Isogeny graphs, modular polynomials, and point counting for higher genus curves

Chloe Martindale

July 7, 2017

These notes are from a talk given in the Number Theory Seminar at INRIA, Nancy, France. The contents of the talk include research from the PhD thesis of the author, which was written under the supervision of Dr Marco Streng, and research from an article which was started at the AGC2016 workshop at UCLA, as joint work together with Sean Ballentine, Aurore Guillevic, Elisa Lorenzo-Garcia, Maike Massierer, Ben Smith, and Jaap Top.

1 Motivation: elliptic curves

In curve-based cryptography, it is important to develop fast algorithms for computing isogenies between curves, for computing endomorphism rings, and for counting points on curves defined over finite fields $\mathbb{F}_p$, where p is a very large prime. All of this research was inspired by previous research into elliptic curves, and so we first recall definitions and results for elliptic curves as a motivation for the higher genus case.

Definition. Suppose that E and E' are elliptic curves over a field k. An *isogeny* $\phi : E \to E'$ is a surjective morphism with finite kernel that sends the identity to the identity.

Remark. Some people consider the constant-zero morphism to be an isogeny, which is not consistent with the above definition. As this morphism will not play a role in our work, we do not include the constant-zero morphism in our definition of isogeny.

Definition. Suppose that $\phi : E \to E'$ is an isogeny of elliptic curves over a field k. This induces an injective morphism of function fields

$$\overline{k}(E') \longrightarrow \overline{k}(E).$$

We define the *degree* of ϕ to be

$$\deg(\phi) = [\overline{k}(E) : \overline{k}(E')].$$

Question 1. *A natural question to ask now is, when are 2 curves isogenous?*

As motivation to answer this question, consider the following: if we can break the discrete logarithm problem on an elliptic E' and efficiently compute an isogeny $E' \to E$, then we can break the discrete logarithm problem on E. We answer this question by looking at *isogeny graphs*.

Definition. An *ℓ-isogeny graph of elliptic curves* as an undirected graph for which each vertex represents a j-invariant (this is an isomorphism invariant) of an elliptic curve over a field k, and an edge between $j(E)$ and $j(E')$ represents an ℓ-isogeny $E \to E'$ defined over k and its dual isogeny $E' \to E$.

Definition. An *ℓ-volcano* is an undirected connected graph whose vertices are partitioned into one or more levels $V_0, \dots, V_d$ such that the following hold:

1. The subgraph on level V_0 is a regular graph of degree at most 2.
2. For $i > 0$, each vertex in V_i has exactly one neighbour in level V_{i-1}, and this accounts for every edge not on the surface.
3. For $i < d$, each vertex in V_i has degree $\ell + 1$.

Example. Here is a 2-volcano with $d = 2$:

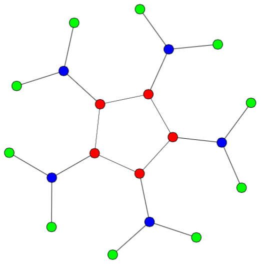

Theorem (Kohel '96)**.** *Let $E/\mathbb{F}_q$ be an ordinary elliptic curve with $j(E) \neq 0, 1728$. Then the connected component of the ℓ-isogeny graph containing $j(E)$ is a ℓ-*volcano.

Remark. The depth is given by $\max\{r \in \mathbb{Z} : \ell^r | [\mathcal{O}_K : \mathbb{Z}[\pi]]\}$, where π is the q-power Frobenius endomorphism of E and $K = \mathbb{Q}(\pi)$. So the depth is as easy to compute as the Frobenius endomorphism. (We'll come back to this later). The length of the cycle and the number of connected components are also easy to compute.

Now with a simple path walking algorithm we can determine if $j(E)$ and $j(E')$ are in the same connected component of the isogeny graph, hence determine if they are isogenous, and if they are, determine the degree of the isogeny (or at least of one of the isogenies).

Question 2. *A natural question to ask at this point is: given an elliptic curve over a field k, and an integer ℓ, can we enumerate all the elliptic curves E' over k such that there exists an isogeny $E \to E'$ of degree ℓ? That is, can we compute the neighbours in the isogeny graph?*

We answer this question using modular polynomials.

Definition. For each $\ell \in \mathbb{Z}_{\geq 2}$, the *modular polynomial of level* ℓ is a non-constant-zero polynomial

$$\Phi_\ell(X, Y) \in \mathbb{C}[X, Y]$$

such that, given any 2 elliptic curves E and E' over $\mathbb{C}$, there exists an isogeny $E \to E'$ of degree ℓ if and only if $\Phi_\ell(j(E), j(E')) = 0$.

Remark. If $\ell \in \mathbb{Z}_{\geq 2}$ is prime, then the degree of $\Phi_\ell(X, Y)$ in both X and Y is given by $\ell + 1$.

Remark. In fact the coefficients of $\Phi_\ell(X, Y)$ are integers, not just complex numbers. This allows us to reduce the coefficients modulo a prime p, so that furthermore, given any elliptic curves E and E' over $\overline{\mathbb{F}_p}$, there exists an isogeny $E \to E'$ of degree ℓ if and only if $\Phi_\ell(j(E), j(E')) \equiv 0 \bmod p$.

For small ℓ, equations for $\Phi_\ell(X, Y)$ can be found for example at LMFDB. These equations can also be thought of as models for the modular curve $X_0(\ell)$ of level ℓ.

We can now answer Question 2: given an elliptic curve $E/\mathbb{F}_p$, compute $j(E) \in \mathbb{F}_p$, and compute the $\mathbb{F}_p$-vauled polynomial in $\Phi_\ell(j(E), Y) \bmod p$ in Y. The roots of this polynomials then give us the j-invariants of each curve $E'/\overline{\mathbb{F}_p}$ for which there exists an isogeny $E \to E'$ of degree ℓ.

Recall that the depth of the isogeny volcano was computed using the q-power Frobenius; this works because isogenous curves over $\mathbb{F}_q$ have the same Frobenius polynomial. Furthermore, if t is the trace of Frobenius of $E/\mathbb{F}_q$, then

$$\#E(\mathbb{F}_q) = 1 + q - t.$$

Therefore, if two elliptic curves over $\mathbb{F}_q$ do not have the same number of $\mathbb{F}_q$-rational points then they are not isogenous.

Question 3. *A natural question at this point is, given $E/\mathbb{F}_p$, what is the most efficient way of counting $\#E(\mathbb{F}_p)$ for a large prime p?*

We answer this with the beautiful theorem of Schoof, Elkies, and Atkin, together with the bound $|t| < 2\sqrt{p}$, which gives us a polynomial time algorithm for computing t, and hence $\#E(\mathbb{F}_p)$.

Theorem (Schoof, Atkin, Elkies)**.** *Let E be an ordinary elliptic curve over $\mathbb{F}_p$ such that $j(E) \neq 0, 1728$, and write the modular polynomial of level ℓ*

$$\Phi_\ell(j(E), Y) = f_1(Y) \cdots f_n(Y)$$

as the product of irreducible polynomials in $\mathbb{F}_p[Y]$. Then (up to ordering) the polynomials f_i satisfy one of the following:

1. *$n = 2$, deg$(f_1) = 1$, and deg$(f_2) = \ell$.*

2. $deg(f_1) = deg(f_2) = 1$, and for every $i > 2$, $deg(f_i) = r > 1$, for some $r \in \mathbb{Z}$.

3. for every i, $deg(f_i) = r$, for some $r \in \mathbb{Z}$.

Furthermore, there exists a primitive r^{th} root of unity $\zeta \in \overline{\mathbb{F}_\ell}$ such that

$$t^2 \equiv (\zeta + \zeta^{-1})^2 p \ mod \ \ell,$$

where in case (1) we set $\zeta = 1$.

We unfortunately do not have time to prove this theorem, although the proof is beautiful and elementary. Schoof has written a report on this theorem, which is referenced as [Sch].

We can now answer Question 3 in the following way: given an elliptic curve E over $\mathbb{F}_p$, we first compute t^2 modulo ℓ for many different primes ℓ (perhaps up to a root of unity). We then use the bound on $|t|$ and the Chinese Remainder Theorem to compute t^2. Lastly, we check the sign of t, for example by multiplying a non-trivial point $P \in E(\mathbb{F}_p)$ by both $1+p-t$ and $1+p+t$. For large p, this is much more efficient than any other known algorithm for counting points.

2 Abelian varieties and isomorphism invariants

We are able to define a discrete logarithm on elliptic curves, classify isogenies of elliptic curves using isogeny graphs, and give polynomial time point counting algorithms on elliptic curves largely due to one property: that there exists a group law. Recall that an elliptic curve (for odd characteristic) is defined by a polynomial

$$y^2 = f(x),$$

where $\deg(f) = 3$. One could ask, what happens if $\deg(f) > 3$? Or what about other algebraic curves? One of the reasons that we so often stick to such a special class of algebraic curves is because of the simple group law. But all is not lost for other algebraic curves: although there is no known group law on the curves themselves, to each algebraic curve C we can associate an *abelian variety* (on which there exists a group law), called the *Jacobian* of C, written $J(C)$, or $\mathrm{Jac}(C)$. In fact, we can do even better, we can assume that the Jacobian is a *principally polarised* abelian variety - which for all purposes of this talk means 'nice'. Furthermore, if C is defined over k, then

$$C(k) \subseteq \mathrm{Jac}(C)(k),$$

so we can study the k-rational points of C by studying the points on the Jacobian, where we have a group law to help us. To generalise what we have seen in this talk for elliptic curves to all algebraic curves, we must answer 4 questions:

Question 4. *(a) How do we associate an isomorphism invariant to a principally polarised abelian variety (e.g. the Jacobian of an algebraic curve)?*

(b) What form do isogeny graphs of principally polarised abelian varieties take?

(c) How do we define a modular polynomial, and can we compute it?

(d) How will these modular polynomials factor, and does the method of Schoof, Elkies, and Atkin generalise in a natural way?

Just as for elliptic curves, we restrict now to the case of *ordinary* principally polarised abelian varieties: that is, if $A/\mathbb{F}_q$ is a principally polarised abelian variety, we assume that there exist non trivial $\overline{\mathbb{F}_q}$-rational $\mathbb{F}_q$-torsion points on A. This is because the results in the non-ordinary case are fundamentally different, although certainly interesting in their own right, and so saved for another research project.
In defining isogeny graphs, and especially modular polynomials, of elliptic curves, the j-invariant (isomorphism invariant) of an elliptic curve played a crucial role, so we first attempt to answer (a). For experts: we first note that if we restrict to genus 2 curves (so to dimension 2), then one might think that the natural choice would be Igusa invariants. However, (b), (c), and (d) look a lot more complicated with this choice, so we use the extra structure given to us by the fact that we are only interested in the problems over finite fields.

Recall: for any abelian variety A over a field $\mathbb{F}_q$, we have endomorphisms on A given by

$$\begin{array}{ccc} A & \longrightarrow & A \\ P & \mapsto & nP \end{array}$$

for every $n \in \mathbb{Z}$, as well as the Frobenius

$$\begin{array}{cccc} \pi : & A & \longrightarrow & A \\ & P & \mapsto & P^q. \end{array}$$

In particular, we have that $\mathbb{Z}[\pi] \subseteq \mathrm{End}(A)$. Futhermore, A is simple and ordinary, then the q-power Frobenius π generates a CM-field $K = \mathbb{Q}(\pi)$, and $\mathrm{End}(A)$ is an order in K. Recall below the definition of a CM-field:

Definition. A *CM-field* K is a totally imaginary degree 2 extension of a totally real number field K_0. For $\mathrm{End}(A)$ an order in K, we say that A has *real multiplication by* K_0

Our isomorphism invariant for principally polarised ordinary abelian varieties is a tuple of functions $(j_1, \ldots, j_d)$ with $d < g + 1$ depending K_0.

These are still relatively nice for principally polarised ordinary abelian varieties of dimension 2 (dimension 1 is the case of elliptic curves), where for characteristic $\neq 2$ every such abelian variety over k is the Jacobian of a curve of the form

$$y^2 = f(x),$$

where $f(x) \in k[x]$ has degree 5 or 6; these curves are called *genus 2 curves.* For a genus 2 curve $C : y_2 = f(x) \in \mathbb{F}_p[x]$ with real multiplication by K_0, we can define rational functions

$$j_1, j_2, j_3 : \mathbb{F}_p[x]|_{\deg=5,6} \longrightarrow \mathbb{A}^3_{\mathbb{F}_p}$$

such that the tuple

$$(j_1(f), j_2(f), j_3(f))$$

determines $\mathcal{J}(C)$ up to (real-multiplication preserving) isomorphism. The only fields K_0 for which these functions are explicitly written down are $K_0 = \mathbb{Q}(\sqrt{5})$ and $K_0 = \mathbb{Q}(\sqrt{8})$, in work by Müller. The equations for j_1 and j_2 are known for all real quadratic fields, thanks to work by Lauter and Yang, but the problem of finding a general equation for j_3 is still open. The existence of functions defining isomorphism invariants (in arbitrary dimension) is proven in the thesis of the author (and perhaps has been done elsewhere, unknown to her).

3 Isogeny graphs of abelian varieties

We now want to look at isogenies for higher genus curves, which we do by studying isogenies of principally polarised ordinary abelian varieties (p.p.o.a.vs). The following definition is the same as the one for elliptic curves:

Definition. A morphism of abelian varieties is an *isogeny* if it preserves the identity, is surjective, and has finite kernel.

The generalisation of an ℓ-isogeny to higher dimension that we use is quite complicated, so we do not give the details here. The interested reader can find the definition in the upcoming thesis of the author [Mar]. We again associate a prime to the isogeny, but now a prime ideal in $\mathcal{O}_{K_0}$ - we study 'μ-isogenies', where μ is a totally positive element of $\mathcal{O}_{K_0}$ which generates a prime ideal in K_0.

Definition. A μ-isogeny graph of p.p.o.a.vs is an undirected graph for which each vertex represents a $(j_1, \dots, j_d)$-invariant of a p.p.o.a.v. over a field $\mathbb{F}_q$, and an edge between $(j_1, \dots, j_d)(A)$ and $(j_1, \dots, j_d)(A')$ represents a μ-isogeny $A \to A'$ defined over $\mathbb{F}_q$ together with its dual isogeny $(A')^\vee \to A^\vee$. (Note that $(j_1, \dots, j_d)(A) = (j_1, \dots, j_d)(A^\vee)$ and $(j_1, \dots, j_d)(A') = (j_1, \dots, j_d)((A')^\vee)$.)

Let I be the graph with one vertex and no edges, let R_1 be a 1-cycle with one edge of weight $\frac{1}{2}$, let R_2 be 2 vertices joined by a single edge, and let C_n be a cycle of length n.

Theorem (M. '17)**.** *Let $A/\mathbb{F}_q$ be a principally polarised ordinary abelian variety and suppose that the only roots of unity in* $\mathrm{End}(A) \otimes \mathbb{Q}$ *are* ± 1*. Then the connected component of the μ-isogeny graph containing A is a* $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)$*-volcano with $V_0 \in \{I, R_1, R_2, C_n\}$.*

Remark. We have formulae for Γ and d given A, and just like the dimension 1 case, the depth d is easy compute using the Frobenius, and is 0 for all but finitely many μ (up to multiplication by units).

Remark. A similar theorem (but not with μ-isogenies) was given for the genus 2 case by Ionica and Thomé in [IT]. Independently, Brooks, Jetchev, and Wesolowski proved a similar statement (in arbitrary dimension) in [BJW].

4 Modular polynomials for genus 2 curves over finite fields

Having answered Question 4(b), we want to know how to compute paths in our isogeny graphs, i.e., we turn to Question 4(c): defining and computing modular polynomials in dimension g. We can currently only implement an algorithm to compute modular polynomials in dimension 2 (i.e. for Jacobians of genus 2 curves), so for simplicity we will work now with $g = 2$. We will use the isomorphism invariants of the previous section, and so we fix the prime p, the real quadratic number field K_0, and an isomorphism invariant (j_1, j_2, j_3) for K_0 throughout. In the elliptic curve case, the modular polynomial of level ℓ told us about isogenies of degree ℓ, otherwise known as ℓ-isogenies. The generalisation of the modular polynomial tells us about μ-isogenies, where μ is a totally positive element K_0 that generates a prime ideal. It is given by the following theorem, which is proven in the upcoming thesis of the author.

Theorem. *There exists an algorithm to compute polynomials*

$$\begin{array}{rcl} G_\mu(X_1, X_2, X_3, Z_1) & \in & \mathbb{Z}[X_1, X_2, X_3, Z_1] \\ H_{\mu,2}(X_1, X_2, X_3, Z_1, Z_2) & \in & \mathbb{Z}[X_1, X_2, X_3, Z_1, Z_2] \\ H_{\mu,3}(X_1, X_2, X_3, Z_1, Z_3) & \in & \mathbb{Z}[X_1, X_2, X_3, Z_1, Z_3] \end{array}$$

with

$$\deg_{Z_1}(G_\mu) = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1, \quad \deg_{Z_2}(H_{\mu,2}) = 1, \quad \deg_{Z_3}(H_{\mu,3}) = 1,$$

such that for 'most' genus 2 curves $C/\mathbb{C}$ with $C : y^2 = f(x)$, and $C'/\mathbb{C}$ with $y^2 = f(x)'$, there exists a μ-isogeny $\mathcal{J}(C) \to \mathcal{J}(C')$ if and only if

$$\begin{array}{c} G_\mu(j_1(f), j_2(f), j_3(f), j_1(f')) = 0 \\ H_{\mu,2}(j_1(f), j_2(f), j_3(f), j_1(f'), j_2(f')) = 0 \\ H_{\mu,3}(J_1(f), j_2(f), j_3(f), j_1(f'), j_3(f')) = 0. \end{array}$$

For the precise definition of 'most', see the upcoming thesis of the author. As in the genus 1 case, we can reduce these polynomials mod p to detect when 2 curves over $\overline{\mathbb{F}_p}$ are μ-isogeneous. The algorithm has been implemented in the cases for which j_1, j_2 and j_3 are known, and the polynomials are computed up to $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) = 19$. Currently the author is working with Marius Vuille on a new algorithm with which we hope to compute modular polynomials up to at least norm 200. As more polynomials are computed, they can be found at `www.martindale.info`.

Remark. Enea Milio has a similar theorem on computing modular polynomials in genus 2 (not with μ-isogenies) in [Mil]. The theoretical complexity of Milio's algorithm is also similar, although through superior implementation skills he has managed to compute modular polynomials up to norm 97.

So, given a genus 2 curve $C/\mathbb{F}_p$, we can enumerate all the (invariants of) genus 2 curves $C'/\overline{\mathbb{F}_p}$ for which there exists a μ-isogeny to $\mathcal{J}(C) \to \mathcal{J}(C')$ in the same way as we did for elliptic curves. That is, given $C/\mathbb{F}_p : y^2 = f(x)$, we can

1. Compute $j_1(f), j_2(f), j_3(f)$.
2. Enumerate the solutions of $G_\mu(j_1(f), j_2(f), j_3(f), Z_1) = 0$, which gives us $j_1(f')$ for every $C' : y^2 = f(x)'$ for which $\mathcal{J}(C')$ is μ-isogenous to $\mathcal{J}(C)$.
3. For each C', find the unique $j_2(f')$ and $j_3(f')$ that satisfy

   $$H_{\mu,2}(j_1(f), j_2(f), j_3(f), j_1(f'), j_2(f')) = 0$$

   and

   $$H_{\mu,3}(J_1(f), j_2(f), j_3(f), j_1(f'), j_3(f')) = 0.$$

If one requires the equation of the curve, in the thesis of the author there are formulae to find the Igusa invariants in terms of j_1, j_2 and j_3 (for $K_0 = \mathbb{Q}(\sqrt{5})$), and we can then use Mestre's algorithm to find the curve. We now answer the remaining part of Question 4, part (d).

5 Schoof's algorithm in genus 2

This section is joint work with Ballentine, Guillevic, Lorenzo-Garcia, Massierer, Smith, and Top. As before, we fix p, K_0, j_1, j_2, and j_3. We again need to recall how the number of $\mathbb{F}_p$-points on a genus 2 curves relates to the Frobenius polynomial:
Let C be a genus 2 curve over $\mathbb{F}_p$; then there exist integers s and t such that the characteristic polynomial of the p-power Frobenius on $\mathcal{J}(C)$ is given by

$$X^4 - tX^3 + (2p+s)X^2 - tpX + p^2.$$

Then in particular, we have the following facts:

1. $\#C(\mathbb{F}_p) = 1 + p - t$,
2. $\#\mathcal{J}(C)(\mathbb{F}_p) = 1 - t + 2p + s - tp + p^2$,
3. $|s| < 4p$, and
4. $|t| < 4\sqrt{p}$.

Given these facts, we hope for a Schoof-style algorithm to compute s and t, and the following theorem gives us just that.

Theorem. *Let $C/\mathbb{F}_p$ be a genus 2 curve, $C : y^2 = f(x)$, such that $\mathcal{J}(C)$ is simple and ordinary and $(\mathrm{End}(\mathcal{J}(C)) \otimes \mathbb{Q})$ is some CM-field K, where the maximal totally real subfield of K is K_0 and the only roots of unity in K are ± 1. Then for a totally positive element $\mu \in \mathcal{O}_{K_0}$ such that $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) = \ell$ is prime, the factorisation of*

$$G_\mu(j_1(f), j_2(f), j_3(f), Z_1) \textit{ mod } p = f_1 \cdots f_n$$

into irreducible polynomials in $\mathbb{F}_p[Z_1]$ satisfies one of

1. $\deg(f_1) = 1$ *and* $\deg(f_2) = \ell$,
2. $\deg(f_1) = 2$ *and for* $i > 1$, $\deg(f_i) = r$, *or*
3. $\deg(f_1) = \deg(f_2) = 1$, *and for* $i > 2$, $\deg(f_i) = r$, *or*
4. *for every* i, $\deg(f_i) = r$.

Furthermore, there exist primitive $2r^{th}$-roots of unity ζ_{2r} and ζ'_{2r} in $\overline{\mathbb{F}_\ell}$ such that for $\eta_{2r} = \zeta_{2r} + \zeta_{2r}^{-1}$ and $\eta'_{2r} = \zeta'_{2r} + \zeta_{2r}^{'-1}$, we have

$$t^2 \equiv (\eta_{2r} + \eta'_{2r})^2 p \bmod \ell,$$

and

$$s \equiv \pm \eta_{2r} \eta'_{2r} p \bmod \ell.$$

Here we define a 'primitive ℓ^{th} root of unity' to be 1.

Hence, our point counting algorithm now becomes, given a curve $C/\mathbb{F}_p$ with real multiplication by K_0 such that the only roots of unity in the endomorphism algebra are ± 1, with $C : y^2 = f(x)$,

1. Compute $j_1(f), j_2(f), j_3(f)$.
2. Compute t^2 and $s \bmod \ell$ for many small ℓ using the theorem above.
3. Find t^2 and s using the Chinese Remainder Theorem, and the bounds on s and t.
4. Check the sign of t with your favourite method (eg. multiplying a random $\mathbb{F}_p$ point in $\mathcal{J}(C)$ by the 2 options for $\#\mathcal{J}(\mathbb{F}_p)$).

References

[BGLMMST] Ballentine, Guillevic, Lorenzo-Garcia, Martindale, Massierer, Smith, and Top, *Isogenies for point counting on genus 2 hyperelliptic curves with maximal real multiplication* `https://arxiv.org/abs/1701.01927` (2017)

[BJW] Brooks, Jetchev, and Wesolowski *Isogeny graphs of ordinary abelian varieties* `https://arxiv.org/abs/1609.09793` (2016)

[DH] Dipippo, Howe, *Real polynomials with all roots on the unit circle and abelian varieties over finite fields*, J. Number Theory 73 (1998), 426450; Corrig., J. Number Theory 83 (2000), 182.

[IT] Ionica and Thomé, *Isogeny graphs with maximal real multiplication*, `https://eprint.iacr.org/2014/230` (2014)

[Mar] Martindale, *Isogeny Graphs, Modular Polynomials, and Applications*, PhD thesis (in preparation), available at `www.martindale.info` (2017)

[Mil] Milio, *A quasi-linear time algorithm for computing modular polynomials in dimension 2*, `https://arxiv.org/abs/1411.0409` (2014)

[Oor] Oort, *Abelian Varieties over Finite Fields*, `http://www.math.nyu.edu/~tschinke/books/finite-fields/final/05_oort.pdf` (2007).

[Sch] Schoof, *Counting points on elliptic curves over finite fields*, Journal de Théorie des Nombres de Bordeaux 7 (1995), 219-254.