Powered by Volcanoes: Three New Algorithms

Andrew V. Sutherland

Massachusetts Institute of Technology

May 12, 2009

http://math.mit.edu/~drew/

arXiv.org/abs/0902.4670 arXiv.org/abs/0903.2785

A 3-volcano of height 2

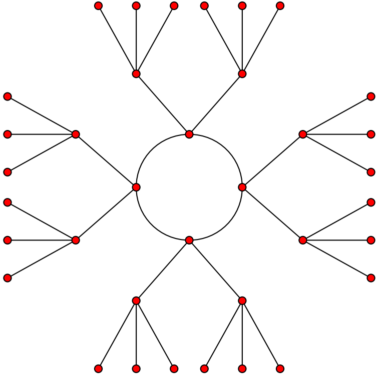

ℓ-volcanoes

An ℓ-*volcano* is a connected undirected graph whose vertices are partitioned into levels $V_0, \ldots, V_h$.

1. The subgraph on V_0 (the *surface*) is a connected d-regular graph, for some $d \leq 2$.

ℓ-volcanoes

An ℓ-*volcano* is a connected undirected graph whose vertices are partitioned into levels $V_0, \ldots, V_h$.

1. The subgraph on V_0 (the *surface*) is a connected d-regular graph, for some $d \leq 2$.
2. For $k > 0$, each $v \in V_k$ has exactly one neighbor in V_{k-1}. All edges not on the surface arise in this manner.
3. For $k < h$, each $v \in V_k$ has degree ℓ+1.

ℓ-volcanoes

An ℓ-*volcano* is a connected undirected graph whose vertices are partitioned into levels $V_0, \ldots, V_h$.

1. The subgraph on V_0 (the *surface*) is a connected d-regular graph, for some $d \leq 2$.
2. For $k > 0$, each $v \in V_k$ has exactly one neighbor in V_{k-1}. All edges not on the surface arise in this manner.
3. For $k < h$, each $v \in V_k$ has degree ℓ+1.

The integers ℓ, h, and $|V_0|$ uniquely determine the shape.

ℓ-isogenies

An *isogeny* $\phi : E_1 \to E_2$ is a morphism that fixes the identity.
It induces a group homomorphism $\phi : E_1(\bar{F}) \to E_2(\bar{F})$.

ℓ-isogenies

An *isogeny* $\phi : E_1 \to E_2$ is a morphism that fixes the identity.
It induces a group homomorphism $\phi : E_1(\bar{F}) \to E_2(\bar{F})$.

The degree of a (separable) isogeny is $|\ker \phi|$.
We are interested in isogenies of prime degree ℓ.
Such an isogeny is necessarily cyclic.

The dual isogeny $\hat{\phi} : E_2 \to E_1$ has the same degree.

The classical modular polynomial Φ_ℓ

The polynomial $\Phi_\ell \in \mathbb{Z}[X, Y]$ has the property

$$\Phi_\ell\bigl(j(E_1), j(E_2)\bigr) = 0 \quad \Longleftrightarrow \quad E_1 \text{ and } E_2 \text{ are } \ell\text{-isogenous.}$$

Note that Φ_ℓ is symmetric in X and Y.

The classical modular polynomial Φ_ℓ

The polynomial $\Phi_\ell \in \mathbb{Z}[X, Y]$ has the property

$$\Phi_\ell\big(j(E_1), j(E_2)\big) = 0 \quad \Longleftrightarrow \quad E_1 \text{ and } E_2 \text{ are } \ell\text{-isogenous.}$$

Note that Φ_ℓ is symmetric in X and Y.

The ℓ-isogeny graph $G_\ell/\mathbb{F}_q$ has vertex set $\{j(E) : E/\mathbb{F}_q\}$ and edges (j_1, j_2) whenever $\Phi_\ell(j_1, j_2) = 0$ (in $\mathbb{F}_q$).

The neighbors of j in G_ℓ are the roots of $\Phi_\ell(X, j) \in \mathbb{F}_q[X]$.

The classical modular polynomial Φ_ℓ

The polynomial $\Phi_\ell \in \mathbb{Z}[X, Y]$ has the property

$$\Phi_\ell\big(j(E_1), j(E_2)\big) = 0 \quad \Longleftrightarrow \quad E_1 \text{ and } E_2 \text{ are } \ell\text{-isogenous.}$$

Note that Φ_ℓ is symmetric in X and Y.

The ℓ-isogeny graph $G_\ell/\mathbb{F}_q$ has vertex set $\{j(E) : E/\mathbb{F}_q\}$ and edges (j_1, j_2) whenever $\Phi_\ell(j_1, j_2) = 0$ (in $\mathbb{F}_q$).

The neighbors of j in G_ℓ are the roots of $\Phi_\ell(X, j) \in \mathbb{F}_q[X]$.

Φ_ℓ is big: $O(\ell^3 \log \ell)$ bits.

The shape of G_ℓ

An elliptic curve is *ordinary* (not *supersingular*) iff its trace is nonzero in $\mathbb{F}_q$. Two curves whose j-invariants lie in the same component of G_ℓ are either both ordinary or both supersingular.

Theorem

The ordinary connected components of G_ℓ are ℓ-volcanoes.
(*assuming $j \neq 0, 1728$*)

Isogenous curves may lie in distint components of G_ℓ.
The components of G_ℓ are a refinement of isogeny classes.

Finding the floor

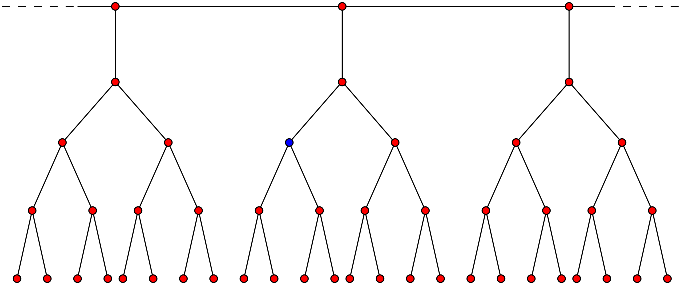

Finding the floor

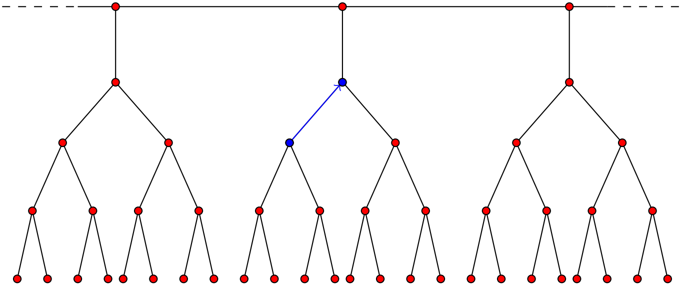

Finding the floor

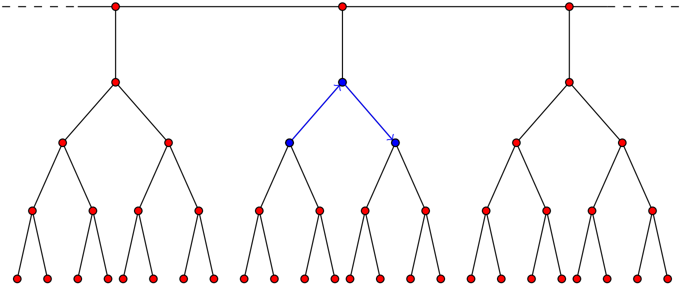

Finding the floor

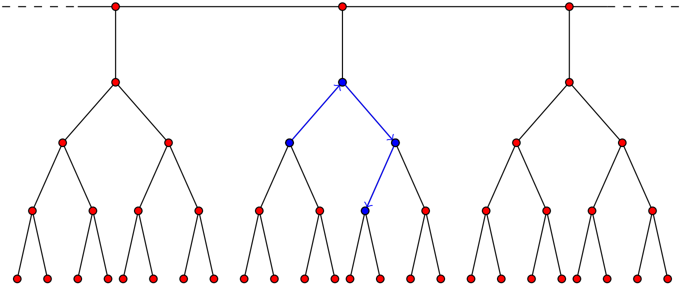

Finding the floor

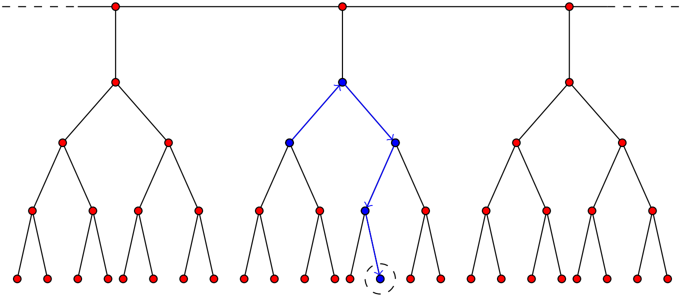

Finding a shortest path to the floor

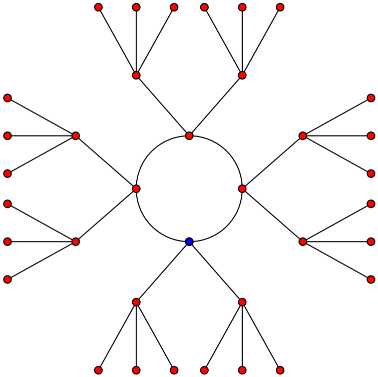

Finding a shortest path to the floor

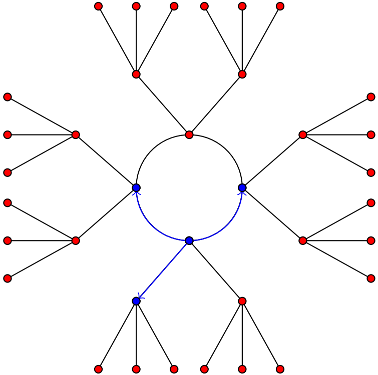

Finding a shortest path to the floor

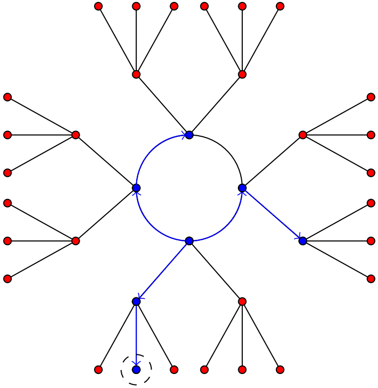

The endomorphism ring End(E)

An endomorphism is an isogeny $\phi : E \to E$.
The multiplication by m map $P \rightsquigarrow mP$ is an example.

The set End(E) of all endomorphisms of E forms a ring which contains a subring isomorphic to $\mathbb{Z}$.

Over $\mathbb{F}_q$ we have $\mathbb{Z} \subsetneq \mathrm{End}(E)$, since

$$\pi : (X, Y) \rightsquigarrow (X^q, Y^q)$$

is not a multiplication by m map.

End(E) for an ordinary elliptic curve

If E is ordinary than End(E) $\cong \mathcal{O}$, where $\mathcal{O}$ is an order in an imaginary quadratic field K.

We may regard π as an element of $\mathcal{O}$ with trace t and norm q. The norm equation for π has the form

$$4q = t^2 - v^2 D_K,$$

where $K = \mathbb{Q}[\sqrt{D_K}]$ and v is the conductor of $\mathbb{Z}[\pi]$.

We have $\mathbb{Z}[\pi] \subseteq \mathcal{O} \subseteq \mathcal{O}_K$, and therefore $\mathcal{O}$ has discriminant $D = u^2 D_K$ for some conductor $u|v$.

The vertical structure of an ℓ-volcano

Theorem (Kohel)

Let $V_0, \ldots, V_h$ be the levels of an ℓ-volcano corresponding to an ordinary component of $G_\ell/\mathbb{F}_q$.

1. *The curves in V_i all have the same endomorphism ring type, with discriminant D_i.*
2. *D_0 has conductor prime to ℓ, and $D_i = \ell^{2i} D_0$.*

The vertical structure of an ℓ-volcano

Theorem (Kohel)

Let $V_0, \ldots, V_h$ be the levels of an ℓ-volcano corresponding to an ordinary component of $G_\ell/\mathbb{F}_q$.

1. *The curves in V_i all have the same endomorphism ring type, with discriminant D_i.*
2. *D_0 has conductor prime to ℓ, and $D_i = \ell^{2i} D_0$.*

This implies $\ell^h \parallel v$, allowing us to determine the height.

The endomorphism ring type of an ordinary elliptic curve E is determined by its level on its ℓ-volcano for each prime $\ell | v$.

The class group action [CM theory]

Suppose $\mathrm{End}(E) \cong \mathcal{O}$, and let $\mathfrak{a}$ an invertible $\mathcal{O}$-ideal.
Let $E[\mathfrak{a}]$ be the points annihilated by all $a \in \mathfrak{a} \subset \mathcal{O} \cong \mathrm{End}(E)$.

There is a separable isogeny $\phi_{\mathfrak{a}} : E \to E/E[\mathfrak{a}]$ with kernel $E[\mathfrak{a}]$, degree $N(\mathfrak{a})$, and $\mathrm{End}(\phi_{\mathfrak{a}}(E)) \cong \mathcal{O}$.

The class group action [CM theory]

Suppose $\mathrm{End}(E) \cong \mathcal{O}$, and let $\mathfrak{a}$ an invertible $\mathcal{O}$-ideal.
Let $E[\mathfrak{a}]$ be the points annihilated by all $a \in \mathfrak{a} \subset \mathcal{O} \cong \mathrm{End}(E)$.

There is a separable isogeny $\phi_{\mathfrak{a}} : E \to E/E[\mathfrak{a}]$ with kernel $E[\mathfrak{a}]$, degree $N(\mathfrak{a})$, and $\mathrm{End}(\phi_{\mathfrak{a}}(E)) \cong \mathcal{O}$.

This defines a group action by the ideal group of $\mathcal{O}$ on the set

$$\mathcal{E}(\mathcal{O}) = \{j(E) : \mathrm{End}(E) \cong \mathcal{O}\},$$

which factors through the class group $\mathrm{cl}(\mathcal{O})$.

The class group action [CM theory]

Suppose $\mathrm{End}(E) \cong \mathcal{O}$, and let $\mathfrak{a}$ an invertible $\mathcal{O}$-ideal.
Let $E[\mathfrak{a}]$ be the points annihilated by all $a \in \mathfrak{a} \subset \mathcal{O} \cong \mathrm{End}(E)$.

There is a separable isogeny $\phi_{\mathfrak{a}} : E \to E/E[\mathfrak{a}]$ with kernel $E[\mathfrak{a}]$, degree $N(\mathfrak{a})$, and $\mathrm{End}(\phi_{\mathfrak{a}}(E)) \cong \mathcal{O}$.

This defines a group action by the ideal group of $\mathcal{O}$ on the set

$$\mathcal{E}(\mathcal{O}) = \{j(E) : \mathrm{End}(E) \cong \mathcal{O}\},$$

which factors through the class group $\mathrm{cl}(\mathcal{O})$.

The above applies over $\mathbb{C}$, but if $E/\mathbb{F}_q$ has $\mathrm{End}(E) \cong \mathcal{O}$, then q is the norm of an element of $\mathcal{O}$ and we may reduce to $\mathbb{F}_q$.

The horizontal structure of an ordinary ℓ-volcano

The degree d of the subgraph on V_0 is $1 + \left(\frac{D_K}{\ell}\right)$.

For $d = 0$ we have $|V_0| = 1$ and for $d = 1$ we have $|V_0| = 2$.

When $d = 2$ there are two $\mathcal{O}$-ideals of norm ℓ, $\mathfrak{a}$ and $\bar{\mathfrak{a}}$, and their ideal classes have order $|V_0|$.

The horizontal structure of an ordinary ℓ-volcano

The degree d of the subgraph on V_0 is $1 + \left(\frac{D_K}{\ell}\right)$.

For $d = 0$ we have $|V_0| = 1$ and for $d = 1$ we have $|V_0| = 2$.

When $d = 2$ there are two $\mathcal{O}$-ideals of norm ℓ, $\mathfrak{a}$ and $\bar{\mathfrak{a}}$, and their ideal classes have order $|V_0|$.

The set $\mathcal{E}(\mathcal{O})$ has size $h(\mathcal{O})$ and is comprised of the surfaces of isomorphic ℓ-volcanoes corresponding to cosets in $\mathrm{cl}(\mathcal{O})$.

And in general, $\mathcal{E}(\mathcal{O})$ is a *torsor* for $\mathrm{cl}(\mathcal{O})$.

The CM method

If $E/\mathbb{F}_q$ has $N = q + 1 - t$ points, with $t \not\equiv 0$ in $\mathbb{F}_q$, then

$$4q = t^2 - v^2 D,$$

where D is the discriminant of $\mathcal{O} \cong \mathrm{End}(E)$. Conversely, any curve with $\mathrm{End}(E) \cong O$ has trace $\pm t$.

The Hilbert class polynomial $H_D \in \mathbb{Z}[X]$ is defined by

$$H_D(X) = \prod_{j \in \mathcal{E}(\mathcal{O})} (X - j).$$

Its roots are the j-invariants of curves with $\mathrm{End}(E) \cong O$.

Given a root of H_D in $\mathbb{F}_q$, we may construct $E/\mathbb{F}_q$ with N points.

Computing $H_D(X)$ with the CRT [ALV '06, BBEL '08]

To compute $H_D \in \mathbb{F}_q[X]$ it suffices to compute H_D modulo many "small" primes p and apply the Chinese Remainder Theorem.

For primes of the form $4p = t_p^2 - v_p^2 D$, H_D splits completely over $\mathbb{F}_p$ and we may compute H_D mod p by finding its roots.

Computing $H_D(X)$ with the CRT [ALV '06, BBEL '08]

To compute $H_D \in \mathbb{F}_q[X]$ it suffices to compute H_D modulo many "small" primes p and apply the Chinese Remainder Theorem.

For primes of the form $4p = t_p^2 - v_p^2 D$, H_D splits completely over $\mathbb{F}_p$ and we may compute H_D mod p by finding its roots.

To find the first root, generate random curves over $\mathbb{F}_p$ until we find one with $\mathrm{End}(E) \cong \mathcal{O}$ (or any E with trace $\pm t$).

To enumerate the other roots, use the group action of $\mathrm{cl}(\mathcal{O})$.

Improvements [S '09]

The CRT approach to computing H_D can be improved:

1. Compute H_D mod P in $O(|D|^{1/2+\epsilon} \log P)$ space.
2. Generate "random" curves with prescribed torsion.
3. Make v_p large (bigger volcanoes are easier to find).
4. Use an optimal presentation of $\mathrm{cl}(\mathcal{O})$ to minimize norms.

An example of a polycyclic presentation

For $D = -79947$, $\mathrm{cl}(D)$ is cyclic of order $h(D) = 100$.
It is generated by the class of an ideal with norm 19.

But $\mathrm{cl}(D)$ is also generated by classes α_2 and α_{13} of ideals of norm 2 and 13. The elements α_2 and α_{13} have orders 20 and 50 and are not independent ($\alpha_{13}^5 = \alpha_2^{18}$).

An example of a polycyclic presentation

For $D = -79947$, $\mathrm{cl}(D)$ is cyclic of order $h(D) = 100$.
It is generated by the class of an ideal with norm 19.

But $\mathrm{cl}(D)$ is also generated by classes α_2 and α_{13} of ideals of norm 2 and 13. The elements α_2 and α_{13} have orders 20 and 50 and are not independent ($\alpha_{13}^5 = \alpha_2^{18}$).

Nevertheless, every $\beta \in \mathrm{cl}(D)$ can be written uniquely as

$$\beta = \alpha_2^{e_2} \alpha_{13}^{e_{13}}$$

with $0 \le e_2 < 20$ and $0 \le e_{13} < 5$.

Using this presentation is about 100 times faster.

Running the rim

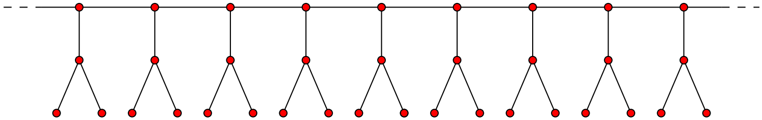

Running the rim

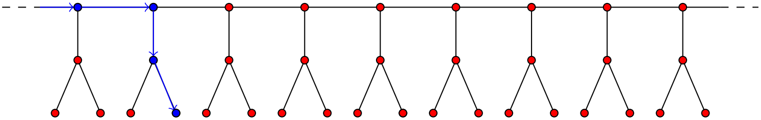

Running the rim

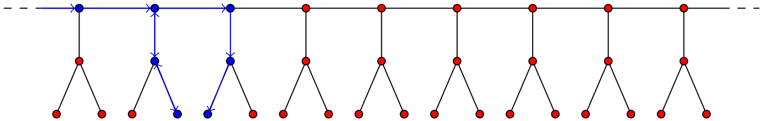

Running the rim

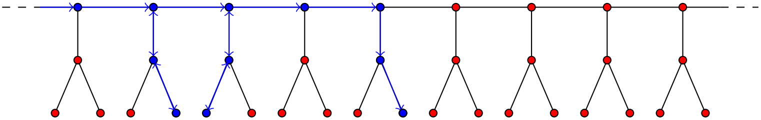

Running the rim

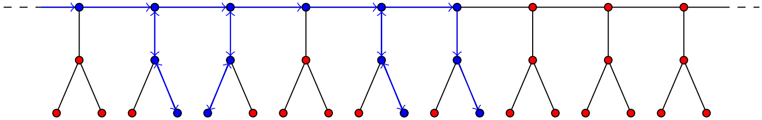

Running the rim

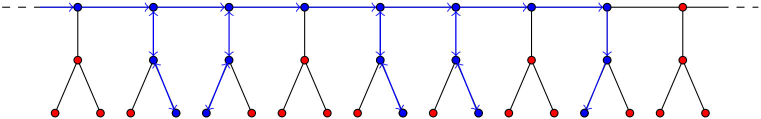

Running the rim

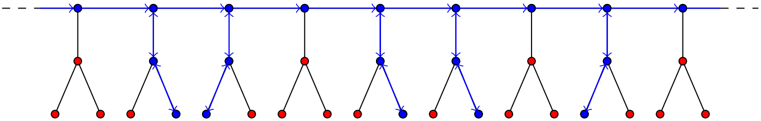

Record-breaking CM constructions

Largest $|D|$

Old Record (June 2008, complex analytic [Enge])

$D = -70,901,505,867$ $\qquad$ $h(D) = 51,244$

New Record (October 2008, CRT method [Enge-S])

$D = -102,197,306,669,747$ $\qquad$ $h(D) = 2,014,236$

Record-breaking CM constructions

Largest $|D|$

Old Record (June 2008, complex analytic [Enge])

$D = -70,901,505,867$ $\quad$ $h(D) = 51,244$

New Record (October 2008, CRT method [Enge-S])

$D = -102,197,306,669,747$ $\quad$ $h(D) = 2,014,236$

Largest $h(D)$

Old Record (January 2006, complex analytic [Enge])

$D = -2,093,236,031$ $\quad$ $h(D) = 100,000$

New Record (April 2009, CRT method, [Bröker-S])

$D = -4,058,817,012,071$ $\quad$ $h(D) = 5,000,000$

Performance comparison

		Analytic $\mathfrak{w}_{3,13}$		CRT $\mathfrak{f}^2$		CRT $\mathfrak{f}$	
$-D$	$h(D)$	height	time	height*	time*	height	time
6961631	5000	9.5k	28	9.5k	4.9	3.8k	2.0
23512271	10000	20k	210	20k	24	8.0k	9.1
98016239	20000	45k	1,800	45k	120	18k	46
357116231	40000	97k	14,000	97k	574	38k	220
2093236031	100000	265k	260,000	265k	4,400	103k	1,600

Complex Analytic vs. CRT method

(2.4 GHz AMD Opteron CPU seconds)

*increased to match the height bound for $\mathfrak{w}_{3,13}$.

Computing End(E) [Bisson-S '09]

Given $E/\mathbb{F}_q$ we may compute t and factor $4q - t^2$ to obtain

$$4q = t^2 - v^2 D_K.$$

The discriminant of End(E) $\cong \mathcal{O}$ is $D = u^2 D_K$ for some $u|v$ To determine End(E) it suffices to compute u.

Computing End(E) [Bisson-S '09]

Given $E/\mathbb{F}_q$ we may compute t and factor $4q - t^2$ to obtain

$$4q = t^2 - v^2 D_K.$$

The discriminant of End(E) $\cong \mathcal{O}$ is $D = u^2 D_K$ for some $u|v$ To determine End(E) it suffices to compute u.

Let $u_1, \ldots, u_n$ be the factors of v. To distinguish u, we seek *relations* that hold in some $\mathrm{cl}(u_i^2 D_K)$ but not others.

We test these relations in the isogeny graph by walking along the surface of various ℓ-volcanoes.

Relations in class groups

A relation R is a pair of vectors $(\ell_1, \ldots, \ell_r)$ and $(e_1, \ldots, e_r)$, with $\ell_i \nmid v$ and $\left(\frac{D_K}{\ell_i}\right) = 1$.

We say R *holds* in $\mathrm{cl}(D)$ if for each i there is an $\alpha_i \in \mathrm{cl}(D)$ containing an ideal of norm ℓ_i such that $\alpha_1^{e_1} \cdots \alpha_r^{e_r} = 1$.

Relations in class groups

A relation R is a pair of vectors $(\ell_1, \ldots, \ell_r)$ and $(e_1, \ldots, e_r)$, with $\ell_i \nmid v$ and $\left(\frac{D_K}{\ell_i}\right) = 1$.

We say R *holds* in $\mathrm{cl}(D)$ if for each i there is an $\alpha_i \in \mathrm{cl}(D)$ containing an ideal of norm ℓ_i such that $\alpha_1^{e_1} \cdots \alpha_r^{e_r} = 1$.

More generally, define the *cardinality* of R in $\mathrm{cl}(D)$ by

$$\#R/\mathrm{cl}(D) = \#\left\{\tau \in \{\pm 1\}^r : \prod \alpha_i^{\tau_i e_i} = 1 \text{ in } \mathrm{cl}(D)\right\}.$$

For $p|v$, let $D_1 = (v/p)^2 D_K$ and $D_2 = p^2 D_K$. We want

$$\#R/\mathrm{cl}(D_1) > \#R/\mathrm{cl}(D_2).$$

Counting relations in the isogeny graph

To compute $\#R/\operatorname{cl}(\mathcal{O})$:

1. Let J_0 be a list consisting of $j(E)$.
2. For i from 1 to r:
   - For each $j \in J_{i-1}$, walk e_i steps in both directions on the surface of the ℓ_i-volcano and append the endpoints to J_i.
3. Output the number of times $j(E)$ occurs in J_r.

Counting relations in the isogeny graph

To compute $\#R/\operatorname{cl}(\mathcal{O})$:

1. Let J_0 be a list consisting of $j(E)$.
2. For i from 1 to r:
   - For each $j \in J_{i-1}$, walk e_i steps in both directions on the surface of the ℓ_i-volcano and append the endpoints to J_i.
3. Output the number of times $j(E)$ occurs in J_r.

To compute $\#R/\operatorname{cl}(\mathcal{O})$ efficiently, we use *smooth* relations, where ℓ_i, e_i, and r are all small.

Record-breaking End(E) computations

Heuristically, we achieve a running time of $L[1/2, \sqrt{3}/2]$.

Over a 200-bit prime field, under 15 minutes.
Over a 256-bit prime field, about 4 hours.

These are worst-case examples (average case is easy).

Record-breaking End(E) computations

Heuristically, we achieve a running time of $L[1/2, \sqrt{3}/2]$.

Over a 200-bit prime field, under 15 minutes.
Over a 256-bit prime field, about 4 hours.

These are worst-case examples (average case is easy).

Kohel's algorithm has complexity $O(q^{1/3})$ (under the GRH).
It cannot feasibly compute End(E) over a cryptographic size field when v contains a large prime factor.

Computing Φ_ℓ with the CRT method [Bröker-Lauter-S]

Choose CRT primes $p \equiv 1 \mod \ell$ with $4p = t^2 - v^2\ell^2 D$.
Suppose we have an ℓ-volcano of height 1 with $|V_0| \geq \ell + 2$.
(we may pick D to ensure this).

Computing Φ_ℓ with the CRT method [Bröker-Lauter-S]

Choose CRT primes $p \equiv 1 \mod \ell$ with $4p = t^2 - v^2\ell^2 D$.
Suppose we have an ℓ-volcano of height 1 with $|V_0| \geq \ell + 2$.
(we may pick D to ensure this).

We can "construct" this volcano without using Φ_ℓ:

1. Use $H_D(X)$ to find the surface.
2. Apply the action of $\mathrm{cl}(D)$ to enumerate the surface.
3. Use Velu's formula to descend to the floor.
4. Apply the action of $\mathrm{cl}(\ell^2 D)$ to enumerate the floor.

From this we can interpolate Φ_ℓ mod p.

Record-breaking Φ_ℓ computations

The time to compute Φ_ℓ is $O(\ell^3 \log^{3+\epsilon} \ell)$ [GRH].
Faster than the best alternative by a factor of $\log \ell$.

Record Φ_ℓ computations (classical)

Computed Φ_ℓ for all $\ell < 3000$, and up to $\ell = 5003$.
Output is generated at a rate of about 5Mb/s.

Previous record: $\ell < 360$ [Rubinstein-Seroussi].

Record-breaking Φ_ℓ computations

The time to compute Φ_ℓ is $O(\ell^3 \log^{3+\epsilon} \ell)$ [GRH].
Faster than the best alternative by a factor of $\log \ell$.

Record Φ_ℓ computations (classical)

Computed Φ_ℓ for all $\ell < 3000$, and up to $\ell = 5003$.
Output is generated at a rate of about 5Mb/s.

Previous record: $\ell < 360$ [Rubinstein-Seroussi].

Record modular polynomial computations (Weber $\mathfrak{f}$)

Computed Φ_ℓ for all $\ell < 10000$ and up to $\ell = 50021$.

Preprint in preparation.

Modular polynomials for $\ell = 7$

Classical:

$$X^8 + Y^8 - X^7Y^7 + 5208X^7Y^6 - 10246068X^7Y^5 + 9437674400X^7Y^4 - 4079701128594X^7Y^3 +$$
$$720168419610864X^7Y^2 - 34993297342013192X^7Y + 104545516658688000X^7 +$$
$$\ldots \text{(2 pages omitted)} \ldots +$$
$$13483958224762213714698012883865296529472356352000000000000000Y^3 +$$
$$1464765079488386840337633731737402825128271675392000000000000000000Y^2$$

Atkin:

$$X^8 - X^7Y + 744X^7 + 196476X^6 + 357X^5Y + 21226520X^5 + 1428X^4Y +$$
$$803037606X^4 - 31647X^3Y + 14547824088X^3 - 204792X^2Y + 138917735740X^2 +$$
$$186955XY + 677600447400X + Y^2 + 2128500Y + 1335206318625$$

Canonical:

$$X^8 + 28X^7 + 322X^6 + 1904X^5 + 5915X^4 + 8624X^3 + 4018X^2 - XY + 748X + 49$$

Weber:

$$X^8 + Y^8 - X^7Y^7 + 7X^4Y^4 - 8XY$$

Powered by Volcanoes: Three New Algorithms

Andrew V. Sutherland

Massachusetts Institute of Technology

May 12, 2009

http://math.mit.edu/~drew/

arXiv.org/abs/0902.4670   arXiv.org/abs/0903.2785