ECC 2010 — Redmond, USA

Faster Implementation of Pairings

Francisco Rodríguez-Henríquez CINVESTAV, IPN, Mexico City, Mexico

Joint work with:

Jean-Luc Beuchat Nicolas Brisebarre Jérémie Detrey Nicolas Estibals Jorge González-Díaz Emmanuel López-Trejo Luis Martínez-Ramos Shigeo Mitsunari Eiji Okamoto Tadanori Teruya LCIS, University of Tsukuba, Japan Arénaire, LIP, ÉNS Lyon, France Caramel, INRIA Nancy Grand-Est, France Caramel, INRIA Nancy Grand-Est, France CINVESTAV, IPN, Mexico City, Mexico Intel Guadalajara Design Center, Mexico CINVESTAV, IPN, Mexico City, Mexico Cybozu Labs, Inc., Tokyo, Japan LCIS, University of Tsukuba, Japan LCIS, University of Tsukuba, Japan Outline of the talk



2 Hardware accelerator for the Tate pairing over supersingular curves

3 Software accelerator for the Tate pairing over supersingular curves



- Let $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ be two additively-written cyclic groups of prime order $\#\mathbb{G}_1 = \#\mathbb{G}_2 = \ell$
- $(\mathbb{G}_{\tau}, \times)$, a multiplicatively-written cyclic group of order $\#\mathbb{G}_{\tau} = \ell$

- Let $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ be two additively-written cyclic groups of prime order $\#\mathbb{G}_1 = \#\mathbb{G}_2 = \ell$
- $(\mathbb{G}_{\tau}, \times)$, a multiplicatively-written cyclic group of order $\#\mathbb{G}_{\tau} = \ell$
- A non-degenerate bilinear pairing is a map

 $\hat{e}: \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_{\tau}$

- Let $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ be two additively-written cyclic groups of prime order $\#\mathbb{G}_1 = \#\mathbb{G}_2 = \ell$
- $(\mathbb{G}_{\tau}, \times)$, a multiplicatively-written cyclic group of order $\#\mathbb{G}_{\tau} = \ell$
- A non-degenerate bilinear pairing is a map

 $\hat{e}: \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_{\tau}$

that satisfies the following conditions:

▶ non-degeneracy: $\hat{e}(P, P) \neq 1_{\mathbb{G}_{\tau}}$ (equivalently $\hat{e}(P, P)$ generates \mathbb{G}_{τ})

- Let $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ be two additively-written cyclic groups of prime order $\#\mathbb{G}_1 = \#\mathbb{G}_2 = \ell$
- $(\mathbb{G}_{\tau}, \times)$, a multiplicatively-written cyclic group of order $\#\mathbb{G}_{\tau} = \ell$
- A non-degenerate bilinear pairing is a map

 $\hat{e}: \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_{\tau}$

- non-degeneracy: $\hat{e}(P, P) \neq 1_{\mathbb{G}_{\tau}}$ (equivalently $\hat{e}(P, P)$ generates \mathbb{G}_{τ})
- ► bilinearity: $\hat{e}(Q_1+Q_2, R) = \hat{e}(Q_1, R) \cdot \hat{e}(Q_2, R)$ $\hat{e}(Q, R_1+R_2) = \hat{e}(Q, R_1) \cdot \hat{e}(Q, R_2)$

- Let $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ be two additively-written cyclic groups of prime order $\#\mathbb{G}_1 = \#\mathbb{G}_2 = \ell$
- $(\mathbb{G}_{\tau}, \times)$, a multiplicatively-written cyclic group of order $\#\mathbb{G}_{\tau} = \ell$
- A non-degenerate bilinear pairing is a map

 $\hat{\mathbf{e}}: \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_{\tau}$

- ▶ non-degeneracy: $\hat{e}(P, P) \neq 1_{\mathbb{G}_{\tau}}$ (equivalently $\hat{e}(P, P)$ generates \mathbb{G}_{τ})
- ► bilinearity: $\hat{e}(Q_1+Q_2, R) = \hat{e}(Q_1, R) \cdot \hat{e}(Q_2, R)$ $\hat{e}(Q, R_1+R_2) = \hat{e}(Q, R_1) \cdot \hat{e}(Q, R_2)$
- computability: ê can be efficiently computed

- Let $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ be two additively-written cyclic groups of prime order $\#\mathbb{G}_1 = \#\mathbb{G}_2 = \ell$
- $(\mathbb{G}_{\tau}, \times)$, a multiplicatively-written cyclic group of order $\#\mathbb{G}_{\tau} = \ell$
- A non-degenerate bilinear pairing is a map

 $\hat{e}: \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_{\tau}$

- non-degeneracy: $\hat{e}(P, P) \neq 1_{\mathbb{G}_{\tau}}$ (equivalently $\hat{e}(P, P)$ generates \mathbb{G}_{τ})
- ▶ bilinearity: $\hat{e}(Q_1+Q_2, R) = \hat{e}(Q_1, R) \cdot \hat{e}(Q_2, R)$ $\hat{e}(Q, R_1+R_2) = \hat{e}(Q, R_1) \cdot \hat{e}(Q, R_2)$
- computability: ê can be efficiently computed
- Immediate property: for any two integers k_1 and k_2

$$\hat{e}(\underline{k_1}Q,\underline{k_2}R) = \hat{e}(Q,R)^{\underline{k_1}\underline{k_2}}$$

- Let $(\mathbb{G}_1, +)$, $(\mathbb{G}_2, +)$ be two additively-written cyclic groups of prime order $\#\mathbb{G}_1 = \#\mathbb{G}_2 = \ell$
- $(\mathbb{G}_{\tau}, \times)$, a multiplicatively-written cyclic group of order $\#\mathbb{G}_{\tau} = \ell$
- A non-degenerate bilinear pairing is a map

 $\hat{\mathbf{e}}: \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_{\tau}$

that satisfies the following conditions:

- non-degeneracy: $\hat{e}(P, P) \neq 1_{\mathbb{G}_{\tau}}$ (equivalently $\hat{e}(P, P)$ generates \mathbb{G}_{τ})
- ▶ bilinearity: $\hat{e}(Q_1+Q_2, R) = \hat{e}(Q_1, R) \cdot \hat{e}(Q_2, R)$ $\hat{e}(Q, R_1+R_2) = \hat{e}(Q, R_1) \cdot \hat{e}(Q, R_2)$
- computability: ê can be efficiently computed
- Immediate property: for any two integers k_1 and k_2

$$\hat{e}(\mathbf{k}_1 Q, \mathbf{k}_2 R) = \hat{e}(Q, R)^{\mathbf{k}_1 \mathbf{k}_2}$$

• When $\mathbb{G}_1 = \mathbb{G}_2$ we say that the pairing is symmetric, otherwise if $\mathbb{G}_1 \neq \mathbb{G}_2$, the pairing is asymmetric.

イロト 不得 トイヨト イヨト 二日

- At first, used to attack supersingular elliptic curves
 - Menezes-Okamoto-Vanstone and Frey-Rück attacks, 1993 and 1994

$$DLP_{G_1} <_P DLP_{G_{\tau}}$$

$$kP \longrightarrow \hat{e}(kP,P) = \hat{e}(P,P)^k$$

 \blacktriangleright for cryptographic applications, we will also require the DLP in \mathbb{G}_{τ} to be hard

- At first, used to attack supersingular elliptic curves
 - Menezes-Okamoto-Vanstone and Frey-Rück attacks, 1993 and 1994

$$\begin{array}{lll} \mathsf{DLP}_{\mathbb{G}_1} & <_\mathsf{P} & \mathsf{DLP}_{\mathbb{G}_\tau} \\ & & \mathsf{kP} & \longrightarrow & \hat{\mathsf{e}}(\mathsf{kP},\mathsf{P}) = \hat{\mathsf{e}}(\mathsf{P},\mathsf{P})^{\mathsf{k}} \end{array}$$

- \blacktriangleright for cryptographic applications, we will also require the DLP in \mathbb{G}_{τ} to be hard
- One-round three-party key agreement (Joux, 2000)

- At first, used to attack supersingular elliptic curves
 - Menezes-Okamoto-Vanstone and Frey-Rück attacks, 1993 and 1994

$$\mathsf{DLP}_{\mathbb{G}_1} <_{\mathsf{P}} \mathsf{DLP}_{\mathbb{G}_{\tau}}$$

 $kP \longrightarrow \hat{e}(kP, P) = \hat{e}(P, P)^k$

- \blacktriangleright for cryptographic applications, we will also require the DLP in \mathbb{G}_{τ} to be hard
- One-round three-party key agreement (Joux, 2000)
- Identity-based encryption
 - Boneh–Franklin, 2001
 - Sakai–Kasahara, 2001

- At first, used to attack supersingular elliptic curves
 - Menezes-Okamoto-Vanstone and Frey-Rück attacks, 1993 and 1994

$$\mathsf{DLP}_{\mathbb{G}_1} <_{\mathsf{P}} \mathsf{DLP}_{\mathbb{G}_{\tau}}$$

 $kP \longrightarrow \hat{e}(kP, P) = \hat{e}(P, P)^k$

- \blacktriangleright for cryptographic applications, we will also require the DLP in \mathbb{G}_{τ} to be hard
- One-round three-party key agreement (Joux, 2000)
- Identity-based encryption
 - Boneh–Franklin, 2001
 - Sakai–Kasahara, 2001
- Short digital signatures
 - Boneh–Lynn–Shacham, 2001
 - Zang–Safavi-Naini–Susilo, 2004

- We first define
 - \mathbb{F}_q , a finite field, with $q = 2^m$ or 3^m
 - *E*, an elliptic curve defined over \mathbb{F}_q
 - ℓ , a large prime factor of $\#E(\mathbb{F}_q)$

- We first define
 - \mathbb{F}_q , a finite field, with $q = 2^m$ or 3^m
 - *E*, an elliptic curve defined over \mathbb{F}_q
 - ℓ , a large prime factor of $\#E(\mathbb{F}_q)$
- $\mathbb{G}_1 = E(\mathbb{F}_q)[\ell]$, the \mathbb{F}_q -rational ℓ -torsion of E:

 $\mathbb{G}_1 = \{ P \in E(\mathbb{F}_q) \mid \ell P = \mathcal{O} \}$

- We first define
 - \mathbb{F}_q , a finite field, with $q = 2^m$ or 3^m
 - *E*, an elliptic curve defined over \mathbb{F}_q
 - ℓ , a large prime factor of $\#E(\mathbb{F}_q)$
- $\mathbb{G}_1 = E(\mathbb{F}_q)[\ell]$, the \mathbb{F}_q -rational ℓ -torsion of E:

 $\mathbb{G}_1 = \{ P \in E(\mathbb{F}_q) \mid \ell P = \mathcal{O} \}$

• $\mathbb{G}_{\tau} = \mu_{\ell}$, the group of ℓ -th roots of unity in $\mathbb{F}_{q^k}^{\times}$: $\mathbb{G}_{\tau} = \{ U \in \mathbb{F}_{q^k}^{\times} \mid U^{\ell} = 1 \}$

- We first define
 - \mathbb{F}_q , a finite field, with $q = 2^m$ or 3^m
 - *E*, an elliptic curve defined over \mathbb{F}_q
 - ℓ , a large prime factor of $\#E(\mathbb{F}_q)$
- $\mathbb{G}_1 = E(\mathbb{F}_q)[\ell]$, the \mathbb{F}_q -rational ℓ -torsion of E:

 $\mathbb{G}_1 = \{ P \in E(\mathbb{F}_q) \mid \ell P = \mathcal{O} \}$

• $\mathbb{G}_{\tau} = \mu_{\ell}$, the group of ℓ -th roots of unity in $\mathbb{F}_{q^k}^{\times}$:

$$\mathbb{G}_{\tau} = \{ U \in \mathbb{F}_{q^k}^{\times} \mid U^{\ell} = 1 \}$$

- k is the embedding degree, the smallest integer such that $\mu_{\ell} \subseteq \mathbb{F}_{a^k}^{\times}$
 - usually large for ordinary elliptic curves
 - bounded in the case of supersingular elliptic curves
 (4 in characteristic 2; 6 in characteristic 3; and 2 in characteristic > 3)

Security considerations for Symmetric Pairings

 $\hat{e}: E(\mathbb{F}_{p^m})[\ell] imes E(\mathbb{F}_{p^m})[\ell] o \mu_\ell \subseteq \mathbb{F}_{p^{km}}^{\times}$

• The discrete logarithm problem should be hard in both \mathbb{G}_1 and \mathbb{G}_{τ}

Security considerations for Symmetric Pairings

 $\hat{e}: E(\mathbb{F}_{p^m})[\ell] imes E(\mathbb{F}_{p^m})[\ell] o \mu_\ell \subseteq \mathbb{F}_{p^{km}}^{\times}$

• The discrete logarithm problem should be hard in both G_1 and $G_{ au}$

Base field (\mathbb{F}_{p^m})	F ₂ <i>m</i>	F ₃ ^{<i>m</i>}
Lower security $(\sim 2^{64})$	<i>m</i> = 239	<i>m</i> = 97
Medium security ($\sim 2^{80}$)	<i>m</i> = 373	<i>m</i> = 163
Higher security ($\sim 2^{128}$)	<i>m</i> = 1103	<i>m</i> = 503

- \mathbb{F}_{2^m} : simpler finite field arithmetic
- \mathbb{F}_{3^m} : smaller field extension

→ Ξ →

3

- Arithmetic over \mathbb{F}_{p^m} :
 - polynomial basis: $\mathbb{F}_{p^m} \cong \mathbb{F}_p[x]/(f(x))$
 - f(x), degree-*m* polynomial irreducible over \mathbb{F}_p

- Arithmetic over F_p^m:
 - polynomial basis: $\mathbb{F}_{p^m} \cong \mathbb{F}_p[x]/(f(x))$
 - f(x), degree-*m* polynomial irreducible over \mathbb{F}_p
- Arithmetic over $\mathbb{F}_{p^{km}}^{\times}$:
 - tower-field representation
 - only arithmetic over the underlying field \mathbb{F}_{p^m}

- Arithmetic over F_p^m:
 - polynomial basis: $\mathbb{F}_{p^m} \cong \mathbb{F}_p[x]/(f(x))$
 - f(x), degree-*m* polynomial irreducible over \mathbb{F}_p
- Arithmetic over $\mathbb{F}_{p^{km}}^{\times}$:
 - tower-field representation
 - only arithmetic over the underlying field \mathbb{F}_{p^m}
- Operations over F_p^m:
 - O(m) additions / subtractions
 - O(m) multiplications
 - O(m) Frobenius maps $(a \mapsto a^p, i.e.$ squarings or cubings)
 - 1 inversion

- Arithmetic over F_p^m:
 - polynomial basis: $\mathbb{F}_{p^m} \cong \mathbb{F}_p[x]/(f(x))$
 - f(x), degree-*m* polynomial irreducible over \mathbb{F}_p
- Arithmetic over $\mathbb{F}_{p^{km}}^{\times}$:
 - tower-field representation
 - only arithmetic over the underlying field \mathbb{F}_{p^m}
- Operations over F_p^m:
 - O(m) additions / subtractions
 - O(m) multiplications
 - O(m) Frobenius maps $(a \mapsto a^p, i.e.$ squarings or cubings)
 - 1 inversion
- A first idea: an all-in-one unified operator:
 - shared resources
 - scalable architecture

Motivations

- High speed is more important than low resources for some cryptographic applications
- Explore the other end of the area vs. time tradeoff:
 - faster but larger than the unified operator
 - what about the area-time product?

Motivations

- High speed is more important than low resources for some cryptographic applications
- Explore the other end of the area vs. time tradeoff:
 - faster but larger than the unified operator
 - what about the area-time product?
- Accelerate the computation by extracting as much parallelism as possible...
- ... Without increasing dramatically the resource requirements

The Tate pairing over E(𝔽_{p^m}) is computed in two main steps
 ê(P, Q)

• The Tate pairing over $E(\mathbb{F}_{p^m})$ is computed in two main steps $\hat{e}(P, Q) = \eta_T(P, Q)$

• The Tate pairing over $E(\mathbb{F}_{p^m})$ is computed in two main steps $\hat{e}(P, Q) = \eta_T(P, Q)^M$

- The Tate pairing over $E(\mathbb{F}_{p^m})$ is computed in two main steps $\hat{e}(P, Q) = \eta_T(P, Q)^M$
- Computation of the η_T pairing
 - via Miller's algorithm: loop of (m+1)/2 iterations
 - ▶ result only defined modulo *N*-th powers in $\mathbb{F}_{p^{km}}^{\times}$, with $N = \#E(\mathbb{F}_{p^m})$

- The Tate pairing over $E(\mathbb{F}_{p^m})$ is computed in two main steps $\hat{e}(P, Q) = \eta_T(P, Q)^M$
- Computation of the η_T pairing
 - via Miller's algorithm: loop of (m+1)/2 iterations
 - ▶ result only defined modulo *N*-th powers in $\mathbb{F}_{p^{km}}^{\times}$, with $N = \#E(\mathbb{F}_{p^m})$
- Final exponentiation by $M = (p^{km} 1)/N$
 - required to obtain a unique value for each congruence class
 - example in characteristic 3 (k = 6 and $N = 3^m + 1 \pm 3^{(m+1)/2}$):

$$M = \frac{3^{6m} - 1}{3^m + 1 \pm 3^{(m+1)/2}} = (3^{3m} - 1)(3^m + 1)(3^m + 1 \mp 3^{(m+1)/2})$$

exploit the special form of the exponent: ad-hoc algorithm

- The Tate pairing over $E(\mathbb{F}_{p^m})$ is computed in two main steps $\hat{e}(P, Q) = \eta_T(P, Q)^M$
- Computation of the η_T pairing
 - via Miller's algorithm: loop of (m+1)/2 iterations
 - ▶ result only defined modulo *N*-th powers in $\mathbb{F}_{p^{km}}^{\times}$, with $N = \#E(\mathbb{F}_{p^m})$
- Final exponentiation by $M = (p^{km} 1)/N$
 - required to obtain a unique value for each congruence class
 - example in characteristic 3 (k = 6 and $N = 3^m + 1 \pm 3^{(m+1)/2}$):

$$M = \frac{3^{6m} - 1}{3^m + 1 \pm 3^{(m+1)/2}} = (3^{3m} - 1)(3^m + 1)(3^m + 1 \mp 3^{(m+1)/2})$$

- exploit the special form of the exponent: ad-hoc algorithm
- Two distinct computational requirements

- The Tate pairing over $E(\mathbb{F}_{p^m})$ is computed in two main steps $\hat{e}(P, Q) = \eta_T(P, Q)^M$
- Computation of the η_T pairing
 - via Miller's algorithm: loop of (m+1)/2 iterations
 - ▶ result only defined modulo *N*-th powers in $\mathbb{F}_{p^{km}}^{\times}$, with $N = \#E(\mathbb{F}_{p^m})$
- Final exponentiation by $M = (p^{km} 1)/N$
 - required to obtain a unique value for each congruence class
 - example in characteristic 3 (k = 6 and $N = 3^m + 1 \pm 3^{(m+1)/2}$):

$$M = \frac{3^{6m} - 1}{3^m + 1 \pm 3^{(m+1)/2}} = (3^{3m} - 1)(3^m + 1)(3^m + 1 \mp 3^{(m+1)/2})$$

- exploit the special form of the exponent: *ad-hoc* algorithm
- Two distinct computational requirements ⇒ use two distinct coprocessors

A (1) × (2) × (3) × (4)

Reduced Tate pairing

Reduced Tate pairing

< 🗇 🕨

э

Reduced Tate pairing



• Input: two points P and Q in $E(\mathbb{F}_{3^m})[\ell]$

э

Reduced Tate pairing



- Input: two points P and Q in $E(\mathbb{F}_{3^m})[\ell]$
- Output: an ℓ -th root of unity in the extension $\mathbb{F}_{3^{6m}}^{\times}$


- Input: two points P and Q in $E(\mathbb{F}_{3^m})[\ell]$
- Output: an ℓ -th root of unity in the extension $\mathbb{F}_{3^{6m}}^{\times}$
- Two very different steps



- Input: two points P and Q in $E(\mathbb{F}_{3^m})[\ell]$
- Output: an ℓ -th root of unity in the extension $\mathbb{F}_{3^{6m}}^{\times}$
- Two very different steps



- Input: two points P and Q in $E(\mathbb{F}_{3^m})[\ell]$
- Output: an ℓ -th root of unity in the extension $\mathbb{F}_{26m}^{\times}$
- Two very different steps

э



- Input: two points P and Q in $E(\mathbb{F}_{3^m})[\ell]$
- Output: an ℓ -th root of unity in the extension $\mathbb{F}_{3^{6m}}^{\times}$
- Two very different steps

- The two operations are purely sequential
- Only one active coprocessor at every moment

- The two operations are purely sequential
- Only one active coprocessor at every moment
- Pipeline the data between the two coprocessors

- The two operations are purely sequential
- Only one active coprocessor at every moment
- Pipeline the data between the two coprocessors
 - both of them are kept busy
 - higher throughput

- The two operations are purely sequential
- Only one active coprocessor at every moment
- Pipeline the data between the two coprocessors
 - both of them are kept busy
 - higher throughput
- Balance the computation time between the two coprocessors

η_{T} pairing algorithm

$$\eta_T: E(\mathbb{F}_{3^m})[\ell] \times E(\mathbb{F}_{3^m})[\ell] \to \mathbb{F}_{3^{6m}}^{\times}$$

- Three tasks per iteration:
 - 1 update the coordinates
 - ② compute the line equation
 - ③ accumulate the new factor
- Total cost: 17 ×, 4 Frobenius/inverse Frobenius and 30 + over \mathbb{F}_{3^m}

η_{T} pairing algorithm

$$\eta_T: E(\mathbb{F}_{3^m})[\ell] \times E(\mathbb{F}_{3^m})[\ell] \to \mathbb{F}_{3^{6m}}^{\times}$$

- Three tasks per iteration:
 - ① update the coordinates
 - ② compute the line equation
 - ③ accumulate the new factor
- Total cost: 17 \times , 4 Frobenius/inverse Frobenius and 30 + over \mathbb{F}_{3^m}
- Cost of the inverse Frobenius: Same as the Frobenius

for $i \leftarrow 0$ to (m-1)/2 do

$$\begin{array}{c} \textcircled{0} \quad \begin{array}{l} x_{P} \leftarrow \sqrt[3]{X_{P}} \quad ; \ y_{P} \leftarrow \sqrt[3]{Y_{P}} \\ x_{Q} \leftarrow x_{Q}^{3} \quad ; \ y_{Q} \leftarrow y_{Q}^{3} \end{array} & \begin{array}{l} 2 \text{ inv. Frobenius} \\ 2 \text{ Frobenius} \end{array} & (\mathbb{F}_{3^{m}}) \\ \hline \end{array} \\ \hline \\ \textcircled{0} \quad \begin{array}{l} t \leftarrow x_{P} + x_{Q} \quad ; \ u \leftarrow y_{P}y_{Q} \\ S \leftarrow -t^{2} \pm u\sigma - t\rho - \rho^{2} \end{array} & \begin{array}{l} 2 \times, 1 + (\mathbb{F}_{3^{m}}) \\ 15 \times, 29 + (\mathbb{F}_{3^{m}}) \end{array} \\ \hline \end{array}$$

end for

10

- Total cost: $17 \times$, 2 Frobenius and inverse Frobenius and $30 + \text{over } \mathbb{F}_{3^m}$ per iteration
 - ► Frobenius/inverse Frobenius and +: cheap and fast operations

- Total cost: 17 \times , 2 Frobenius and inverse Frobenius and 30 + over \mathbb{F}_{3^m} per iteration
 - ► Frobenius/inverse Frobenius and +: cheap and fast operations
 - critical operation: ×
- Need for a fast parallel multiplier: Karatsuba

- Total cost: 17 \times , 2 Frobenius and inverse Frobenius and 30 + over \mathbb{F}_{3^m} per iteration
 - Frobenius/inverse Frobenius and +: cheap and fast operations
 - critical operation: ×
- Need for a fast parallel multiplier: Karatsuba



 $A^{H}B^{L} + A^{L}B^{H} = (A^{H} + A^{L})(B^{H} + B^{L}) - A^{H}B^{H} - A^{L}B^{L}$



- fully parallel: all sub-products are computed in parallel
- pipelined architecture: higher clock frequency, one product per cycle



- fully parallel: all sub-products are computed in parallel
- pipelined architecture: higher clock frequency, one product per cycle
- sub-products recursively implemented as Karatsuba-Ofman multipliers



- fully parallel: all sub-products are computed in parallel
- pipelined architecture: higher clock frequency, one product per cycle
- sub-products recursively implemented as Karatsuba-Ofman multipliers
- support for other variants: odd-even split, 3-way split, ...



- fully parallel: all sub-products are computed in parallel
- pipelined architecture: higher clock frequency, one product per cycle
- sub-products recursively implemented as Karatsuba-Ofman multipliers
- support for other variants: odd-even split, 3-way split, ...
- final reduction modulo the irreducible polynomial f.

• η_T coprocessor based on a single large multiplier:

- parallel Karatsuba architecture
- 7-stage pipeline
- one product per cycle

• η_T coprocessor based on a single large multiplier:

- parallel Karatsuba architecture
- 7-stage pipeline
- one product per cycle
- Challenge: keep the multiplier busy at all times

• η_T coprocessor based on a single large multiplier:

- parallel Karatsuba architecture
- 7-stage pipeline
- one product per cycle
- Challenge: keep the multiplier busy at all times
- Careful scheduling to avoid pipeline bubbles (idle cycles):
 - ensure that multiplication operands are always available
 - avoid memory congestion issues
- We managed to accomplish that: our processor computes Miller loop in just $17 \cdot (m+3)/2$ clock cycles (considering the initialization phase)

A parallel operator for the η_{T} pairing



(日) (同) (三) (三)

• Compute $\hat{e}(P, Q)$ as $\eta_T(P, Q)^M$ with $\eta_T(P, Q) \in \mathbb{F}_{3^{6m}}^{\times}$ and

$$M = (3^{3m} - 1) (3^m + 1) (3^m + 1 \mp 3^{(m+1)/2})$$

• Compute $\hat{e}(P, Q)$ as $\eta_T(P, Q)^M$ with $\eta_T(P, Q) \in \mathbb{F}_{3^{6m}}^{\times}$ and

$$M = (3^{3m} - 1) (3^m + 1) (3^m + 1 \mp 3^{(m+1)/2})$$

Operations over F_{3^m}: 73 ×, 3m + 3 Frobenius, 3m + 175 +, and 1 inversion

• Compute $\hat{e}(P, Q)$ as $\eta_T(P, Q)^M$ with $\eta_T(P, Q) \in \mathbb{F}_{3^{6m}}^{\times}$ and

$$M = (3^{3m} - 1) (3^m + 1) (3^m + 1 \mp 3^{(m+1)/2})$$

Operations over F_{3^m}: 73 ×, 3m + 3 Frobenius, 3m + 175 +, and 1 inversion (~ log m × and m − 1 Frobenius)

• Compute $\hat{e}(P, Q)$ as $\eta_T(P, Q)^M$ with $\eta_T(P, Q) \in \mathbb{F}_{3^{6m}}^{\times}$ and

$$M = (3^{3m} - 1) (3^m + 1) (3^m + 1 \mp 3^{(m+1)/2})$$

- Operations over F_{3^m}: 73 ×, 3m + 3 Frobenius, 3m + 175 +, and 1 inversion (~ log m × and m − 1 Frobenius)
- Cost of the η_T pairing:
 - ► (m+1)/2 iterations
 - ▶ 17 ×, 10 Frobenius and 30 + over \mathbb{F}_{3^m} per iteration

• Compute $\hat{e}(P, Q)$ as $\eta_T(P, Q)^M$ with $\eta_T(P, Q) \in \mathbb{F}_{3^{6m}}^{\times}$ and

$$M = (3^{3m} - 1) (3^m + 1) (3^m + 1 \mp 3^{(m+1)/2})$$

- Operations over F_{3^m}: 73 ×, 3m + 3 Frobenius, 3m + 175 +, and 1 inversion (~ log m × and m − 1 Frobenius)
- Cost of the η_T pairing:
 - ► (m+1)/2 iterations
 - ▶ 17 ×, 10 Frobenius and 30 + over \mathbb{F}_{3^m} per iteration
- The final exponentiation is much cheaper than the η_T pairing
- Challenge for the final exponentiation:
 - computation in the same time as the η_T pairing
 - using as few resources as possible

- Design the smallest architecture possible supporting all the required operations over F₃^m
- purely sequential scheduling

- Design the smallest architecture possible supporting all the required operations over F₃^m
- purely sequential scheduling
- Although some parallelism is required.

- Design the smallest architecture possible supporting all the required operations over F₃^m
- purely sequential scheduling
- Although some parallelism is required.
- We found out that the usage of the inverse Frobenius operator is advantageous for computing the final exponentiation (as long as the irreducible polynomials are inverse-Frobenius friendly)

- Design the smallest architecture possible supporting all the required operations over F₃^m
- purely sequential scheduling
- Although some parallelism is required.
- We found out that the usage of the inverse Frobenius operator is advantageous for computing the final exponentiation (as long as the irreducible polynomials are inverse-Frobenius friendly)
- New coprocessor with two arithmetic units:
 - a standalone multiplier, based on a parallel-serial scheme
 - a unified operator supporting addition/subtraction, inverse Frobenius map and inverse double Frobenius map

A coprocessor for the final exponentiation



3. 3

Agenda

Context

- Hardware accelerator for the Tate pairing over supersingular curves
 Implementation Results in Hardware
- **3** Software accelerator for the Tate pairing over supersingular curves
 - Computing the non-reduced pairing
 - Final exponentiation
 - Implementation results
- Optimal Ate Pairing over Barreto-Naehrig Curves
 Barreto-Naehrig Curves

Hardware accelerators



1 / 49)

э

Hardware implementation notes

• Our Xilinx FPGA implementation, significantly improved the computation time of all the hardware pairing coprocessors for supersingular curves previously published

Hardware implementation notes

- Our Xilinx FPGA implementation, significantly improved the computation time of all the hardware pairing coprocessors for supersingular curves previously published
- (a bit Surprisingly) our architecture also enjoys the best area/time trade-off performance among supersingular pairing accelerators

Hardware implementation notes

- Our Xilinx FPGA implementation, significantly improved the computation time of all the hardware pairing coprocessors for supersingular curves previously published
- (a bit Surprisingly) our architecture also enjoys the best area/time trade-off performance among supersingular pairing accelerators
- However, because we exceeded the FPGA's capacity, we could only achieve up to 109 bits of security
Hardware implementation notes

- Our Xilinx FPGA implementation, significantly improved the computation time of all the hardware pairing coprocessors for supersingular curves previously published
- (a bit Surprisingly) our architecture also enjoys the best area/time trade-off performance among supersingular pairing accelerators
- However, because we exceeded the FPGA's capacity, we could only achieve up to 109 bits of security
- Although it was not discussed here, we also implemented the Tate pairing over char 2. Experimentally, we observed that our char 2 and char 3 accelerators achieve almost the same time performance

Hardware implementation notes

- Our Xilinx FPGA implementation, significantly improved the computation time of all the hardware pairing coprocessors for supersingular curves previously published
- (a bit Surprisingly) our architecture also enjoys the best area/time trade-off performance among supersingular pairing accelerators
- However, because we exceeded the FPGA's capacity, we could only achieve up to 109 bits of security
- Although it was not discussed here, we also implemented the Tate pairing over char 2. Experimentally, we observed that our char 2 and char 3 accelerators achieve almost the same time performance
- In the design process of our char 2 accelerator we found the following undocumented family of square-root friendly irreducible pentanomials: $f(x) = x^m + x^{m-d} + x^{m-2d} + x^d + 1.$
- all technical details of these designs can be found in the preprint manuscripts eprint 2009/122 and eprint 2009/398

くほと くほと くほと

/ 49)

Agenda

Context

Hardware accelerator for the Tate pairing over supersingular curves
 Implementation Results in Hardware

Software accelerator for the Tate pairing over supersingular curves
 Computing the non-reduced pairing

- Final exponentiation
- Implementation results
- Optimal Ate Pairing over Barreto-Naehrig Curves
 Barreto-Naehrig Curves

• η_T pairing: shorter loop

for $i \leftarrow 0$ to (m-1)/2 do

end for

3

- η_T pairing: shorter loop
- Based on Miller's algorithm:

for $i \leftarrow 0$ to (m-1)/2 do $x_P \leftarrow \sqrt[3]{x_P}$; $y_P \leftarrow \sqrt[3]{y_P}$ $x_Q \leftarrow x_Q^3$; $y_Q \leftarrow y_Q^3$ $t \leftarrow x_P + x_Q$ $u \leftarrow y_P y_Q$ $S \leftarrow -t^2 \pm u\sigma - t\rho - \rho^2$ $R \leftarrow R \cdot S$

- η_T pairing: shorter loop
- Based on Miller's algorithm:
 - 1 update of point coordinates

for $i \leftarrow 0$ to (m-1)/2 do

Francisco Rodríguez-Henríquez

- η_T pairing: shorter loop
- Based on Miller's algorithm:
 - 1 update of point coordinates
 - 2 computation of line equation

for $i \leftarrow 0$ to (m-1)/2 do (1) $\begin{array}{c} x_P \leftarrow \sqrt[3]{x_P} & ; y_P \leftarrow \sqrt[3]{y_P} \\ x_Q \leftarrow x_Q^3 & ; y_Q \leftarrow y_Q^3 \end{array} \begin{array}{c} 2 \sqrt[3]{\cdot} \\ 2 (\cdot)^3 \end{array}$ (2) $\begin{array}{c} t \leftarrow x_P + x_Q & ; u \leftarrow y_P y_Q \\ S \leftarrow -t^2 \pm u\sigma - t\rho - \rho^2 \end{array} \begin{array}{c} 2 \times , 2 + \\ R \leftarrow R \cdot S \end{array}$

- η_T pairing: shorter loop
- Based on Miller's algorithm:
 - 1 update of point coordinates
 - 2 computation of line equation
 - ③ accumulation of the new factor

for $i \leftarrow 0$ to (m-1)/2 do

- η_T pairing: shorter loop
- Based on Miller's algorithm:
 - 1 update of point coordinates
 - 2 computation of line equation
 - ③ accumulation of the new factor
- Multiplication is critical
- Comb right-to-left multiplier over F_{3^m}

for $i \leftarrow 0$ to (m-1)/2 do

- η_T pairing: shorter loop
- Based on Miller's algorithm:
 - 1 update of point coordinates
 - 2 computation of line equation
 - ③ accumulation of the new factor
- Multiplication is critical
- Comb right-to-left multiplier over \mathbb{F}_{3^m}
- Sparse multiplication over $\mathbb{F}_{3^{6m}}$

for $i \leftarrow 0$ to (m-1)/2 do

- η_T pairing: shorter loop
- Based on Miller's algorithm:
 - 1 update of point coordinates
 - 2 computation of line equation
 - ③ accumulation of the new factor
- Multiplication is critical
- Comb right-to-left multiplier over \mathbb{F}_{3^m}
- Sparse multiplication over F_{36m}
 - $15 \times \text{ and } 29 + \text{ over } \mathbb{F}_{3^m}$ (Beuchat *et al.*, ARITH 18)

for $i \leftarrow 0$ to (m-1)/2 do

1)
$$\begin{array}{c} x_{P} \leftarrow \sqrt[3]{x_{P}} & ; y_{P} \leftarrow \sqrt[3]{y_{P}} & 2 \sqrt[3]{\cdot} \\ x_{Q} \leftarrow x_{Q}^{3} & ; y_{Q} \leftarrow y_{Q}^{3} & 2 (\cdot)^{3} \end{array} \\ \\ \begin{array}{c} t \leftarrow x_{P} + x_{Q} & ; u \leftarrow y_{P}y_{Q} \\ S \leftarrow -t^{2} \pm u\sigma - t\rho - \rho^{2} \end{array} & 2 \times , 2 + \\ \\ \begin{array}{c} 3 \\ R \leftarrow R \cdot S \end{array} & 15 \times , 29 + \end{array} \end{array}$$

- η_T pairing: shorter loop
- Based on Miller's algorithm:
 - 1 update of point coordinates
 - 2 computation of line equation
 - ③ accumulation of the new factor
- Multiplication is critical
- Comb right-to-left multiplier over \mathbb{F}_{3^m}
- Sparse multiplication over F_{36m}
 - $15 \times \text{ and } 29 + \text{ over } \mathbb{F}_{3^m}$ (Beuchat *et al.*, ARITH 18)
 - $12 \times \text{ and } 59 + \text{ over } \mathbb{F}_{3^m}$ (Gorla *et al.*, SAC 2007)

for $i \leftarrow 0$ to (m-1)/2 do



3



3



3

イロト イポト イヨト イヨト



3

イロト 不得下 イヨト イヨト



3

Agenda

Context

Hardware accelerator for the Tate pairing over supersingular curvesImplementation Results in Hardware

Software accelerator for the Tate pairing over supersingular curves
 Computing the non-reduced pairing

- Final exponentiation
- Implementation results

Optimal Ate Pairing over Barreto-Naehrig Curves
 Barreto-Naehrig Curves

• Final exponentiation consists of raising $\hat{e}(P, Q)$ to the exponent,

$$M = \frac{2^{4m} - 1}{N} = (2^{2m} - 1) \cdot (2^m + 1 - \nu 2^{(m+1)/2}),$$

where $\nu = (-1)^b$ when $m \equiv 1, 7 \pmod{8}$ and $\nu = (-1)^{1-b}$ in all other cases.

Highly sequential computation, Very heterogeneous

• Final exponentiation consists of raising $\hat{e}(P, Q)$ to the exponent,

$$M = \frac{2^{4m} - 1}{N} = (2^{2m} - 1) \cdot (2^m + 1 - \nu 2^{(m+1)/2}),$$

where $\nu = (-1)^b$ when $m \equiv 1, 7 \pmod{8}$ and $\nu = (-1)^{1-b}$ in all other cases.

- Highly sequential computation, Very heterogeneous
- We perform this operation according to a slightly optimized version:
 - ► Raising to the (2^m + 1)-th power. Raising the outcome of Miller's algorithm to the (2^{2m} 1)-th power produces an element U ∈ F_{2^{4m}} of order 2^{2m} + 1. This property allows one to save a multiplication over F_{2^{4m}} when raising U to the (2^m + 1)-th power.

• Final exponentiation consists of raising $\hat{e}(P, Q)$ to the exponent,

$$M = \frac{2^{4m} - 1}{N} = (2^{2m} - 1) \cdot (2^m + 1 - \nu 2^{(m+1)/2}),$$

where $\nu = (-1)^b$ when $m \equiv 1, 7 \pmod{8}$ and $\nu = (-1)^{1-b}$ in all other cases.

- Highly sequential computation, Very heterogeneous
- We perform this operation according to a slightly optimized version:
 - ► Raising to the (2^m + 1)-th power. Raising the outcome of Miller's algorithm to the (2^{2m} 1)-th power produces an element U ∈ F_{2^{4m}} of order 2^{2m} + 1. This property allows one to save a multiplication over F_{2^{4m}} when raising U to the (2^m + 1)-th power.
 - ► Raising to the 2^{m+1}/₂-th power. raising an element of F_{24m} to the 2ⁱ-th power involves 4i squarings and at most four additions over F_{2m}

• Final exponentiation consists of raising $\hat{e}(P, Q)$ to the exponent,

$$M = \frac{2^{4m} - 1}{N} = (2^{2m} - 1) \cdot (2^m + 1 - \nu 2^{(m+1)/2}),$$

where $\nu = (-1)^b$ when $m \equiv 1, 7 \pmod{8}$ and $\nu = (-1)^{1-b}$ in all other cases.

- Highly sequential computation, Very heterogeneous
- We perform this operation according to a slightly optimized version:
 - ► Raising to the (2^m + 1)-th power. Raising the outcome of Miller's algorithm to the (2^{2m} 1)-th power produces an element U ∈ F_{2^{4m}} of order 2^{2m} + 1. This property allows one to save a multiplication over F_{2^{4m}} when raising U to the (2^m + 1)-th power.
 - ► Raising to the 2^{m+1}/₂-th power. raising an element of F_{24m} to the 2ⁱ-th power involves 4i squarings and at most four additions over F_{2m}

Finite field arithmetic

• Target: multi-core architectures

-

э

Finite field arithmetic

- Target: multi-core architectures
- Arithmetic over \mathbb{F}_{2^m} and \mathbb{F}_{3^m} : SSE instruction set

Finite field arithmetic

- Target: multi-core architectures
- Arithmetic over \mathbb{F}_{2^m} and \mathbb{F}_{3^m} : SSE instruction set
- Timings are given in clock cycles and were measured on an Intel Core 2 processor working at 2.4 GHz.

	Field	xp	√x	Mult
Aranha et al. CT-RSA'10	F ₂₁₂₂₃	160	166	4030
Our work CANS'10	$F_{2^{1223}}$	480	749	5438
Our work CAINS 10	F ₃₅₀₉	900	974	4128

Agenda

Context

Hardware accelerator for the Tate pairing over supersingular curvesImplementation Results in Hardware

Software accelerator for the Tate pairing over supersingular curves
 Computing the non-reduced pairing
 Final exponentiation

Implementation results

Optimal Ate Pairing over Barreto-Naehrig Curves
 Barreto-Naehrig Curves

Implementation results

- Timings achieved on an Intel Core2 are given in millions of clock cycles
- Windows XP 64-bit SP2 environment

Implementation results

- Timings achieved on an Intel Core2 are given in millions of clock cycles
- Windows XP 64-bit SP2 environment

	Curve	Security [bits]	# of cores	Freq. [GHz]	Calc. time [Mcycles]	
	$E(\mathbb{F}_{2^{1223}})$	128	1	2.4	18.76	
Aranha et al.	$E(\mathbb{F}_{2^{1223}})$	128	2	2.4	10.08	
CT-RSA'10	$E(\mathbb{F}_{2^{1223}})$	128	4	2.4	5.72	
O	$E(\mathbb{F}_{3^{509}})$	128	1	2.4	18.2	
CANS'10	$E(\mathbb{F}_{3^{509}})$	128	2	2.4	10.34	
	$E(\mathbb{F}_{3^{509}})$	128	4	2.4	7.06	

• Significantly faster implementation (for a while)

- Significantly faster implementation (for a while)
- How many cores?

- Significantly faster implementation (for a while)
- How many cores?
 - acceleration always less than the ideal speedup factor
 - best choice: dual-core implementation

- Significantly faster implementation (for a while)
- How many cores?
 - acceleration always less than the ideal speedup factor
 - best choice: dual-core implementation
- Characteristic 3 performs better than characteristic 2

- Significantly faster implementation (for a while)
- How many cores?
 - acceleration always less than the ideal speedup factor
 - best choice: dual-core implementation
- Characteristic 3 performs better than characteristic 2
 - at least on Intel Core2 and Intel Core i7

- Significantly faster implementation (for a while)
- How many cores?
 - acceleration always less than the ideal speedup factor
 - best choice: dual-core implementation
- Characteristic 3 performs better than characteristic 2
 - at least on Intel Core2 and Intel Core i7
 - next generation of processors: built-in carry-less 64-bit multiplier
 - the battle is not over!

Agenda

Context

- Hardware accelerator for the Tate pairing over supersingular curves
 Implementation Results in Hardware
- Software accelerator for the Tate pairing over supersingular curves
 Computing the non-reduced pairing
 Final exponentiation
 - Implementation results

Optimal Ate Pairing over Barreto-Naehrig Curves Barreto-Naehrig Curves

Barreto-Naehrig Curves

Defined by the equation $E: y^2 = x^3 + b$, where $b \neq 0$. Their embedding degree k is equal to 12. The characteristic p of the prime field, the group order r, and the trace of Frobenius t_r of the curve are parametrized as follows:

$$p(t) = 36t^{4} + 36t^{3} + 24t^{2} + 6t + 1,$$

$$r(t) = 36t^{4} + 36t^{3} + 18t^{2} + 6t + 1,$$

$$t_{r}(t) = 6t^{2} + 1,$$
(1)

where $t \in \mathbb{Z}$ is an arbitrary integer such that p = p(t) and r = r(t) are both prime numbers.

For efficiency purposes, t must have a low Hamming weight. In this work we used,

$$t = 2^{62} - 2^{54} + 2^{44}$$
Barreto-Naehrig Curves

Let E[r] denote the *r*-torsion subgroup of *E* and π_p be the Frobenius endomorphism $\pi_p: E \to E$ given by $\pi_p(x, y) = (x^p, y^p)$. We define,

- $\mathbb{G}_1 = E(\mathbb{F}_p)[r]$,
- $\mathbb{G}_2 \subseteq E(\mathbb{F}_{p^{12}})[r],$
- $\mathbb{G}_{\tau} = \mu_r \subset \mathbb{F}_{p^{12}}^*$ (*i.e.* the group of *r*-th roots of unity).
- The optimal ate pairing on the BN curve E is given as,

$$\begin{array}{rcl} a_{\mathrm{opt}} : \mathbb{G}_{2} \times \mathbb{G}_{1} & \longrightarrow & \mathbb{G}_{3} \\ & (Q, P) & \longmapsto & \left(f_{6t+2, Q}(P) \cdot I_{[6t+2]Q, \pi_{p}(Q)}(P) \cdot \right. \\ & & \left. I_{[6t+2]Q + \pi_{p}(Q), -\pi_{p}^{2}(Q)}(P) \right)^{\frac{p^{12}-1}{r}}, \end{array}$$

 In practice, pairing computations can be restricted to points P and Q' that belong to E(𝔽_p) and E'(𝔽_{p²}), respectively, where, E'/𝔽_{p²}: y² = x³ + b/ξ.

Optimal ate pairing algorithm

Input: $P \in \mathbb{G}_1 \vee Q \in \mathbb{G}_2$. **Output:** $a_{opt}(Q, P)$. 1. Write s = 6t + 2 as $s = \sum_{i=0}^{L-1} s_i 2^i$, where $s_i \in \{-1, 0, 1\}$; 2. $T \leftarrow Q$. $f \leftarrow 1$: 3 for i = 1 - 2 to 0 do 4. $f \leftarrow f^2 \cdot I_T \tau(P); T \leftarrow 2T;$ 5. **if** $s_i = -1$ then $f \leftarrow f \cdot I_{T-Q}(P); T \leftarrow T-Q;$ 6. 7. else if $s_i = 1$ then $f \leftarrow f \cdot I_T \circ (P); T \leftarrow T + Q;$ 8 end if 9 10. end for 11. $Q_1 \leftarrow \pi_p(Q); \quad Q_2 \leftarrow \pi_{p^2}(Q);$ 12. $f \leftarrow f \cdot I_{T,Q_1}(P)$; $T \leftarrow T + Q_1$; 13. $f \leftarrow f \cdot I_T = Q_2$; $T \leftarrow T - Q_2$; 14 $f \leftarrow f^{(p^{12}-1)/r}$. 15. **return** *f*: • • = • • = •

э.













 $f = g + hw \in \mathbb{F}_{p^{12}},$ with $g, h \in \mathbb{F}_{p^6}.$

but also

$$g = g_0 + g_1 v + g_2 v^2$$
,
 $h = h_0 + h_1 v + h_2 v^2$,
where $g_i, h_i \in \mathbb{F}_{p^2}$,
for $i = 1, 2, 3$.



hence, we can write $f \in \mathbb{F}_{p^{12}}$ as

$$F = g + hw = g + hw = g_0 + h_0W + g_1W^2 + h_1W^3 + g_2W^4 + h_2W^5$$

- Let (a, m, s, i), (ã, m, ŝ, i), and (A, M, S, I) denote the cost of field addition, multiplication, squaring, and inversion in F_p, F_{p²}, and F_{p⁶}, respectively
- we sometimes need to compute the multiplication in the base field by the constant coefficient $\beta \in \mathbb{F}_p$ of the irreducible binomial $f(u) = u^2 - \beta$. We refer to this operation as m_β
- we sometimes need to compute the multiplication of an arbitrary element in 𝔽_{p²} times the constant ξ = u ∈ 𝔽_p at a cost of one multiplication by the constant β. We refer to this operation as m_ξ, but it is noticed that the cost of m_ξ is essentially the same of that of m_β.

Computational costs of the tower extension field arithmetic

Field	Add/Sub	Mult	Squaring	Inversion
\mathbb{F}_{p^2}	ã = 2a	$\tilde{m} = 3m + 3a + m_{\beta}$	$\tilde{s}=2m+3a+m_{eta}$	$\widetilde{i} = 4m + m_{eta} + 2a + i$
\mathbb{F}_{p^6}	3ã	$6 ilde{m}+2m_{\xi}+15 ilde{a}$	$2\tilde{m}+3\tilde{s}+2m_{\xi}+8\tilde{a}$	$9 ilde{m}+3 ilde{s}+4m_{\xi}\+4 ilde{a}+ ilde{i}$
$\mathbb{F}_{p^{12}}$	6ã	$18 ilde{m} + 6 m_{\xi} + 60 ilde{a}$	$12\tilde{m} + 4m_{\xi} + 45\tilde{a}$	$\begin{array}{c} 25\tilde{m}+9\tilde{s}+12m_{\xi}\\ +61\tilde{a}+\tilde{i} \end{array}$
$\mathbb{G}_{\Phi_6}(\mathbb{F}_{p^2})$	6ã	$18\tilde{m}+6m_{\xi}+60\tilde{a}$	$9 ilde{s}+4m_{\xi} \ +30 ilde{a}$	Conjugate

• The bit-length of 6t + 2 is L = 65

э

- The bit-length of 6t + 2 is L = 65
 - This implies that we require 64 point doubling in the Miller loop.

- The bit-length of 6t + 2 is L = 65
 - This implies that we require 64 point doubling in the Miller loop.
- The Hamming weight of 6t + 2 is 7

- The bit-length of 6t + 2 is L = 65
 - This implies that we require 64 point doubling in the Miller loop.
- The Hamming weight of 6t + 2 is 7
 - This implies that we require 6 point addition/subtraction in the Miller loop.

- The bit-length of 6t + 2 is L = 65
 - This implies that we require 64 point doubling in the Miller loop.
- The Hamming weight of 6t + 2 is 7
 - This implies that we require 6 point addition/subtraction in the Miller loop.
- The low Hamming weight of *t* allows us to save arithmetic operations in the hard part of the final exponentiation

Mille Loop Cost

$$\begin{array}{llllllll} \text{Miller Loop} &=& 64 \cdot \left(28\tilde{m} + 8\tilde{s} + 100\tilde{a} + 4m + 6m_{\beta}\right) + \\ && 6 \cdot \left(20\tilde{m} + 7\tilde{s} + 64\tilde{a} + 4m + 2m_{\beta}\right) + \\ && 40\tilde{m} + 14\tilde{s} + 128\tilde{a} + 14m + 4m_{\beta} \\ &=& 1952\tilde{m} + 568\tilde{s} + 6912\tilde{a} + 294m + 400m_{\beta}. \end{array}$$

2

<ロ> (日) (日) (日) (日) (日)

$$e = rac{p^{12}-1}{r} = (p^6-1)\cdot(p^2+1)\cdotrac{p^4-p^2+1}{r}.$$

$$e = rac{p^{12}-1}{r} = (p^6-1)\cdot(p^2+1)\cdotrac{p^4-p^2+1}{r}.$$

- Raising to f^(p⁶-1) = f̄ ⋅ f⁻¹ costs one conjugation, one inversion and one multiplication over 𝔽_{p¹²}.
- After this step, f becomes an element of the cyclotomic group $\mathbb{G}_{\Phi_6}(\mathbb{F}_{p^2}).$

$$e = rac{p^{12}-1}{r} = (p^6-1)\cdot(p^2+1)\cdotrac{p^4-p^2+1}{r}.$$

- Raising to f^(p⁶-1) = f̄ ⋅ f⁻¹ costs one conjugation, one inversion and one multiplication over 𝔽_{p¹²}.
- After this step, f becomes an element of the cyclotomic group $\mathbb{G}_{\Phi_6}(\mathbb{F}_{p^2}).$
- Raising to the power $p^2 + 1$ costs 5 multiplications over \mathbb{F}_p , and one multiplication over $\mathbb{F}_{p^{12}}$.

$$e = rac{p^{12}-1}{r} = (p^6-1)\cdot(p^2+1)\cdotrac{p^4-p^2+1}{r}.$$

- Raising to f^(p⁶-1) = f̄ ⋅ f⁻¹ costs one conjugation, one inversion and one multiplication over 𝔽_{p¹²}.
- After this step, f becomes an element of the cyclotomic group $\mathbb{G}_{\Phi_6}(\mathbb{F}_{p^2}).$
- Raising to the power p² + 1 costs 5 multiplications over F_p, and one multiplication over F_{p¹²}.
- Raising to the power $m^{(p^4-p^2+1)/r}$ is referred as the hard part of the final exponentiation

Hard part of the final exponentiation

We used the addition chain proposed by Scott et al. at Pairing'09

 m^{t} , $m^{t^{2}}$, $m^{t^{3}}$, m^{p} , $m^{p^{2}}$, $m^{p^{3}}$, $m^{(tp)}$, $m^{(t^{2}p)}$, $m^{(t^{3}p)}$, $m^{(t^{2}p^{2})}$,

Hard part of the final exponentiation

We used the addition chain proposed by Scott et al. at Pairing'09

$$m^{t}$$
, $m^{t^{2}}$, $m^{t^{3}}$, m^{p} , $m^{p^{2}}$, $m^{p^{3}}$, $m^{(tp)}$, $m^{(t^{2}p)}$, $m^{(t^{3}p)}$, $m^{(t^{2}p^{2})}$,

• Taking advantage of the Frobenius, we can easily compute, m^p , m^{p^2} , m^{p^3} , $m^{(tp)}$, $m^{(t^2p)}$, $m^{(t^3p)}$, y $m^{(t^2p^2)}$ at a cost of 35 multiplications in the base field \mathbb{F}_p .

Hard part of the final exponentiation

We used the addition chain proposed by Scott et al. at Pairing'09

$$m^{t}$$
, $m^{t^{2}}$, $m^{t^{3}}$, m^{p} , $m^{p^{2}}$, $m^{p^{3}}$, $m^{(tp)}$, $m^{(t^{2}p)}$, $m^{(t^{3}p)}$, $m^{(t^{2}p^{2})}$,

- Taking advantage of the Frobenius, we can easily compute, m^p , m^{p^2} , m^{p^3} , $m^{(tp)}$, $m^{(t^2p)}$, $m^{(t^3p)}$, y $m^{(t^2p^2)}$ at a cost of 35 multiplications in the base field \mathbb{F}_p .
- The most costly part of this procedure consists on the computation of m^t , $m^{t^2} = (m^t)^t$, $m^{t^3} = (m^{t^2})^t$.
- Since t = 2⁶² 2⁵⁴ + 2⁴⁴, these exponentiations can be computed at a cost of 62 ⋅ 3 = 186 cyclotomic squarings plus 2 ⋅ 3 = 6 multiplications over F_{p¹²}.

Final exponentiation computational cost

Exp. Final =
$$(25\tilde{m} + 9\tilde{s} + 12m_{\beta} + 61\tilde{a} + \tilde{i}) + (18\tilde{m} + 6m_{\beta} + 60\tilde{a}) + (18\tilde{m} + 6m_{\beta} + 60\tilde{a}) + 10m + 13 \cdot (18\tilde{m} + 6m_{\beta} + 60\tilde{a}) + 4 \cdot (9\tilde{s} + 4m_{\beta} + 30\tilde{a}) + 70m + 186 \cdot (9\tilde{s} + 4m_{\beta} + 30\tilde{a}) + 6 \cdot (18\tilde{m} + 6m_{\beta} + 60\tilde{a}) = 403\tilde{m} + 1719\tilde{s} + 7021\tilde{a} + 80m + 898m_{\beta} + \tilde{i}.$$

E + 4 E +

э

A Comparison of arithmetic operations required by the computation of the ate pairing variants.

		ñ	ŝ	ã	ĩ	m_{ξ}
Hankerson et al. (5.). (Miller Loop	2277	356	6712	1	412
D ato poiring	Final Exp.	1616	1197	8977	1	1062
R-ate pairing	Total	3893	1553	15689	2	1474
Nachrig et al. (Miller Loop	2022	590	7140		410
Optimal ato pairing	Final Exp.	678	1719	7921	1	988
Optimal ate pairing	Total	2700	2309	15061	1	1398
This work and	Miller Loop	1954	568	6912		400
Optimal at a pairing 2010	Final Exp	443	1719	7021	1	898
Optimal ate pairing	Total	2397	2287	13933	1	1298

Library Implementation

- We use the mul operation included in the x86-64 instruction set. It multiplies two 64-bit unsigned integers in about 3 clock cycles on Intel Core i7 and AMD Opteron processors
- An element $x \in \mathbb{F}_p$ is represented as $x = (x_3, x_2, x_1, x_0)$, where $x_i, 0 \le i \le 3$, are 64-bit integers
- Multiplication and inversion over \mathbb{F}_p are accomplished according to the well-known Montgomery multiplication and Montgomery inversion algorithms, respectively
- The 256-bit integer multiplication and Montgomery reduction are computed in 55 and 100 clock cycles, respectively

Cycle counts of multiplication over \mathbb{F}_{p^2} , squaring over \mathbb{F}_{p^2} , and optimal ate pairing on different machines

	Our results				
	Core i7ª	Opteron ^b	Core 2 Duo ^c	Athlon 64 X2 ^d	
Multiplication over \mathbb{F}_{p^2}	435	443	558	473	
Squaring over \mathbb{F}_{p^2}	342	355	445	376	
Miller loop	1,330,000	1,360,000	1,680,000	1,480,000	
Final exponentiation	1,000,000	1,040,000	1,270,000	1,150,000	
Optimal ate pairing	2,330,000	2,400,000	2,950,000	2,630,000	

^a Intel Core i7 860 (2.8GHz), Windows 7, Visual Studio 2008 Professional

^b Quad-Core AMD Opteron 2376 (2.3GHz), Linux 2.6.18, gcc 4.4.1

^c Intel Core 2 Duo T7100 (1.8GHz), Windows 7, Visual Studio 2008 Professional

^d Athlon 64 X2 Dual Core 6000+(3GHz), Linux 2.6.23, gcc 4.1.2

^e Intel Core 2 Quad Q6600 (2394MHz), Linux 2.6.28, gcc 4.3.3

Comparison Table

A comparison of cycles and timings required by the computation of the ate pairing variants. The frequency is given in GHz and the timings are in milliseconds.

	Alg.	Architecture	Cycles	Freq.	Calc. time
Arapha at al. ICT DCA 2010		Intel Xeon 45nm (1 core)	17,400,000	2.0	8.70
Aranna et al. [CI-RSA 2010]	η_T	Intel Xeon 45nm (8 cores)	3,020,000	2.0	1.51
Developt at al. (Cases and	ητ	Intel Core i7 (1 core)	15,138,000	2.0	5.22
Beuchat et al. [CANS 2009]		Intel Core i7 (8 cores)	5,423,000	2.9	1.87
Hankerson et al.	R-ate	Intel Core 2	10,000,000	2.4	4.10
Naehrig et al. eprint 2010/526, April.6.2010	a _{opt}	Intel Core 2 Quad Q6600	4,470,000	2.4	1.80
Fan et al. CHES'09	"R-ate"	130 nm ASIC	59,976	.204	2.91
This Work eprint 2010/526, jun.17.2010	a _{opt}	Intel Core i7	2,330,000	2.8	0.83
Aranha et al. eprint 2010/526, oct.19.2010	aopt	Intel Core i7	1,703,000	2.8	0.608

A B M A B M

• records do not last long in software!

• However, in hardware the performance records hold longer: Fan *et al.* CHES'09 and Beuchat *et al.* CHES'09 are still the fastest hardware accelerators for ordinary and supersingular curves, respectively.

- However, in hardware the performance records hold longer: Fan *et al.* CHES'09 and Beuchat *et al.* CHES'09 are still the fastest hardware accelerators for ordinary and supersingular curves, respectively.
- Future projects/open problems,
 - To target higher security levels in software implementation of pairings (e.g, 192 bits of security)

- However, in hardware the performance records hold longer: Fan *et al.* CHES'09 and Beuchat *et al.* CHES'09 are still the fastest hardware accelerators for ordinary and supersingular curves, respectively.
- Future projects/open problems,
 - To target higher security levels in software implementation of pairings (e.g, 192 bits of security)
 - ► To design a hardware accelerator faster than any software library for asymmetric pairings over BN curves at 128-bit of security

- However, in hardware the performance records hold longer: Fan *et al.* CHES'09 and Beuchat *et al.* CHES'09 are still the fastest hardware accelerators for ordinary and supersingular curves, respectively.
- Future projects/open problems,
 - To target higher security levels in software implementation of pairings (e.g, 192 bits of security)
 - ► To design a hardware accelerator faster than any software library for asymmetric pairings over BN curves at 128-bit of security
 - to implement efficient pairing-based protocols in software and/or hardware

- However, in hardware the performance records hold longer: Fan *et al.* CHES'09 and Beuchat *et al.* CHES'09 are still the fastest hardware accelerators for ordinary and supersingular curves, respectively.
- Future projects/open problems,
 - To target higher security levels in software implementation of pairings (e.g, 192 bits of security)
 - ► To design a hardware accelerator faster than any software library for asymmetric pairings over BN curves at 128-bit of security
 - to implement efficient pairing-based protocols in software and/or hardware
Thank you for your attention

Questions?

Francisco Rodríguez-Henríquez

Faster Implementation of Pairings (49 / 49)

→ Ξ →

3