involve
a journal of mathematics

Numerical evidence on the uniform distribution of power residues for elliptic curves

Jeffrey Hatley and Amanda Hittson

mathematical sciences publishers

2009 Vol. 2, No. 3

INVOLVE 2:3(2009)

Numerical evidence on the uniform distribution of power residues for elliptic curves

Jeffrey Hatley and Amanda Hittson

(Communicated by Nigel Boston)

Elliptic curves are fascinating mathematical objects which occupy the intersection of number theory, algebra, and geometry. An elliptic curve is an algebraic variety upon which an abelian group structure can be imposed. By considering the ring of endomorphisms of an elliptic curve, a property called complex multiplication may be defined, which some elliptic curves possess while others do not. Given an elliptic curve E and a prime p, denote by N_p the number of points on E over the finite field $\mathbb{F}_p$. It has been conjectured that given an elliptic curve E without complex multiplication and any modulus M, the primes for which N_p is a square modulo p are uniformly distributed among the residue classes modulo M. This paper offers numerical evidence in support of this conjecture.

1. Introduction

Let $F(x, y) = y^2 - a_1xy - a_3y - x^3 - a_2x^2 + a_4x - a_6$ where the $a_i \in \mathbb{C}$. Consider the set of points

$$E = \{(x, y) \in \mathbb{C}^2 : F(x, y) = 0\}.$$

We also wish to associate with the set E a special point $\mathbb{O} \in E$, called the *point at infinity*, an idea which is made rigorous by projective geometry and whose existence is justified below. Provided there are no points (x, y) such that

$$\left.\frac{\partial F}{\partial x}\right|_{(x,y)} = \left.\frac{\partial F}{\partial y}\right|_{(x,y)} = 0,$$

we say that F is *nonsingular*, and when F is nonsingular, we call E an *elliptic curve*. Through some substitutions, the equations defining elliptic curves can be put into the form $F(x, y) = y^2 - x^3 - ax - b$ for some $a, b \in \mathbb{C}$, which is called *Weierstrass normal form*.

MSC2000: 11Y99.

Keywords: elliptic curve, power residue, uniform distribution, computational algebraic geometry.

This research was sponsored by grant number NSF-DMS 0552577 and the REU program, while the authors were undergraduates at the College of New Jersey (Hatley) and Bryn Mawr College (Hittson).

If one were to graph the real points $(x, y) \in \mathbb{R}^2$ of E, the graph would form an ordinary-looking plane curve which is symmetric about the x-axis:

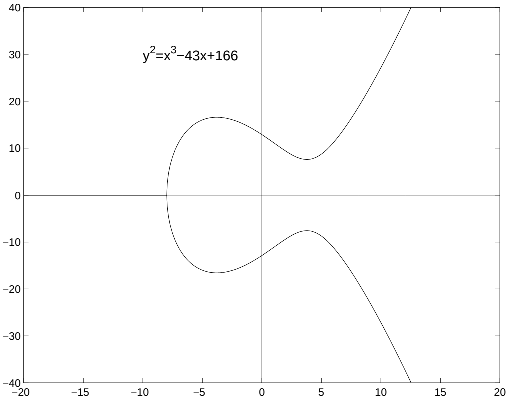

However, elliptic curves are far from ordinary and are the focus of much research due to some remarkable properties they possess. In particular, an addition law can be defined for points on the curve under which the points form an abelian group with identity element $\mathbb{O}$. The addition law can be described in the following geometric way: given two points on the curve, p_1 and p_2, draw the line connecting these two points:

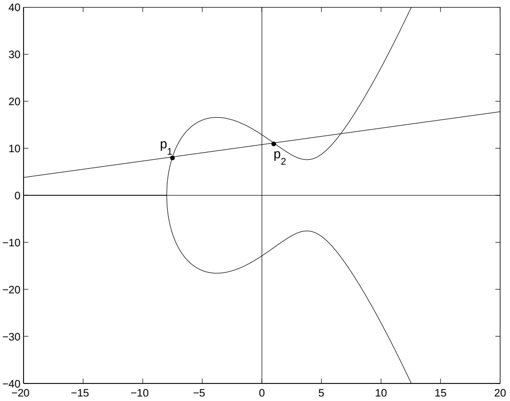

We define $p_3 = p_1 * p_2$ to be the third point of intersection between this line and the curve E. (If $p_1 = p_2$, we take the tangent line to the curve E at that point, which exists since E is nonsingular.)

This third point is guaranteed to exist by the following theorem.

Theorem (Bézout). *Let C and D be two curves in the projective space $\mathbb{CP}_2$ of degrees n and m, respectively, which have no common component* (*that is, there do not exist curves C_1, D_1, and E of degrees at least* 1 *such that $C = C_1 \cup E$ and*

$D = D_1 \cup E$). *Then C and D have precisely nm points of intersection counting multiplicities.*

This theorem is about curves in projective space; projective space and its role in the study of elliptic curves is discussed below. For a detailed discussion of projective space, multiplicity, and Bézout's theorem, see [Kirwan 1992].

Bézout's theorem states that two algebraic curves of degrees n and m intersect in exactly nm points (counting multiplicity), provided the curves do not share a common component. Now, the degree of a curve is simply the degree of the homogeneous polynomial defining it. (Homogeneous polynomials are discussed below.) Since E is defined by a polynomial of degree 3, E is of degree 3, and since a line is of degree 1, Bézout's theorem implies that the two curves should intersect in exactly 3 points; hence precisely one such point p_3 is guaranteed to exist. Then $p_1 + p_2$ is defined to be the reflection of p_3 about the x-axis:

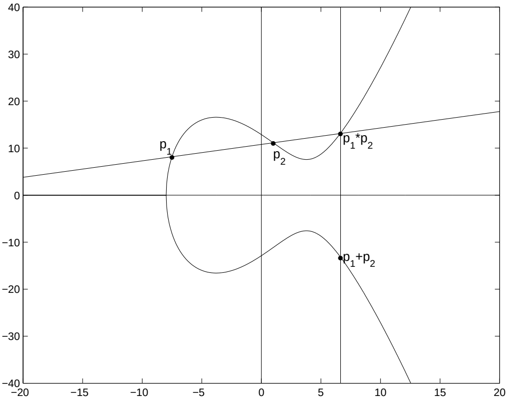

This method of addition is frequently referred to as the *chord and tangent* method. This addition law can be stated succinctly in the following way: three points on E sum to the identity, $\mathcal{O}$, if and only if they are collinear. With this formulation, and the understanding that $\mathcal{O}$ exists "at the top of the y-axis" as discussed below, we see that $p_1 * p_2 = -(p_1 + p_2)$, and the reflection of this point across the y-axis, which we defined to be $p_1 + p_2$, is the third point of E on the line between $p_1 * p_2$ and $\mathcal{O}$.

It is now necessary to consider $\mathcal{O}$, the point at infinity. For our purposes, it suffices to think of $\mathcal{O}$ as the point at which all vertical lines in the $x - y$ plane intersect. Consider the line connecting $p_1 * p_2$ and $p_1 + p_2$. This is a vertical line, and clearly only intersects E twice. However, Bézout's Theorem assures us that there are three points of intersection. In this case, that third point is $\mathcal{O}$.

To make this more rigorous, we consider the homogenized form of the curve E defined by

$$E : F(x, y) = y^2 - x^3 - ax - b.$$

A *homogeneous curve of degree n* defined by an equation in three variables x, y, z is one in which, for every monomial $ax^i y^j z^k$, we have $i+j+k=n$. To homogenize F, we make the substitutions $x = X/Z$ and $y = Y/Z$ and multiply F by Z^3, obtaining the curve

$$\bar{E} : F(X, Y, Z) = Y^2 Z - X^3 - aXZ^2 - bZ^3.$$

Furthermore, we consider the curve in the *complex projective space of degree two*, denoted $\mathbb{CP}_2$, where two 3-tuples $[x, y, z]$ and $[x', y', z']$ in $\mathbb{CP}_2$ are considered equivalent if they differ by a constant multiple; that is, we have the equivalence relation

$$[x, y, z] \sim [x', y', z'] \Longleftrightarrow [x, y, z] = \lambda[x', y', z']$$

for some complex, nonzero constant $\lambda \neq 0$. We do not consider $[0, 0, 0]$ to be an element of $\mathbb{CP}_2$. To summarize, we have

$$\mathbb{CP}_2 = \big\{[x, y, z] : x, y, z \in \mathbb{C} - \{[0, 0, 0]\}\big\} \big/ \sim .$$

Because of this equivalence, we see that as long as $Z \neq 0$, we may divide by Z and obtain our original curve, since in projective space

$$[X, Y, Z] = \frac{1}{Z}[X, Y, Z] = [x, y, 1],$$

and plugging this point into our homogeneous equation yields

$$E : F(X, Y, 1) = y^2 - x^3 - ax - b.$$

If $Z = 0$, then we have

$$F(X, Y, 0) = X^3.$$

Since points on the curve are those for which $F(X, Y, Z) = 0$, we must have $X = 0$, and Y is free to take any nonzero value, hence we obtain the homogeneous point $[0, Y, 0] = [0, 1, 0]$ corresponding to the line $Z = 0$; this is the point at infinity, $\mathbb{O}$. In projective space, all curves intersect, including parallel lines, which intersect at $\mathbb{O}$. Bézout's theorem, stated above, is actually a statement about projective space, but we extend its implications to $\mathbb{C}^2$ using $\mathbb{O}$.

As an illustration, note that if we have two parallel projective lines, $y = \alpha x + \beta_1 z$ and $y = \alpha x + \beta_2 z$, where $\beta_1 \neq \beta_2$, then to find their intersection, we solve these two equations simultaneously. These lines coincide when $z = 0$, so just as we claimed, these lines intersect at the point at infinity.

To make the addition law rigorous, let us now describe it algebraically. We wish to find the formula for $p_1 + p_2$, the sum any of two points on our curve, E, which is given by the equation $y^2 = x^3 + ax + b$. There are several cases to consider.

CASE 1: Let $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$, and suppose $p_1 \neq p_2$ and neither point is equal to $\mathbb{O}$. To find $p_1 + p_2$, we must first find $p_1 * p_2 = p_3 = (x_3, y_3)$, the third point of intersection of the line between p_1 and p_2 with E. The line between p_1 and p_2 is given by

$$y = \lambda x + \nu, \quad \text{where } \lambda = \frac{y_2 - y_1}{x_2 - x_1} \quad \text{and } \nu = y_1 - \lambda x_1.$$

Now, we wish to find the points where this line intersects our curve, so we make the following substitution:

$$y^2 = (\lambda x + \nu)^2 = x^3 + ax + b.$$

Subtracting gives us

$$0 = x^3 - \lambda^2 x^2 + (a - 2\lambda)x + (b - \nu^2).$$

The x-coordinates of the three points of intersection of the line and E are given by the roots of this equation. Factoring, we obtain

$$x^3 - \lambda^2 x^2 + (a - 2\lambda)x + (b - \nu^2) = (x - x_1)(x - x_2)(x - x_3),$$

and by identifying coefficients, we know that the sum of the roots is equal to the negative of the coefficient of x^2; that is,

$$x_1 + x_2 + x_3 = \lambda^2 \Rightarrow x_3 = \lambda^2 - x_1 - x_2.$$

By the equation of our line, this allows us to find y_3:

$$y_3 = \lambda x_3 + \nu.$$

Thus, we have $p_3 = (x_3, y_3)$. Now, the definition of our group law tells us that since p_1, p_2, and p_3 are collinear, their sum must be equal to the group identity, which is $\mathbb{O}$. Thus

$$p_1 + p_2 + p_3 = \mathbb{O}.$$

But this implies that $p_3 = -(p_1 + p_2)$. By the definition of the inverse of a group element, we know that

$$(p_1 + p_2) + p_3 = (p_1 + p_2) - (p_1 + p_2) = \mathbb{O}.$$

Using the definition of our group law again, as well as the fact that $\mathbb{O}$ is the identity element, we see that

$$(p_1 + p_2) + p_3 + \mathbb{O} = \mathbb{O},$$

which, geometrically, means that $p_1 + p_2$ is the third point of intersection of E with the line between p_3 and $\mathbb{O}$; but by our definition of $\mathbb{O}$, this is just a vertical line, and since E is symmetric about the x-axis, we conclude that p_1+p_2 is simply the reflection of p_3 about the x-axis. Thus,

$$p_1+p_2=-p_3=(x_3,-y_3),$$

where x_3 and y_3 are given by

$$x_3=\lambda^2-x_1-x_2, \quad y_3=\lambda x_3+\nu. \tag{1}$$

CASE 2: Let $p_1=p_2=(x_1,y_1)\neq\mathbb{O}$. Then the line "connecting" p_1 and p_2 is simply the line tangent to E at p_1. Since E is given by

$$y^2=x^3+ax+b,$$

implicit differentiation yields

$$\frac{\partial y}{\partial x}=\frac{3x^2-a}{2y}.$$

Then the formulas in (1) hold, by the exact same arguments, with $\lambda=\partial y/\partial x$. Note that λ blows up if $y=0$; that is, the tangent line is vertical. So if $p=(x,0)$, then $p+p=2p=\mathbb{O}$. Thus, the points of order two on E are precisely those with y-coordinate equal to zero.

CASE 3: If $p_1=\mathbb{O}$, then since $\mathbb{O}$ is the identity element of the group, we have

$$p_1+p_2=\mathbb{O}+p_2=p_2.$$

As an example, consider the curve E defined by

$$E=\{(x,y)\in\mathbb{C}^2 : y^2=x^3+17\}.$$

It is easy to check that the points $p_1=(2,5)$ and $p_2=(-1,4)$ are on the curve. Applying the addition formulas given above, we see that $\lambda=\frac{1}{3}$, $\nu=\frac{13}{3}$, and $p_1+p_2=(-\frac{8}{9},\frac{109}{27})$, which is also easily verified to be on the curve.

To summarize, under the chord-and-tangent addition law and the inclusion of $\mathbb{O}$, E forms an abelian group with identity element $\mathbb{O}$, where $p_1+p_2+p_3=\mathbb{O}$ if and only if $p_1,p_2,p_3\in E$ are the three points of intersection of some line with E. A wonderful introduction to elliptic curves can be found in [Silverman and Tate 1992].

Frequently, one prefers to look at the set

$$E(\mathbb{Q})=\{(x,y)\in\mathbb{Q}^2 : y^2=x^3+ax+b \text{ with } a,b\in\mathbb{Q}\}\cup\{\mathbb{O}\}$$

of rational points on the curve. Here again, $\mathbb{O}$ is the point at infinity with projective coordinates $\mathbb{O}=[0,1,0]$, and under the chord-and-tangent method of addition, the algebraic description of which is given in (1), $E(\mathbb{Q})$ forms an abelian group with identity element $\mathbb{O}$. In fact, this is a subgroup of the original group E, and interestingly, the Mordell–Weil theorem tells us that it is finitely-generated. More generally, given a field K, one might look at the group

$$E(K) = \{(x, y) \in K^2 : y^2 = x^3 + ax + b \text{ with } a, b \in K\} \cup \{\mathbb{O}\}$$

of K-*rational* points on the curve. Once again, $\mathbb{O} = [0, 1, 0]$ is the point at infinity. If $K = \mathbb{F}_p$ is the finite field with p elements, where p is an odd prime, then $E(K)$ is called the *reduction modulo p of E*. An examination of our derivation of the algebraic formulas for the addition law, given in (1), reveals that the formulas hold in any field, provided the field has characteristic other than 2. Under this addition law, $E(K)$ forms an abelian group with identity element $\mathbb{O}$.

Let p be an odd prime. Given an elliptic curve E reduced modulo p, it is common to ask how many points N_p lie on the curve (equivalently, what the order is of the group given by the curve). Clearly this number is finite, since for each of the p possible values of x there are only two possible values of y, plus the point at infinity $\mathbb{O}$; hence there are at most $2p + 1$ points on the curve. A better estimate for N_p might be derived in the following way. An element a of $\mathbb{F}_p$ is said to be a *quadratic residue* if there exists a nonzero $b \in \mathbb{F}_p$ such that $b^2 \equiv a \pmod{p}$. In $\mathbb{F}_p$, there are exactly $(p-1)/2$ quadratic residues. Finding points (x, y) on E amounts to finding those values of x such that $x^3 + ax + b$ is a quadratic residue modulo p; hence we might expect $x^3 + ax + b$ to be a square modulo p about half of the time. Since each such square yields the two pairs (x, y) and $(x, -y)$, we should expect about $p - 1$ such points. We might also have $x^3 + ax + b = 0$, in which case we get the point $(x, 0)$. Finally, there is the point at infinity. Adding these points up, we get an estimate of $N_p = p + 1$. Of course, this is a heuristic argument, so we should expect an error term. A theorem due to Hasse and Weil bounds this error term:

Theorem (Hasse, Weil). *If E is an elliptic curve defined over the finite field $\mathbb{F}_p$, then the number of points on E with coordinates in $\mathbb{F}_p$ is $p + 1 - a_p$, where the "error term" a_p satisfies $|a_p| \le 2\sqrt{p}$.*

Returning to our previous example, let us look at the curve

$$\bar{E} = \{(x, y) \in \mathbb{F}_5 : y^2 = x^3 + 2\}.$$

In this case, brute force suffices to show that

$$E = \{(2, 0), (3, 2), (3, 3), (4, 1), (4, 4), \mathbb{O}\},$$

hence $N_p = 6$. Since $p + 1 = 6$, the error term from the Hasse–Weil theorem is in this case $a_p = 0$, and our heuristic argument gave us N_p exactly. This does not happen in general, however.

In this paper, we are particularly interested in the set

$$Q_E = \{\text{odd primes } p \in \mathbb{Z} : N_p \text{ is a quadratic residue modulo } p\}.$$

We have just shown that for the elliptic curve E which we have been considering, $N_5 = 6 \equiv 1 \equiv 1^2 \pmod 5$, so $5 \in Q_E$. Note that, in general, for two different curves E_1 and E_2 we have $Q_{E_1} \neq Q_{E_2}$.

Recall that an endomorphism of a group G is a homomorphism $\phi : G \to G$. Now, since an elliptic curve $E/\mathbb{C}$ defined over $\mathbb{C}$ forms a group, it is natural to study $\operatorname{End}(E)$, the ring of endormorphisms of E. (To be precise, we actually only look at rational endomorphisms, those which are defined by rational functions with entries in $\mathbb{C}$. These are also called *isogenies*.) For each integer n, the multiplication-by-n map $\phi_n : E \to E$ defined by $\phi_n(x, y) = n(x, y)$ (where $n(x, y)$ represents repeated chord-and-tangent addition) defines an endomorphism of E, hence $\phi_n \in \operatorname{End}(E)$. For most curves, these are the only endomorphisms; however, some curves do have additional endormorphisms, and these curves are said to have *complex multiplication*, or simply CM. Curves for which $\operatorname{End}(E) \cong \mathbb{Z}$ are said to be non-CM.

Returning to our example curve E defined by the polynomial

$$y^2 = x^3 + 17,$$

let $\phi : E \to E$ be the homomorphism defined by

$$\phi(x, y) = \left(\frac{-1+\sqrt{-3}}{2}x, -y\right).$$

This is not a multiplication-by-n map, and since $(\frac{-1+\sqrt{-3}}{2})^3 = 1$ and $(-y)^2 = y$, if $(x, y) \in E$ then also $\phi(x, y) \in E$; hence E has CM.

Conjecture. *Let E be a non-CM curve, fix a modulus M, and let $r_1, \dots, r_s$ denote the residue classes modulo M such that $\gcd(r_i, M) = 1$ for each i. Now, look at the reduction of E modulo p for each p in the set $P_n = \{3, 5, \dots, p_n\}$ of the first n odd primes, calculate N_p for each p, and let Q_E be defined as before. Let*

$$R_i = \{p \in Q_E \cap P_n \mid p \equiv r_i \pmod M\}.$$

Let $\#R_i$ denote the cardinality of this set. Then the residues of the elements of Q_E modulo M are uniformly distributed among the r_i; that is, for every $1 \le i, k \le s$ we have

$$\lim_{n\to\infty} \frac{\#R_i}{\#R_j} = 1.$$

This conjecture, suggested to Martin by Ramakrishna, is based on [Weston 2005], which investigates a similar problem involving the distribution of power residues for $a_p \equiv N_p - 1 \pmod p$.

To test this conjecture, we wrote a program using William Stein's project Sage which takes as input an elliptic curve E, a range of moduli M, and a large number B. For each modulus M, the program looks at the reduction of E modulo p for every prime $p < B$ and computes the number of points N_p on the reduced elliptic curve. Finally, it takes those p such that N_p is a quadratic residue modulo p and looks at their residues modulo M. The output data consists of 1) the number of primes congruent to r_i for each residue class r_i relatively prime to M, and 2) the maximum percent deviation of the size of each $\#R_i$ from the expected size (were the N_p distributed uniformly among the r_i).

In this paper, we present data from our program for several curves without CM, arguing that the trends in the data as B increases strongly suggest the truth of the conjecture.

2. The program

2.1. ***What the programs do.*** To get the data we needed, we wrote a program that takes an elliptic curve and a few other parameters, computes a subset of the prime numbers, called Q_E, and, given a modulus M, computes a count for each residue. The count is the number of elements in Q_E that, when reduced modulo M, have that residue. The program also computes the largest percent difference from the expected value.

More specifically, given an elliptic curve E, a range of prime numbers P, and a range of moduli M, the program computes the set

$$Q_E = \{p \in P \mid N_p \text{ is a quadratic residue modulo } p\}.$$

Then for all $m \in M$, the program computes for each residue r,

$$\#R = \#\{p \in P \mid p \equiv r \pmod{m}\}.$$

This information is written to a file. Then the program computes some statistical information. First, it computes the expected count for each residue. If a residue, r, is not relatively prime to the modulus, m, then $p \in Q_E$ will have $p \equiv r \pmod{m} \Leftrightarrow r \in Q_E$. This means that $\#R = 0, 1$ for all residues, r, that are not relatively prime to m. So we are really only interested in residues relatively prime to the modulus. If the conjecture were true for non-CM elliptic curves, we would expect that

$$\#R_i = \frac{\#Q_E}{\varphi(m)},$$

where $\varphi(m) = \#\{r \in \mathbb{Z} \mid 0 \le r < m \text{ and } \gcd(r, m) = 1\}$ denotes the Euler-phi function and $r_1, \dots, r_s$ are the residues relatively prime to m. For each modulus, m, the program computes the percent difference of the actual count, $\#R_i$, from the expected count, $C = (\#Q_E)/\varphi(m)$. So the percent difference for each residue is

$|1-\#R_i/C|\cdot 100$. The program then picks out the largest percent difference. We are interested in the largest percent difference because it tells us how far off the actual count is from the expected count. Finally, the program writes to a file each modulus and its corresponding largest percent difference.

For example, consider the elliptic curve E defined by

$$F(x,y)=y^2+xy+y-x^3-4x+6.$$

Then given

$$P=\{p\in\mathbb{Z}^+ \mid p<10^6,\ p \text{ prime}\},$$
$$M=\{m\in\mathbb{Z}\mid 3\le m<301\},\ \text{and}$$
$$Q_E=\{p\in P\mid N_p \text{ is a quadratic residue modulo } p\},$$

the program computed that $\#Q_E=40593$. Now consider a specific modulus, say $m=9$. Then the residues relatively prime to m are $r_1=1, r_2=2, r_3=4, r_4=5, r_5=7, r_6=8$, so $\varphi(m)=6$. Thus the expected count for each residue is

$$C=\frac{\#Q_E}{\varphi(m)}=\frac{40593}{6}=6765.5.$$

In fact, the program computed that the actual counts are

$$\#R_1=6551,\quad \#R_2=6876,\quad \#R_3=6802,$$
$$\#R_4=6850,\quad \#R_5=6632,\quad \#R_6=6882.$$

So there are 6551 primes p less than one million such that N_p is a quadratic residue modulo p and $p\equiv 1 \pmod 9$. In fact, of these six counts, $\#R_1$ is the farthest off from the expected count $C=6765.5$. So $\#R_1$ will yield the biggest percent difference, which is

$$\left|1-\frac{\#R_1}{C}\right|\cdot 100\approx 3.17.$$

2.2. ***Outline of programs and efficiency.*** We organized the programs to try to maximize efficiency. The original programs we wrote were slow. Using them, we could not have computed nearly the same amount of data that we did with our newer version. By timing the original programs, we were able to see that the process that took the most time was generating and storing the set Q_E. Since this set was not even an output of our programs, we decided to generate one part of Q_E at a time. It worked like this:

```
while counter < limit:
    compute part of $Q_E$ beginning with counter
    read in previously computed data (if any)
    compute counts for this part of $Q_E$
```

```
        save these data
        delete current part of $Q_E$
compute percentages for all data
```

Note that counter and limit are simply variables to keep track of which piece of Q_E has been computed.

This while loop lends itself to division into two programs: 1) a program to return a part of the set Q_E and 2) a program to maintain a count of the residue classes in the set Q_E over the various moduli in some set M. After running initial tests on various curves for primes strictly less than one hundred thousand, which ran in under five minutes, and then for primes strictly less than one million, which ran for more than two hours, we realized that the second run of these tests was recomputing data. This led us to modify the second program to look for output files generated by previous runs of the program and start with an updated count and then proceed from there. This also meant that if the program crashed in the middle of running, we wouldn't lose all previous data. Finally, there are several driver programs that run the tests for various elliptic curves, various ranges of moduli and various ranges of prime numbers.

2.3. ***How the main programs work.*** The main program, called residueCounter, is the second program described in the preceding paragraph. It takes as parameters a starting number and upper bound to specify which range of prime integers to look at, a starting modulus and ending modulus to specify which range of moduli to test, a list of five integers to specify an elliptic curve, and finally a boolean to specify whether or not to look for files containing data that this program can use. The program stores all of the output in a dictionary, called dataDict. The keys of the dictionary are the moduli in the range specified by the parameters. Given a modulus, m, dataDict[m] evaluates to a list of lists. This list is a count for each residue of m.

The program starts by deciding whether or not to look for old files based on the boolean passed as a parameter.

```
if True:
        find all files generated by previous runs of residueCounter
        pick the most relevant file
        initialize dataDict to include all of the data computed from this file
        exclude the primes already checked by this file
if False:
        initialize a blank dataDict
```

Note that the most relevant file is the file whose range of moduli match that of the current program and of the files whose ranges match; the one with the highest upper bound has the most data and hence will save the most time.

The program needs to find the set Q_E and update dataDict to include the count from this set. However, as mentioned in the previous section, the program runs too slowly to do this all at once. So, instead it runs a loop that calls a different function, the first one mentioned above, to return a piece of Q_E, update dataDict to get all the counts for this piece, and then repeat this over and over until dataDict has all the data needed.

More precisely, the program has a variable current prime to keep track of what part of Q_E has been retrieved. It then runs the following while loop:

```
while current prime is strictly less than the upper bound:
    call function
    this function returns a new current prime and a piece of $Q_E$
    update dataDict for this piece of $Q_E$
    get rid of that piece of $Q_E$ to clear up memory space
    try to run loop again
```

The program then calculates the largest percent difference for each modulus, as described in Section 2.1. Finally, it writes this information to two different files, one for the data, one for the statistical information. The file names are keyed to include the elliptic curve, the upper bound, the range of moduli and what kind of file it is.

The function called in the last while loop is the one that actually deals with the elliptic curves. The elliptic curves we are looking at have the form $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_5$ for some $a_1, a_2, a_3, a_4, a_6 \in \mathbb{Z}$. This way we can look at the curves reduced modulo p for odd primes p. A standard way to reference a specific curve is by a list of the coefficients $[a_1, a_2, a_3, a_4, a_6]$.

This function takes as parameters an elliptic curve in the form $[a_1, a_2, a_3, a_4, a_6]$, the current prime, how many primes to check and the upper bound. It then calculates the conductor of the elliptic curve. Every elliptic curve has a unique integer associated with it, called the conductor. Essentially, the conductor tells you which primes to avoid; every prime divisor of the conductor has what is known as ‘bad reduction’, meaning that the reduction of E modulo these primes is singular and therefore not an elliptic curve. Accordingly, this function skips all of these primes. Also, we are only interested in primes p such that N_p is a quadratic residue modulo p. Luckily, there is a quick way to find out if N_p is a quadratic residue. A result due to Euler says that a necessary and sufficient condition for $a \in \mathbb{F}_p$ to be a quadratic residue is:

$$a^{(p-1)/2} \equiv 1 \pmod{p}.$$

The numbers N_p are calculated in Sage via the command `cardinality` using either the Schoof–Elkies–Atkins algorithm or the baby-step-giant-step algorithm of Mestre and Shanks. Sage decides which algorithm to use by heuristically determining which will be more computationally efficient.

That said, here is a sketch of the program:

```
compute conductor of elliptic curve
initialize blank list to store primes of interest
while (current prime < upper bound) and (counter < how many primes to check):
    if current prime doesn't divide conductor:
        if the coefficients define an elliptic curve:
            if $N_p^{(p-1)/2} \equiv 1 \pmod{p}$:
                add current prime to list
set current prime to the next prime
```

3. The data

Let us recall the conjecture on which we are working.

Conjecture. *Given a non-CM elliptic curve E, a modulus M, and a list Q_E of all the primes p for which the number of points N_p on the reduction of E modulo p are quadratic residues modulo p, the elements of Q_E are uniformly distributed among the residue classes of M.*

Due to the nature of the conjecture, the data we collect is bound to have some experimental error, since we can only look at a finite subset of Q_E at a time. When this subset of Q_E is large with respect to the modulus M, we should expect less error, and when this subset is small with respect to M, we should expect greater variance in the distribution of the primes, and hence greater error.

When we ran our program, to obtain our subset of Q_E we took all of the primes in Q_E below some fixed bound B. We then looked at their distribution among the residue classes of moduli from 3 to 300. The preceding discussion suggests that by increasing the size of B, we should expect to see our experimental error decrease. Consider Figure 1, which shows the data for $F(x, y) = y^2 + y - x^3 + x$, a curve lacking complex multiplication. The bars in the graph indicate the maximum

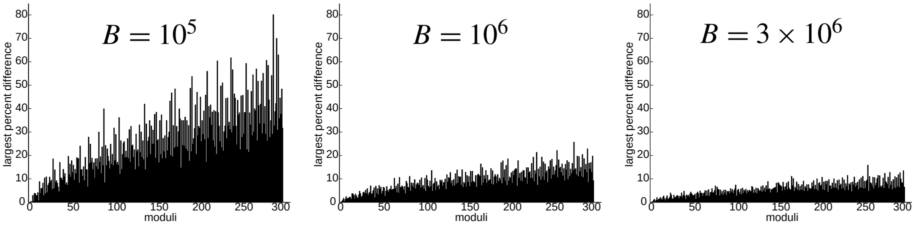

Figure 1. Prime distribution on the non-CM curve $y^2 + y = x^3 - x$ of conductor 37 and rank 1, for different values of the bound B.

percent deviation from the expected (uniform) distribution among the moduli; a higher bar indicates higher deviation. The first important thing to note is that, in all three figures, we can see that the error does indeed increase as the modulus increases. The second thing to note is that as we increase B (and thus the size of our subset of Q_E), the error decreases for every modulus. The three parts of the figure show the error when $B = 10^5$, 10^6, and 3×10^6. As B increases, the graphs "flatten out", indicating a decrease in the error. This suggests that the deviation is simply "experimental error".

We tested our program on several non-CM curves of varying ranks and conductors. The data for some of these curves is presented on this and the next two pages. The data for all of the curves tested strongly suggest that the conjecture is correct. However, we have no proof of it at this time.

All of our data and program code is available upon request.

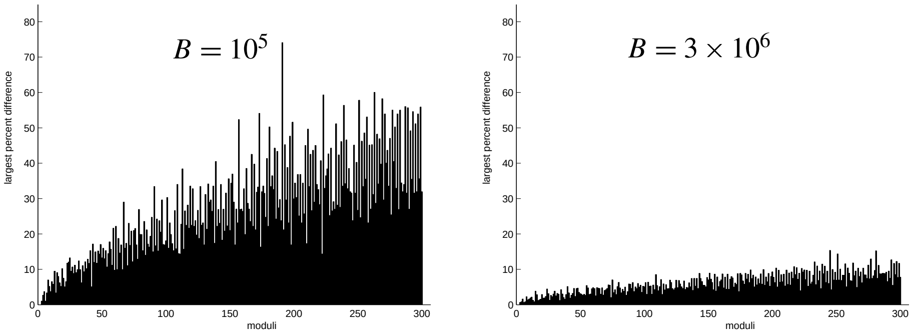

$$y^2 + y = x^3 - x^2 - 10x - 20$$

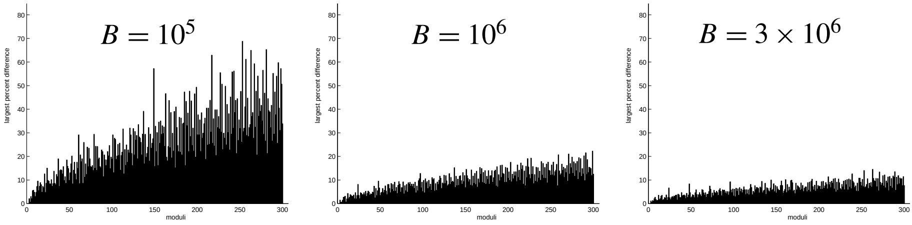

$$y^2 + xy + y = x^3 + 4x - 6$$

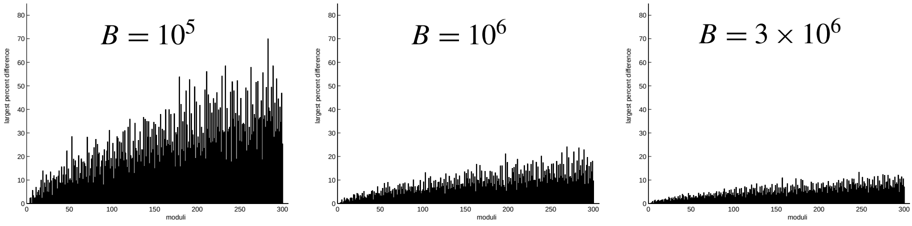

$$y^2 + y = x^3 + x^2 - 2x$$

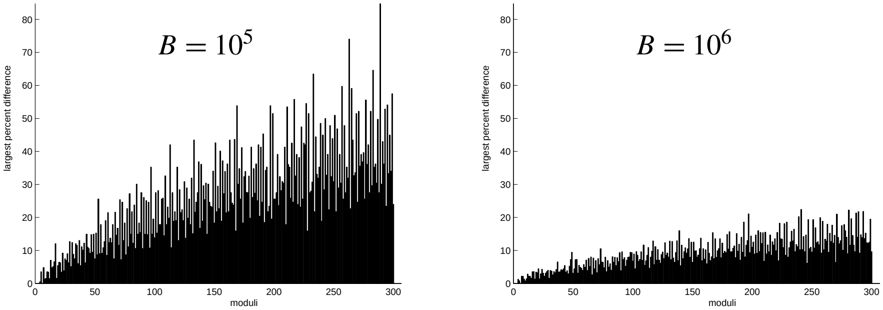

$$y^2 + y = x^3 - 7x + 6$$

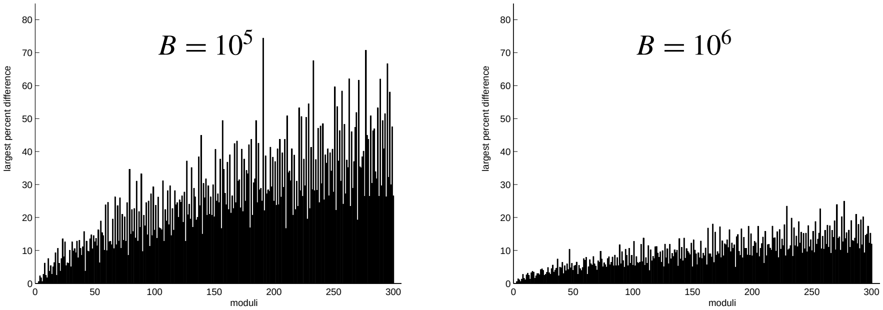

$$y^2 + xy = x^3 - x^2 - 79x + 289$$

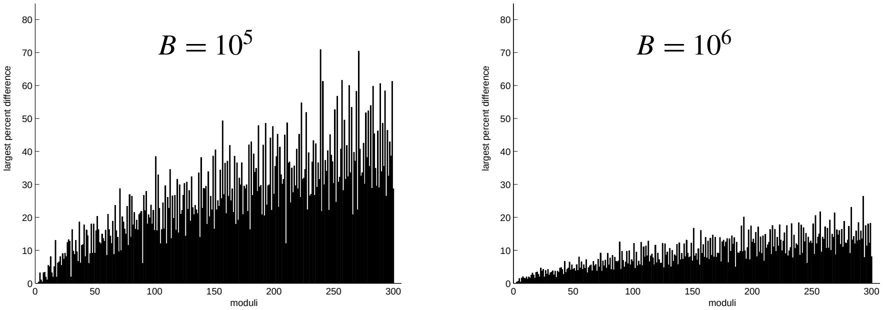

$$y^2 + y = x^3 - 79x + 342$$

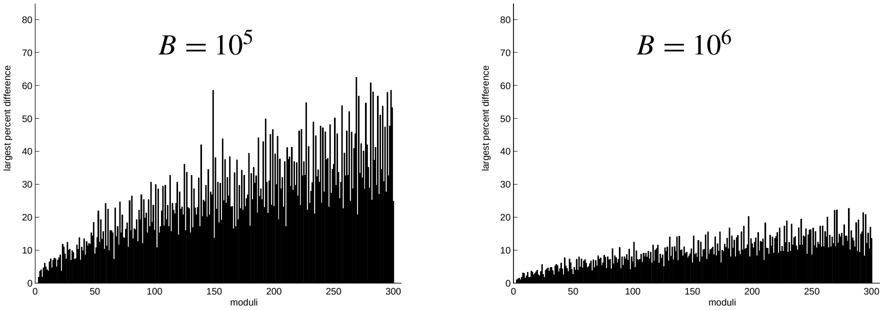

$$y^2 + xy = x^3 + x^2 - 2582x + 48720$$

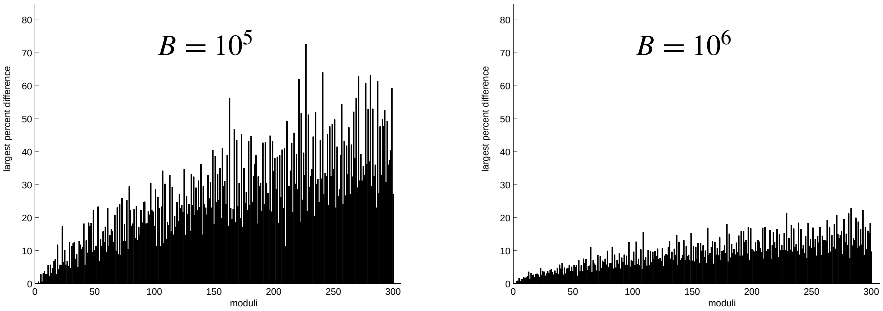

$$y^2 = x^3 - 10012x + 346900$$

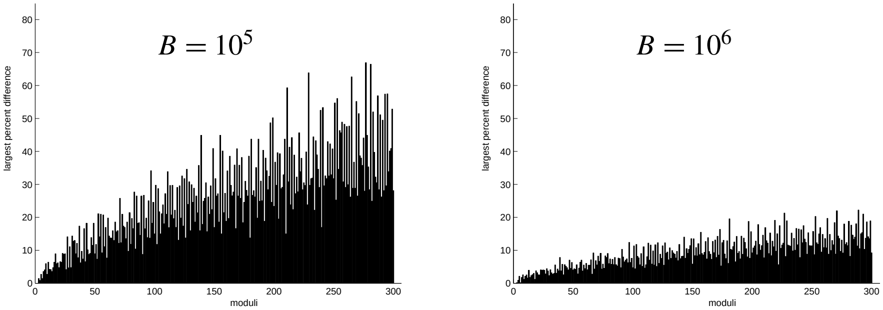

$$y^2 + y = x^3 - 23737x + 960366$$

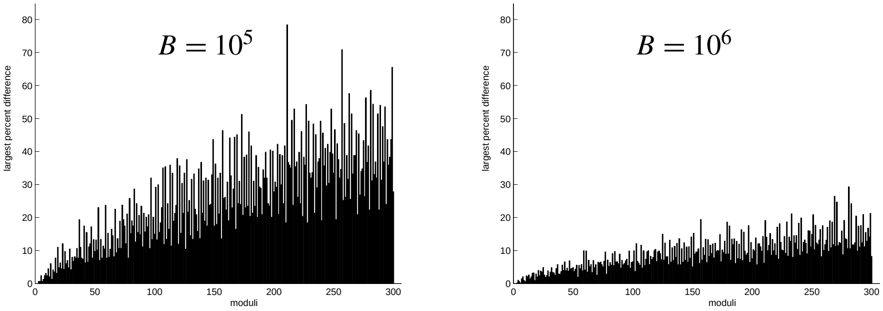

$$y^2 + y = x^3 + x^2 - 3529920x + 2567473020$$

4. Future research ideas

Several related problems might be considered in the future. It would be beneficial to find an efficient way to run a program like ours for larger primes for further data collection. The data presented in is paper is for values of B at the limit of our program's reasonable run time. We plan to modify our program to test the conjecture regarding a_p rather than N_p; furthermore, Weston [2005] conjectures that a similar result holds for primes for which the a_p are higher power residues. A slight modification of our program, or any similar program, would allow for data collection for these cases.

Acknowledgement

This work was completed at the James Madison University Summer 2008 REU Program. We thank our mentor, Jason Martin, for all his help and guidance with this project.

References

[Kirwan 1992] F. Kirwan, *Complex algebraic curves*, London Mathematical Society Student Texts **23**, Cambridge University Press, Cambridge, 1992. MR 93j:14025 Zbl 0744.14018

[Silverman and Tate 1992] J. H. Silverman and J. Tate, *Rational points on elliptic curves*, Springer, New York, 1992. MR 93g:11003 Zbl 0752.14034

[Weston 2005] T. Weston, "Power residues of Fourier coefficients of modular forms", *Canad. J. Math.* **57**:5 (2005), 1102–1120. MR 2006e:11058 Zbl 1105.11011

Received: 2008-09-21 Revised: 2009-04-01 Accepted: 2009-04-11

hatley2@tcnj.edu *Department of Mathematics, University of Massachusetts,*
Amherst, MA 01003-9305, United States

ahittson@brynmawr.edu *Department of Mathematics, University of Wisconsin,*
480 Lincoln Drive, Madison, WI 53706-1388, United States