_0 + 16r_1 + \dots + 16^{63}r_{63}$ , with

$$r_i \in \{-8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7\}$$

- ▶ Precompute  $16^i|r_i|B$  for  $i = 0, \dots, 63$  and  $|r_i| \in \{1, \dots, 8\}$ , in a lookup table at compile time
- ▶ Compute  $R = \sum_{i=0}^{63} 16^i r_i B$

# Fast signing

- ▶ Main computational task: Compute  $R = rB$
- ▶ First compute  $r \bmod \ell$ , write it as  $r_0 + 16r_1 + \dots + 16^{63}r_{63}$ , with
$$r_i \in \{-8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7\}$$
- ▶ Precompute  $16^i|r_i|B$  for  $i = 0, \dots, 63$  and  $|r_i| \in \{1, \dots, 8\}$ , in a lookup table at compile time
- ▶ Compute  $R = \sum_{i=0}^{63} 16^i r_i B$
- ▶ 64 table lookups, 64 conditional point negations, 63 point additions

# Fast signing

- ▶ Main computational task: Compute  $R = rB$
- ▶ First compute  $r \bmod \ell$ , write it as  $r_0 + 16r_1 + \dots + 16^{63}r_{63}$ , with
$$r_i \in \{-8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7\}$$
- ▶ Precompute  $16^i|r_i|B$  for  $i = 0, \dots, 63$  and  $|r_i| \in \{1, \dots, 8\}$ , in a lookup table at compile time
- ▶ Compute  $R = \sum_{i=0}^{63} 16^i r_i B$
- ▶ 64 table lookups, 64 conditional point negations, 63 point additions
- ▶ Wait, lookups?

# Fast signing

- ▶ Main computational task: Compute  $R = rB$
- ▶ First compute  $r \bmod \ell$ , write it as  $r_0 + 16r_1 + \dots + 16^{63}r_{63}$ , with
$$r_i \in \{-8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7\}$$
- ▶ Precompute  $16^i|r_i|B$  for  $i = 0, \dots, 63$  and  $|r_i| \in \{1, \dots, 8\}$ , in a lookup table at compile time
- ▶ Compute  $R = \sum_{i=0}^{63} 16^i r_i B$
- ▶ 64 table lookups, 64 conditional point negations, 63 point additions
- ▶ Wait, lookups?
- ▶ In each lookup load all 8 relevant entries from the table, use arithmetic to obtain the desired one



# Fast signing

- ▶ Main computational task: Compute  $R = rB$
- ▶ First compute  $r \bmod \ell$ , write it as  $r_0 + 16r_1 + \dots + 16^{63}r_{63}$ , with
$$r_i \in \{-8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7\}$$
- ▶ Precompute  $16^i|r_i|B$  for  $i = 0, \dots, 63$  and  $|r_i| \in \{1, \dots, 8\}$ , in a lookup table at compile time
- ▶ Compute  $R = \sum_{i=0}^{63} 16^i r_i B$
- ▶ 64 table lookups, 64 conditional point negations, 63 point additions
- ▶ Wait, lookups?
- ▶ In each lookup load all 8 relevant entries from the table, use arithmetic to obtain the desired one
- ▶ Signing takes 88,328 cycles on an Intel Westmere CPU
- ▶ Key generation takes about 6,000 cycles more (read from `/dev/urandom`)

## Fast verification

- ▶ First part: point decompression, compute  $x$  coordinate  $x_R$  of  $R$  as

$$x_R = \pm \sqrt{(y_R^2 - 1)/(dy_R^2 + 1)}$$

- ▶ Looks like a square root and an inversion is required

## Fast verification

- ▶ First part: point decompression, compute  $x$  coordinate  $x_R$  of  $R$  as

$$x_R = \pm \sqrt{(y_R^2 - 1)/(dy_R^2 + 1)}$$

- ▶ Looks like a square root and an inversion is required
- ▶ As  $q \equiv 5 \pmod{8}$  for each square  $\alpha$  we have  $\alpha^2 = \beta^4$ , with  $\beta = \alpha^{(q+3)/8}$
- ▶ Standard: Compute  $\beta$ , conditionally multiply by  $\sqrt{-1}$  if  $\beta^2 = -\alpha$

## Fast verification

- ▶ First part: point decompression, compute  $x$  coordinate  $x_R$  of  $R$  as

$$x_R = \pm \sqrt{(y_R^2 - 1)/(dy_R^2 + 1)}$$

- ▶ Looks like a square root and an inversion is required
- ▶ As  $q \equiv 5 \pmod{8}$  for each square  $\alpha$  we have  $\alpha^2 = \beta^4$ , with  $\beta = \alpha^{(q+3)/8}$
- ▶ Standard: Compute  $\beta$ , conditionally multiply by  $\sqrt{-1}$  if  $\beta^2 = -\alpha$
- ▶ Decompression has  $\alpha = u/v$ , merge square root with inversion:

$$\beta = (u/v)^{(q+3)/8}$$

## Fast verification

- ▶ First part: point decompression, compute  $x$  coordinate  $x_R$  of  $R$  as

$$x_R = \pm \sqrt{(y_R^2 - 1)/(dy_R^2 + 1)}$$

- ▶ Looks like a square root and an inversion is required
- ▶ As  $q \equiv 5 \pmod{8}$  for each square  $\alpha$  we have  $\alpha^2 = \beta^4$ , with  $\beta = \alpha^{(q+3)/8}$
- ▶ Standard: Compute  $\beta$ , conditionally multiply by  $\sqrt{-1}$  if  $\beta^2 = -\alpha$
- ▶ Decompression has  $\alpha = u/v$ , merge square root with inversion:

$$\begin{aligned} \beta &= (u/v)^{(q+3)/8} = u^{(q+3)/8} v^{q-1-(q+3)/8} \\ &= u^{(q+3)/8} v^{(7q-11)/8} = uv^3 (uv^7)^{(q-5)/8}. \end{aligned}$$

## Fast verification

- ▶ First part: point decompression, compute  $x$  coordinate  $x_R$  of  $R$  as

$$x_R = \pm \sqrt{(y_R^2 - 1)/(dy_R^2 + 1)}$$

- ▶ Looks like a square root and an inversion is required
- ▶ As  $q \equiv 5 \pmod{8}$  for each square  $\alpha$  we have  $\alpha^2 = \beta^4$ , with  $\beta = \alpha^{(q+3)/8}$
- ▶ Standard: Compute  $\beta$ , conditionally multiply by  $\sqrt{-1}$  if  $\beta^2 = -\alpha$
- ▶ Decompression has  $\alpha = u/v$ , merge square root with inversion:

$$\begin{aligned} \beta &= (u/v)^{(q+3)/8} = u^{(q+3)/8} v^{q-1-(q+3)/8} \\ &= u^{(q+3)/8} v^{(7q-11)/8} = uv^3 (uv^7)^{(q-5)/8}. \end{aligned}$$

- ▶ Second part: computation of  $SB - H(\underline{R}, \underline{A}, M)A$
- ▶ Double-scalar multiplication using signed sliding windows
- ▶ Different window sizes for  $B$  (compile time) and  $A$  (run time)

## Fast verification

- ▶ First part: point decompression, compute  $x$  coordinate  $x_R$  of  $R$  as

$$x_R = \pm \sqrt{(y_R^2 - 1)/(dy_R^2 + 1)}$$

- ▶ Looks like a square root and an inversion is required
- ▶ As  $q \equiv 5 \pmod{8}$  for each square  $\alpha$  we have  $\alpha^2 = \beta^4$ , with  $\beta = \alpha^{(q+3)/8}$
- ▶ Standard: Compute  $\beta$ , conditionally multiply by  $\sqrt{-1}$  if  $\beta^2 = -\alpha$
- ▶ Decompression has  $\alpha = u/v$ , merge square root with inversion:

$$\begin{aligned} \beta &= (u/v)^{(q+3)/8} = u^{(q+3)/8} v^{q-1-(q+3)/8} \\ &= u^{(q+3)/8} v^{(7q-11)/8} = uv^3 (uv^7)^{(q-5)/8}. \end{aligned}$$

- ▶ Second part: computation of  $SB - H(\underline{R}, \underline{A}, M)A$
- ▶ Double-scalar multiplication using signed sliding windows
- ▶ Different window sizes for  $B$  (compile time) and  $A$  (run time)
- ▶ Verification takes  $< 280,000$  cycles

## Faster batch verification

- ▶ Verify a batch of  $(M_i, A_i, R_i, S_i)$ , where  $(R_i, S_i)$  is the alleged signature of  $M_i$  under key  $A_i$



## Faster batch verification

- ▶ Verify a batch of  $(M_i, A_i, R_i, S_i)$ , where  $(R_i, S_i)$  is the alleged signature of  $M_i$  under key  $A_i$
- ▶ Choose independent uniform random 128-bit integers  $z_i$
- ▶ Compute  $H_i = H(\underline{R_i}, \underline{A_i}, M_i)$

## Faster batch verification

- ▶ Verify a batch of  $(M_i, A_i, R_i, S_i)$ , where  $(R_i, S_i)$  is the alleged signature of  $M_i$  under key  $A_i$
- ▶ Choose independent uniform random 128-bit integers  $z_i$
- ▶ Compute  $H_i = H(\underline{R}_i, \underline{A}_i, M_i)$
- ▶ Verify the equation

$$\left(-\sum_i z_i S_i \bmod \ell\right) B + \sum_i z_i R_i + \sum_i (z_i H_i \bmod \ell) A_i = 0$$

## Faster batch verification

- ▶ Verify a batch of  $(M_i, A_i, R_i, S_i)$ , where  $(R_i, S_i)$  is the alleged signature of  $M_i$  under key  $A_i$
- ▶ Choose independent uniform random 128-bit integers  $z_i$
- ▶ Compute  $H_i = H(\underline{R}_i, \underline{A}_i, M_i)$
- ▶ Verify the equation

$$\left(-\sum_i z_i S_i \bmod \ell\right) B + \sum_i z_i R_i + \sum_i (z_i H_i \bmod \ell) A_i = 0$$

- ▶ Use Bos-Coster algorithm for multi-scalar multiplication

## Faster batch verification

- ▶ Verify a batch of  $(M_i, A_i, R_i, S_i)$ , where  $(R_i, S_i)$  is the alleged signature of  $M_i$  under key  $A_i$
- ▶ Choose independent uniform random 128-bit integers  $z_i$
- ▶ Compute  $H_i = H(\underline{R}_i, \underline{A}_i, M_i)$
- ▶ Verify the equation

$$\left(-\sum_i z_i S_i \bmod \ell\right)B + \sum_i z_i R_i + \sum_i (z_i H_i \bmod \ell)A_i = 0$$

- ▶ Use Bos-Coster algorithm for multi-scalar multiplication
- ▶ Verifying a batch of 64 signatures takes 8.55 million cycles (134,000 cycles/signature)

# The Bos-Coster algorithm

- ▶ Computation of  $Q = \sum_1^n s_i P_i$

# The Bos-Coster algorithm

- ▶ Computation of  $Q = \sum_1^n s_i P_i$
- ▶ Idea: Assume  $s_1 > s_2 > \dots > s_n$ . Recursively compute  $Q = (s_1 - s_2)P_1 + s_2(P_1 + P_2) + s_3P_3 \dots + s_nP_n$
- ▶ Each step requires the two largest scalars, one scalar subtraction and one point addition
- ▶ Each step “eliminates” expected  $\log n$  scalar bits

# The Bos-Coster algorithm

- ▶ Computation of  $Q = \sum_1^n s_i P_i$
- ▶ Idea: Assume  $s_1 > s_2 > \dots > s_n$ . Recursively compute  $Q = (s_1 - s_2)P_1 + s_2(P_1 + P_2) + s_3P_3 \dots + s_nP_n$
- ▶ Each step requires the two largest scalars, one scalar subtraction and one point addition
- ▶ Each step “eliminates” expected  $\log n$  scalar bits
- ▶ Requires fast access to the two largest scalars: put scalars into a heap
- ▶ Crucial for good performance: fast heap implementation

# The Bos-Coster algorithm

- ▶ Computation of  $Q = \sum_1^n s_i P_i$
- ▶ Idea: Assume  $s_1 > s_2 > \dots > s_n$ . Recursively compute  $Q = (s_1 - s_2)P_1 + s_2(P_1 + P_2) + s_3P_3 \dots + s_nP_n$
- ▶ Each step requires the two largest scalars, one scalar subtraction and one point addition
- ▶ Each step “eliminates” expected  $\log n$  scalar bits
- ▶ Requires fast access to the two largest scalars: put scalars into a heap
- ▶ Crucial for good performance: fast heap implementation
- ▶ Typical heap root replacement: start at the root, swap down for a variable amount of times



# The Bos-Coster algorithm

- ▶ Computation of  $Q = \sum_1^n s_i P_i$
- ▶ Idea: Assume  $s_1 > s_2 > \dots > s_n$ . Recursively compute  $Q = (s_1 - s_2)P_1 + s_2(P_1 + P_2) + s_3P_3 \dots + s_nP_n$
- ▶ Each step requires the two largest scalars, one scalar subtraction and one point addition
- ▶ Each step “eliminates” expected  $\log n$  scalar bits
- ▶ Requires fast access to the two largest scalars: put scalars into a heap
- ▶ Crucial for good performance: fast heap implementation
- ▶ Typical heap root replacement: start at the root, swap down for a variable amount of times
- ▶ Floyd’s heap: swap down to the bottom, swap up for a variable amount of times, advantages:
  - ▶ Each swap-down step needs only one comparison (instead of two)
  - ▶ Swap-down loop is more friendly to branch predictors

# Results



## Results

- ▶ New fast and secure signature scheme
- ▶ (Slow) C and Python reference implementations
- ▶ Fast AMD64 assembly implementations
- ▶ All software in the public domain and included in eBATS
- ▶ Software to be included in the NaCl library
- ▶ Paper to be presented at CHES 2011

<http://ed25519.cr.yp.to/>

Questions?

