

COMPUTING RATIONAL POINTS
ON RANK 1 ELLIPTIC CURVES
VIA L -SERIES AND CANONICAL HEIGHTS

JOSEPH H. SILVERMAN

ABSTRACT. Let E/\mathbb{Q} be an elliptic curve of rank 1. We describe an algorithm which uses the value of $L'(E, 1)$ and the theory of canonical heights to efficiently search for points in $E(\mathbb{Q})$ and $E(\mathbb{Z}_S)$. For rank 1 elliptic curves E/\mathbb{Q} of moderately large conductor (say on the order of 10^7 to 10^{10}) and with a generator having moderately large canonical height (say between 13 and 50), our algorithm is the first practical general purpose method for determining if the set $E(\mathbb{Z}_S)$ contains non-torsion points.

INTRODUCTION

Let E/\mathbb{Q} be an elliptic curve given in minimal Weierstrass form by an equation

$$(1) \quad E : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6,$$

and let $b_2, b_4, b_6, b_8, c_4, c_6, \Delta, j$ be the usual associated quantities [15, III §1]. The Mordell-Weil group $E(\mathbb{Q})$ is finitely generated and the set of S -integral points $E(\mathbb{Z}_S)$ is finite, and there is a vast literature devoted to the determination of generators for $E(\mathbb{Q})$ and elements of $E(\mathbb{Z}_S)$. In this paper we will concentrate on the case that $E(\mathbb{Q})$ has rank 1, and we will describe a new algorithm which can be used either to search for a generator of $E(\mathbb{Q})$ or to determine if $E(\mathbb{Z}_S)$ contains non-torsion points.

There are 5 general methods known for searching for points in $E(\mathbb{Q})$:

(1) Brute Force Search Algorithm. In this method one loops over $d = 1, 2, \dots$ and $a = 0, \pm 1, \pm 2, \dots$ and checks if a/d^2 is the x -coordinate of a point in $E(\mathbb{Q})$. If $E(\mathbb{Q})$ contains a rational point with (logarithmic) height $\log D$, then the running time is $O(D^{3/2})$.

(2) Sieve Assisted Search Algorithm. For each d , one uses congruence conditions to eliminate many of the potential a values. The running time is still $O(D^{3/2})$, but as a practical matter the run time may be reduced by a factor of 1000 or more. See [6, §3.5] for the basic idea, although the method there has been substantially improved by Cremona [7].

Received by the editor May 8, 1996 and, in revised form, March 3, 1997.
1991 *Mathematics Subject Classification.* Primary 11G05, 11Y50.
Key words and phrases. Elliptic curve, canonical height.
Research partially supported by NSF DMS-9424642.

(3) Homogeneous Space Search Algorithm. In this method one looks for homogeneous spaces C_1, \dots, C_t and covering maps $\phi_i : C_i \rightarrow E$ of degree m so that $E(\mathbb{Q})$ is equal to the union of the $\phi_i(C_i(\mathbb{Q}))$'s. The gain in the method lies in the fact that the height of the smallest point in $C_i(\mathbb{Q})$ will tend to be approximately the m^{th} root of the height of the smallest point in $E(\mathbb{Q})$, potentially a huge savings even for $m = 2, 3$, or 4 . The problem with using descent comes from the difficulty in finding the curves C_i . If $E(\mathbb{Q})$ contains a point of order 2 (or 3), it is easy to find these curves for degree 2, and feasible for degrees 4 and 8. See for example [3] and [4] for some spectacular computations on the curves $y^2 = x^3 + px$. However, if $E(\mathbb{Q})$ has no torsion, then descent will only succeed if the relevant homogeneous spaces happen to be defined by equations with comparatively small coefficients. For degree 4 descents, as described in [6], the expected search time is $O(D^{1/2})$ after the homogeneous space has been found.

(4) Heegner Point Algorithm. This method is only applicable in case $E(\mathbb{Q})$ has rank 1. Briefly, one picks out certain special points on $X_0(N)$, where N is the conductor of E , and uses the modular parametrization $X_0(N) \rightarrow E$ to obtain points in $E(\mathbb{Q})$. (See [10] for basic information about Heegner points.) The Heegner point method is completely effective in principle, but in practice the series defining the modular parametrization converges slowly if the conductor of E is large. More precisely, the Heegner point method requires computation of a series with $O(N)$ terms, so it is quite practical if N is on the order of 10^5 , but seems quite impractical in general if N is greater than 10^8 . However, we should mention that if E is a (large) twist of a small conductor curve E' , then Elkies [9] and Zagier [21] explained how to compute Heegner points quite efficiently. For example, Elkies finds a non-torsion generator on $1063y^2 = x^3 - x$ having $h(x) \approx 120$, and Liverance used Zagier's method to find a Heegner point on $x^3 + y^3 = 1354$ having height $h(x) > 3000$.

(5) Canonical Height Search Algorithm. This is the new algorithm which we will describe in detail in this paper. The algorithm can be used as a straight search algorithm, in which case it has a projected run time of $O(D + \sqrt{N})$, or it can be used to determine if the set $E(\mathbb{Z}_S)$ contains non-torsion points, in which case the run time is essentially

$$O\left(\sqrt{N} + \prod_{p \in S} \log_p D\right).$$

In particular, it takes $O(\sqrt{N})$ steps to determine if $E(\mathbb{Z})$ contains non-torsion points, a running time which is independent of the actual size D of the generator of $E(\mathbb{Q})$ and is much faster than any other known method.

The validity of the Canonical Height Search Algorithm depends on two assumptions, namely that E is modular and that $L'(E, 1)$ is equal to the Birch-Swinnerton-Dyer value up to multiplication by the square of an integer. By the work of Wiles [20], Taylor-Wiles [18], and Diamond [8], we know that most E/\mathbb{Q} are modular; and assuming the modularity, the work of Gross-Zagier [10] and Kolyvagin [11] gives the desired value of $L'(E, 1)$. We will implicitly be using these results when we apply the Canonical Height Search Algorithm to prove that $E(\mathbb{Z}_S)$ contains no non-torsion points without actually finding a generator for $E(\mathbb{Q})$. Of course, if the Canonical Height Search Algorithm produces a candidate rational point P in $E(\mathbb{Q})$,

we can verify that P is in $E(\mathbb{Q})$ directly without using any conjectures or deep theorems. Finally, we mention that the Canonical Height Search Algorithm is practical in the sense that the “big- O ” constants are reasonably small.

The basic idea underlying the Canonical Height Search Algorithm is simple to explain. First we compute the number $H = L'(E, 1)T^2/2\Omega c$, which should equal the canonical height $\hat{h}(P)$ of a rational point $P \in E(\mathbb{Q})$. Here T , c , and Ω are the usual quantities appearing in the Birch-Swinnerton-Dyer coefficient. The canonical height has a decomposition

$$\hat{h}(P) = \hat{\lambda}_\infty(P) + \log d(P) + \hat{\lambda}_{\text{bad}}(P),$$

where $x(P) = a(P)/d(P)^2$, $\hat{\lambda}_{\text{bad}}(P)$ is the contribution from the primes of bad reduction, and $\hat{\lambda}_\infty(P)$ is the archimedean local height. It turns out that for any given E , the number $\hat{\lambda}_{\text{bad}}(P)$ lies in a short list of values, say $\hat{\lambda}_{\text{bad}}(P) \in \Lambda_{\text{bad}}$. Further, the archimedean local height is a (real) analytic function $\hat{\lambda}_\infty : \mathbb{C}/L \rightarrow \mathbb{R} \cup \{\infty\}$, where the lattice L is chosen so that $E(\mathbb{C}) \cong \mathbb{C}/L$. So the algorithm proceeds as follows. For each hypothetical denominator d and each possible bad contribution $\lambda \in \Lambda_{\text{bad}}$, we set $\hat{\lambda}_\infty(z) = \hat{h}(P) - \log d - \lambda$ and “solve” for z . Of course, for any $r \in \mathbb{R}$, the inverse image $\hat{\lambda}_\infty^{-1}(r)$ consists of a curve in \mathbb{C}/L . However, we have the additional information that $E(\mathbb{Q}) \subset E(\mathbb{R})$, and $E(\mathbb{R})$ itself consists of one or two circles in $E(\mathbb{C})$. So the (real) curve $\hat{\lambda}_\infty^{-1}(r)$ will intersect $E(\mathbb{R})$ in only a few points, and these can be found explicitly. This gives a few points $(x, y) \in E(\mathbb{R})$, and we then check whether d^2x and d^3y are close to being integers. If they are, say $d^2x \approx a$ and $d^3y \approx b$, then we check whether $(a/d^2, b/d^3)$ is indeed a rational point on E . If so, we are done, and if not, then we go back and check another $\hat{\lambda}_{\text{bad}}(P)$ or another d . This concludes our brief description of how the Canonical Height Search Algorithm is used to find a generator for $E(\mathbb{Q})$.

In order to prove that (say) $E(\mathbb{Z})$ contains no non-torsion points, it suffices to use the Canonical Height Search Algorithm with $d = 1$. If no point is found, then we know that the point $P \in E(\mathbb{Q})$ with $\hat{h}(P) = H$ is not an integral point. Of course, this does not prove that $E(\mathbb{Z})$ consists only of torsion points, since P might be a multiple of a $Q \in E(\mathbb{Z})$, say $P = MQ$. So we proceed as follows. For each $m = 1, 2, \dots$, we apply the Canonical Height Search Algorithm with $d = 1$ to the height value H/m^2 . If $P = MQ$ with $Q \in E(\mathbb{Z})$, then the algorithm will eventually find Q . Otherwise, after checking all m 's up to, say, m_0 , we perform a brute force (or sieve-assisted) search on E for all integer points having canonical height less than H/m_0^2 . Even if H is very large, say between 100 and 1000, taking $m_0 \approx \sqrt{10H}$ will lead to a very small search provided the coefficients of the curve E are not too huge. A similar method can be used to verify that $E(\mathbb{Z}_S)$ consists entirely of torsion points, provided the set S of primes is not too large.

Karl Rubin [13] has described a method to construct a point $P_p \in E(\mathbb{Q}) \otimes \mathbb{Z}_p$ in the case that E has complex multiplication and $\text{ord}_{s=1} L(E/\mathbb{Q}, s) = 1$. This point is not in $E(\mathbb{Q})$, but using p -adic L -series, he is able to show how to p -adically approximate a point $Q \in E(\mathbb{Q})$. He gives an example using $y^2 = x^3 - 49x$ and constructs the point $(-49/25, -1176/125)$ by approximating it modulo the ideal $(2+i)^8\mathbb{Z}[i]$. It is not clear whether Rubin's method would be practical for computing points with large (archimedean) height of the sort considered in this paper.

Doug Ulmer [19] has given a construction in the same spirit as Rubin for the universal elliptic curve defined over a modular function field. He explicitly constructs points defined over a certain completion of the field and gives a conjectural description of how one might use these to construct rational points. He does not give any explicit examples, and again it is not clear to what extent his method is practical.

[In fairness to both Rubin and Ulmer, it should be noted that the main aim of their papers is to construct points which can be used for various theoretical purposes. Both papers contain extremely interesting results.]

The organization of this paper is as follows. In Section 1 we give a step-by-step description of the Canonical Height Search Algorithm, starting with a straight search version and following with a version to determine if $E(\mathbb{Z}_S)$ contains non-torsion points. This serves as an overview and is essentially a guide for implementation. In Section 2 we go over each step of the algorithm in more detail, providing justification, further implementation suggestions, and references. One of the steps of the algorithm requires computation of certain power series in q , where $|q| < 1$. In general, $|q|$ will be fairly small and this step will cause no problems. However, if it happens that $|j(E)| \gg 1$ and $c_6(E) < 0$, then $|q|$ may be close to 1 and the computation will not be feasible. In Section 3 we describe a modification of this step which allows q to be replaced with a small q' , albeit at the cost of introducing various additional complications. Section 4 contains two estimates needed by the algorithm, followed by a number of remarks, including a suggestion on how the Canonical Height Search Algorithm might be modified to help with certain curves of rank 2, and a brief description on how one can combine the Canonical Height and the Homogeneous Space Algorithms to produce, in principle, an algorithm with $O(D^{1/4})$ search time. Finally, in Section 5, we give numerical examples illustrating the Canonical Height Search Algorithm.

Acknowledgements. I would like to thank Horst Zimmer, whose questions in Toronto led me to reconsider the problem of efficiently computing rational points, John Cremona for sending me lists of test cases and for much helpful algorithmic advice, Noam Elkies, Hendrik Lenstra, Richard Pinch, and Nigel Smart for their helpful suggestions, and Don Zagier for explaining that the Heegner Point Method is far more practical than I had thought and suggesting that the real utility of the Canonical Height Method would be for dealing with points of small denominator. I would also like to thank John Tate for asking me, many years ago, if it might be possible to recover a rational point from its canonical height. This paper is a belated partial answer to his question.

1. STEP-BY-STEP DESCRIPTION OF THE ALGORITHM

In this section we give a detailed step-by-step description of the Canonical Height Search Algorithm. The first version we present will (eventually) find a non-torsion point in $E(\mathbb{Q})$ when $\text{rank } E(\mathbb{Q}) = 1$. The second version can be used to determine if $E(\mathbb{Z}_S)$ contains non-torsion points, again when $\text{rank } E(\mathbb{Q}) = 1$. In particular, the second version provides a method for proving that $E(\mathbb{Z}_S)$ consists entirely of torsion points (or is empty) in many situations where there are currently no other practical algorithms.

This section will merely state the main steps in the algorithms. In the next section we will discuss each of these steps, giving some theoretical justification,

references, and implementation suggestions. We assume throughout that E is given by a minimal Weierstrass equation (1).

We begin with the algorithm to compute a non-torsion point in $E(\mathbb{Q})$.

The Canonical Height Search Algorithm — Version I
Computing a Point in $E(\mathbb{Q})$ When $\text{rank } E(\mathbb{Q}) = 1$.

- (1) Compute the quantities $b_2, b_4, b_6, b_8, c_4, c_6, \Delta, j$ attached to the Weierstrass equation (1)
- (2) Compute

$$N = \text{Conductor of } E,$$

$$\varepsilon = \text{Sign of Functional Equation of } E.$$

If $\varepsilon = 1$, terminate with message “Analytic rank is even.”

- (3) Compute

$$L'(E, 1) = 2 \sum_{n \geq 1} \frac{a_n}{n} E_1 \left(\frac{2\pi n}{\sqrt{N}} \right), \quad \text{where } E_1(x) = \int_x^\infty \frac{e^{-t}}{t} dt$$

and the a_n 's are the L -series coefficients for E/\mathbb{Q} , $L(E/\mathbb{Q}, s) = \sum a_n n^{-s}$. If $L'(E, 1) \approx 0$, terminate with the message “Analytic rank ≥ 3 .”

- (4) Compute the real period $\Omega = \int_{E(\mathbb{R})} dx/(2y + a_1x + a_3)$, the order of the torsion subgroup $T = E(\mathbb{Q})_{\text{tors}}$, and the Tamagawa number $c = \prod_{p|\Delta} c_p$, where $c_p = [E(\mathbb{Q}_p) : E_0(\mathbb{Q}_p)]$.
- (5) Compute the value $H = L'(E, 1)T^2/2\Omega c$. This will be the canonical height of a Heegner point P in $E(\mathbb{Q})$.
- (6) Make a list Λ_{bad} of the possible contributions to the canonical height coming from primes of bad reduction. In practice, Λ_{bad} tends to be fairly short.
- (7) Compute a real period ω_1 and a complex period ω_2 for E , and let $\tau = \omega_2/\omega_1$ and $q = \exp(2\pi i\tau)$. These should be chosen so that $q \in \mathbb{R}$, $|q| < 1$, and

$$c_4 = \left(\frac{2\pi}{\omega_1} \right)^4 \left(1 + 240 \sum_{n \geq 1} \frac{n^3 q^n}{1 - q^n} \right) \quad \text{and} \quad c_6 = \left(\frac{2\pi}{\omega_1} \right)^6 \left(1 - 504 \sum_{n \geq 1} \frac{n^5 q^n}{1 - q^n} \right).$$

(If $c_6 < 0$, then $|q|$ may be close to 1, which will seriously impair the efficiency of the algorithm. We will discuss below a modification using twists to deal with this situation.)

- (8) Define a modified Weierstrass \wp -function

$$\mathcal{P}(z) = -\frac{b_2}{12} + \left(\frac{2\pi i}{\omega_1} \right)^2 \left(\frac{1}{12} + \frac{u}{(1-u)^2} + \sum_{n=1}^\infty q^n \left(u \left(\frac{1}{(1-q^n u)^2} + \frac{1}{(q^n - u)^2} \right) - \frac{2}{(1-q^n)^2} \right) \right),$$

where $u = \exp(2\pi iz)$.

- (9) LOOP $1 \leq d \leq d_{\text{max}}$
- (10) LOOP $\lambda \in \Lambda_{\text{bad}}$
- (11) Solve

$$S \prod_{n \geq 1} ((1 - q^n)^2 + q^n S^2) = |q|^{-1/12} \exp(\lambda + \log d - H)$$

for S with $0 \leq S \leq 2$.

(12) Compute

$$z = \frac{1}{\pi} \sin^{-1} \left(\frac{S}{2} \right),$$

$a =$ closest integer to $d^2\mathcal{P}(z)$.

(13) IF $4a^3 + b_2a^2d^2 + 2b_4ad^4 + b_6d^6$ is a perfect square, say equal to b^2 , THEN terminate with the rational point $(a/d^2, (b - a_1ad - a_3d^3)/2d^3)$ on E .

(14) IF $\Delta > 0$ ($E(\mathbb{R})$ Has Two Components)

(15) Set $Q = \sqrt{q}$ (note $\Delta > 0$ implies that $q > 0$). Solve

$$\prod_{n \geq 1} ((1 - Q^{2n-1})^2 + Q^{2n-1}S^2) = Q^{1/12} \exp(\lambda + \log d - H)$$

for S with $0 \leq S \leq 2$.

(16) Compute

$$z = \frac{1}{\pi} \sin^{-1} \left(\frac{S}{2} \right) + \frac{1}{2}\tau,$$

$a =$ closest integer to $d^2\mathcal{P}(z)$.

(17) IF $4a^3 + b_2a^2d^2 + 2b_4ad^4 + b_6d^6$ is a perfect square, say equal to b^2 , THEN terminate with the rational point $(a/d^2, (b - a_1ad - a_3d^3)/2d^3)$ on E .

(18) ENDIF ($E(\mathbb{R})$ Has Two Components)

(19) END Λ_{bad} LOOP

(20) END d LOOP

(21) Terminate with message “No rational point found.”

We now describe the modifications needed to search for non-torsion points in $E(\mathbb{Z}_S)$ for a finite set of primes S . In particular, if $S = \emptyset$, then the algorithm will determine if $E(\mathbb{Z})$ contains non-torsion points. We mention that the validity of the output of this algorithm is dependent on knowing that the elliptic curve E/\mathbb{Q} is modular. For example, this condition will be satisfied if E/\mathbb{Q} has good or multiplicative reduction at 3 and 5 (see [20], [18], [8]).

The Canonical Height Search Algorithm — Version II

Searching For Non-torsion Points in $E(\mathbb{Z}_S)$ When $\text{rank } E(\mathbb{Q}) = 1$.

(a) Do Steps (1)–(8) of the previous algorithm. In particular, H is the canonical height of a non-torsion point in $E(\mathbb{Q})$. Save this original value of H as H_0 .

(b) LOOP $1 \leq m \leq \lceil \sqrt{10H_0} \rceil$

(c) Set $H = H_0/m^2$.

(d) Compute d_{\max} according to the formula

$$\log(d_{\max}) = H + \frac{1}{24}h(j) + \frac{1}{12} \log |\Delta| + \frac{1}{12} \log^+ |j| + \frac{1}{2} \log^+ |b_2/12| + 1.32.$$

(e) LOOP $1 \leq d \leq d_{\max}$ WITH $d \in \mathbb{Z} \cap \mathbb{Z}_S^*$. (That is, d is composed of a product of primes in S .)

(f) Do Steps (10)–(19) of the previous algorithm. In particular, if $E(\mathbb{Q})$ has a point with x -denominator d^2 and canonical height H , then the algorithm will find this point and terminate.

(g) END d LOOP

(h) END m LOOP

- (i) Compute d_{\max} according to the following formula (notice the H has disappeared, so this value will generally be fairly small):

$$\log(d_{\max}) = \frac{1}{24}h(j) + \frac{1}{12} \log |\Delta| + \frac{1}{12} \log^+ |j| + \frac{1}{2} \log^+ |b_2/12| + 1.42.$$

- (j) Search for non-torsion points in $E(\mathbb{Q})$ with x -coordinate a/d^2 satisfying $|a| \leq d_{\max}^2$ and $1 \leq d \leq d_{\max}$. If d_{\max} isn't too large, this can be done with a brute-force search; otherwise it may be preferable to use a sieve search. If a non-torsion point is discovered, then return the point and terminate.
- (k) Terminate with the message “ $E(\mathbb{Z}_S)$ contains no non-torsion points.”

Remark 1.1. It is possibly worth mentioning that for any fixed bound d_0 , the Canonical Height Search Algorithm II can also be used to rapidly find all points in $E(\mathbb{Q})$ of the form $(a/d^2, b/d^3)$ with $1 \leq d \leq d_0$. One need merely replace Step (e) with a loop over $1 \leq d \leq d_0$.

2. DISCUSSION OF THE ALGORITHM

In this section we will discuss in more detail each of the main steps described in Section 1. This will include some theoretical justification, implementation remarks, and references. For the convenience of the reader, we will also give the appropriate functions in PARI [1] for computing various quantities.

Step 1. The formulas to compute $b_2, b_4, b_6, b_8, c_4, c_6, \Delta, j$ are given in [5, §7.1.3], [6, §3.1], [15, III §1]. [PARI: $b_2, b_4, b_6, b_8, c_4, c_6, \Delta, j$ are the 6th through 13th components of $\mathbf{e} = \mathbf{initell}([a_1, a_2, a_3, a_4, a_6])$.]

Step 2. The conductor N of E may be computed using Tate’s algorithm and Ogg’s formula, or via the L -series of E . The sign ε of the functional equation is most easily computed using the L -series. For Tate’s algorithm, see [5, § 7.5.1], [6, §3.2], or [16, IV §9]. For the L -series method, see [5, §7.5.3]. [PARI: $\mathbf{N} = \mathbf{globalred}(\mathbf{e})$ [1].]

Step 3. Methods to compute the a_n coefficients of the L -series

$$L(E/\mathbb{Q}, s) = \sum a_n n^{-s}$$

are given in [5, §§7.4.3, 7.5.3] and [6, §2.9]. The exponential integral $E_1(x) = \int_x^\infty dt/te^t$ can be efficiently computed as described in [5, §5.6.2] and [6, §2.13]. Note that although it may be necessary to take a large number of terms in the series for $L'(E, 1)$ in order to get (say) 50 or 100 digits of accuracy, this calculation only needs to be done once. We will give an explicit error estimate below—see Proposition 4.1. [PARI: $\mathbf{anell}(\mathbf{e}, \mathbf{n}_0)$ gives a vector with the first n_0 of the a_n ’s, while $\mathbf{ake11}(\mathbf{e}, \mathbf{n})$ gives just a_n . The exponential integral is $\mathbf{eint1}(\mathbf{x})$.]

Step 4. The real period Ω can be computed using the AGM method [5, §7.4.1, Algorithm 7.4.7], [6, §3.7]. The torsion subgroup, or more precisely its order T , can be computed using the method described in [5, §7.5.2, Algorithm 7.5.5] and [6, §3.3]. The local Tamagawa numbers c_p whose product is c are computed as a by-product of Tate’s algorithm [5, § 7.5.1], [6, §3.2], [16, IV §9]. [PARI: The real period Ω equals $\mathbf{e}[15]$ if $\Delta < 0$, and Ω equals $2\mathbf{e}[15]$ if $\Delta > 0$. The torsion subgroup has order $\mathbf{T} = \mathbf{torsell}(\mathbf{e})$ [1]. The Tamagawa number is $\mathbf{c} = \mathbf{globalred}(\mathbf{e})$ [3].]

Step 5. Set $H = L'(E, 1)T^2/2\Omega c$, using the values computed in Steps 3 and 4.

Non-archimedean Local Heights ($n = \text{ord}_p(\Delta_E)$)	
Reduction Type	$\hat{\lambda}_{p,\text{bad}}/\log(p)$
I_n	$(n/12) - (i/2) + (i^2/2n)$ for some $0 \leq i \leq n/2$
I_m^*	$n/12$ or $m/12 + \{(n-m-6)/12\}$ or $-m/24 + \{(n-m-6)/12\}$
II	$n/12$
III	$n/12$ or $\{(n-3)/12\}$
IV	$n/12$ or $\{(n-4)/12\}$
IV^*	$n/12$ or $\{(n-8)/12\}$
III^*	$n/12$ or $\{(n-9)/12\}$
II^*	$n/12$

Step 6. We need to make a list of all possible contributions to the canonical height coming from primes of bad reduction. (More precisely, we want the possible contributions other than those coming from primes dividing the denominator of the x -coordinate, since primes dividing the denominator of x will be accounted for in our choice of d .) For example, if $p|\Delta$ and if P does not reduce to the singular point modulo p , then the local height $\hat{\lambda}_p(P)$ has a $\frac{1}{12} \text{ord}_p(\Delta) \log p$ attached to it [16, VI.4.1]. If P does reduce to the singular point modulo p , then there is a small list of possibilities for $\hat{\lambda}_p(P)$ depending on the particular reduction type and the particular component of the Néron model hit by P . The reduction type for each $p|\Delta$ can be computed using Tate's algorithm [5, § 7.5.1], [6, §3.2], [16, IV §9]. (A small amount of time can be saved by observing that if $p^2 \nmid \Delta$, then the p -contribution is always $\frac{1}{12} \log p$.) Then the possible p -contributions to $\hat{\lambda}_{\text{bad}}(P)$ can be read off of Table 1, which appeared originally in [14]. (Remark. If $p \geq 5$, then the quantities in braces in Table 1 are equal to 0.) Take each of the possibilities for each p dividing Δ and add them up to create the list Λ_{bad} of possible contributions to $\hat{h}(P)$ coming from primes of bad reduction.

[PARI: The command `factor(abs(e[12]))[,1]` gives a list of primes of bad reduction. For each prime p of bad reduction, the reduction type can be computed using `t=localred(e,p)[2]`, where the values

$$t = 1, 2, 3, 4, -1, -2, -3, -4, 4 + N, -4 - N$$

correspond respectively to the reduction types

$$I_0, II, III, IV, I_0^*, II^*, III^*, IV^*, I_N, I_N^*]$$

Step 7. The periods ω_1 and ω_2 can be computed using AGM's as described in [5, §7.4.1, Algorithm 7.4.7] or [6, §3.7]. The resulting $\tau = \omega_2/\omega_1$ has real part equal to either 0 or $-1/2$. That these values are correct can be checked by comparing the series listed in Step 7 with the values of c_4 and c_6 . We observe that there are two special cases. First, if $c_6 = 0$, then $\tau = i$ (respectively $\tau = (1+i)/2$) if $c_4 < 0$ (respectively $c_4 > 0$). Second, if $c_4 = 0$, then $\tau = (1+i\sqrt{3})/2$ (respectively $\tau = 1/2 + i/2\sqrt{3}$) if $c_6 > 0$ (respectively $c_6 < 0$).

If $c_6 \geq 0$, then $|q|$ will be quite small. Precisely, if $c_6 > 0$ and $\Delta > 0$, then $0 < q < e^{-2\pi} \approx 0.001867$, while if $c_6 > 0$ and $\Delta < 0$, then $0 > q > -e^{-\sqrt{3}\pi} \approx -0.004333$.

Thus if $c_6 \geq 0$, all q -series converge rapidly and one can perform multi-precision computations easily.

However, if $c_6 < 0$ and $|j| \gg 1$, then the corresponding q value will be quite close to ± 1 , which will slow the algorithm considerably. A possible cure for this problem is to move to the \mathbb{C}/\mathbb{R} twist of the elliptic curve. This has the effect of replacing c_6 by $-c_6$, yielding a better q value. But it also has the effect of switching the real and (purely) imaginary loci, so various modifications need to be made, especially in Steps 11 and 15, as will be described in Section 3. In practice, it is probably not worth moving to the twist if $|q| < 1/2$.

[PARI: The periods ω_1 and ω_2 are $e[15]$ and $e[16]$.]

Step 8. Our modified Weierstrass function \mathcal{P} is related to the usual \wp (as defined in [5, Proposition 7.4.4], for example) by the formula

$$\mathcal{P}(z, q) = -\frac{b_2}{12} + \wp(\omega_1 z; \omega_1 \mathbb{Z} + \omega_2 \mathbb{Z}).$$

For high speed implementation, it might be better to compute the Weierstrass \wp function in terms of the Weierstrass σ function, since the σ function converges quadratically. (See [5, Remark after Algorithm 7.4.5].) [PARI: This is essentially the function `pointell`. See Step (12) below.]

Step 9. Choose a value for d_{\max} . The algorithm will search for a point in $E(\mathbb{Q})$ whose x -coordinate has the form a/d^2 for some integer $1 \leq d \leq d_{\max}$.

Step 10. Loop over the (hopefully small number of) elements in the set Λ_{bad} .

Step 11. There are two issues to discuss at this step. First, how should one solve efficiently for S ? Second, why is one solving for S ? The first is easily dealt with. Let $f(S)$ be the function

$$f(S) = S \prod_{n \geq 1} ((1 - q^n)^2 + q^n S^2).$$

Taking the logarithmic derivative, we find that

$$\begin{aligned} f'(S) &= f(S) \left\{ \frac{1}{S} + \sum_{n \geq 1} \frac{2q^n S}{(1 - q^n)^2 + q^n S^2} \right\} \\ &= \left(\prod_{n \geq 1} ((1 - q^n)^2 + q^n S^2) \right) \left(1 + 2S^2 \sum_{n \geq 1} \frac{q^n}{(1 - q^n)^2 + q^n S^2} \right). \end{aligned}$$

This formula for $f'(S)$ is rapidly converging (provided $|q|$ is small, of course), so we can solve an equation $f(S) = c$ efficiently using Newton's method. A good starting point is $S = c/(1 - q)^2$, and then iterate $S \leftarrow S - (f(S) - c)/f'(S)$. (We also note that $f'(S) > 0$ for all real $S > 0$, so there will be only one solution.)

Next we explain why the function $f(S)$ appears. What we are doing is searching for a point P whose canonical height has the known value $\hat{h}(P) = H$. This canonical height can be decomposed as

$$\hat{h}(P) = \hat{\lambda}_\infty(P) + \log d(P) + \hat{\lambda}_{\text{bad}}(P).$$

At this step we are assuming that the denominator $d(P)$ of the x -coordinate of P is equal to our loop variable d , and that the contribution $\hat{\lambda}_{\text{bad}}(P)$ coming from the bad primes is equal to the value λ chosen from Λ_{bad} . So we are assuming that

$$H = \hat{\lambda}_{\infty}(P) + \log d + \lambda,$$

where the only unknown quantity in this equation is $\hat{\lambda}_{\infty}(P)$.

Next we consider the explicit formula for the archimedean local height as a (real) analytic function $\hat{\lambda}_{\infty} : \mathbb{C}/L \rightarrow \mathbb{R} \cup \{\infty\}$. Here $L = \mathbb{Z} + \tau\mathbb{Z}$ is a lattice attached to E from Step 7, and we have fixed an isomorphism $\mathbb{C}/L \cong E(\mathbb{C})$ defined over \mathbb{R} . For any $z \in \mathbb{C}$, we write $u = e^{2\pi iz}$, and we let $B_2(t) = t^2 - t + \frac{1}{6}$ be the 2nd Bernoulli polynomial. Then $\hat{\lambda}_{\infty}$ is given by the formula

$$\hat{\lambda}(z) = -\frac{1}{2}B_2\left(\frac{\log|u|}{\log|q|}\right) \log|q| - \log|1-u| - \sum_{n \geq 1} \log|(1-q^n u)(1-q^n u^{-1})|.$$

(This formula is due to Néron and Tate. See [16, VI.3.4].)

We are searching for points in $E(\mathbb{Q})$, so in particular for points in $E(\mathbb{R})$. The real locus of \mathbb{C}/L consists of either one or two circles. More precisely, in all cases it contains the circle \mathbb{R}/\mathbb{Z} ; and if $\Delta > 0$ (equivalently $q > 0$, $\tau \in i\mathbb{R}$), then it also contains the circle $(\mathbb{R}/\mathbb{Z}) + \tau/2$. For this step, we will concentrate on the set \mathbb{R}/\mathbb{Z} , which corresponds to the identity component of $E(\mathbb{R})$. Notice if $z \in \mathbb{R}/\mathbb{Z}$, then $u = e^{2\pi iz}$ has absolute value 1. We compute

$$\begin{aligned} 1-u &= u^{1/2}(u^{-1/2} - u^{1/2}) \\ &= -2iu^{1/2} \sin(\pi z), \\ (1-q^n u)(1-q^n u^{-1}) &= (1-q^n)^2 - q^n(u^{1/2} - u^{-1/2})^2 \\ &= (1-q^n)^2 + 4q^n \sin^2(\pi z). \end{aligned}$$

So if we write $S = 2 \sin(\pi z)$, then for points in $z \in \mathbb{R}/\mathbb{Z}$, the archimedean local height is given by (remember $|u| = 1$)

$$\hat{\lambda}_{\infty}(z) = -\frac{1}{12} \log|q| - \log|S| - \sum_{n \geq 1} \log|(1-q^n)^2 + q^n S^2|.$$

In terms of the function $f(S)$ defined above, this can be rewritten as

$$f(S) = |q|^{-1/12} \exp(-\hat{\lambda}_{\infty}(z)).$$

We are looking for a point with $\hat{\lambda}_{\infty}(z) = H - \log d - \lambda$, so we want to solve the equation

$$f(S) = |q|^{-1/12} \exp(-H + \log d + \lambda)$$

for S . This is exactly the task specified in Step 11.

Step 12. In Step 11 we found a value of S so that $S = 2 \sin(\pi z)$, where $z \in \mathbb{R}/\mathbb{Z}$ corresponds to the desired point of $\mathbb{R}/\mathbb{Z} \cong E(\mathbb{R})$. In this step we use the known value of S to compute $z = (1/\pi) \sin^{-1}(S/2)$. Then we use an explicit isomorphism $\mathbb{C}/L \cong E(\mathbb{C})$ to compute the corresponding point on E . More precisely, the modified Weierstrass function in Step 8 gives the x -coordinate on E , so the number $\mathcal{P}(z)$ gives the x -coordinate of a candidate rational point. This rational point should have x -denominator d , so we compute the closest integer a to the real number $d^2 \mathcal{P}(z)$. [PARI: The point on E corresponding to z is equal to `pointell(e, e[15]*z)`.]

Step 13. Having found a candidate a/d^2 for the x -coordinate of a rational point on E , we plug this into the equation for E and check if it actually leads to a rational point. An alternative procedure is to use a modified Weierstrass derivative series \mathcal{P}' so that $(\mathcal{P}, \mathcal{P}') : \mathbb{C}/L \rightarrow E(\mathbb{C})$. Then one could let b be the integer closest to $d^3\mathcal{P}'(z)$ and check if $(a/d^2, b/d^3)$ is in $E(\mathbb{Q})$. This avoids checking if a large number is a perfect square, at the cost of evaluating another power series. Also, unfortunately, since b will tend to be much larger than a , this alternative method requires greater precision in the floating point calculations.

Step 14. If $\Delta > 0$ (equivalently, $\tau \in i\mathbb{R}$ or $q > 0$), then $E(\mathbb{R})$ has two components, and we need to check for a rational point of the known height on the non-identity component. (If there is a rational torsion point T on the non-identity component, then Steps 14–18 can be skipped. See Remark 4.1 in Section 4 for a discussion.)

Step 15. The justification for this step is much the same as in Step 11, but now we're looking for a point $z = \xi + \frac{1}{2}\tau$, where $\xi \in \mathbb{R}$ and $\tau \in i\mathbb{R}$. Let $\zeta = e^{2\pi i\xi}$ and $Q = q^{1/2} = e^{\pi i\tau}$. Notice that $u = e^{2\pi iz} = Q\zeta$. We compute

$$\begin{aligned} (1-u) \prod_{n \geq 1} (1-q^n u)(1-q^n u^{-1}) &= (1-Q\zeta) \prod_{n \geq 1} (1-Q^{2n+1}\zeta)(1-Q^{2n-1}\zeta^{-1}) \\ &= \prod_{n \geq 1} (1-Q^{2n-1}\zeta)(1-Q^{2n-1}\zeta^{-1}) \\ &= \prod_{n \geq 1} ((1-Q^{2n-1})^2 - Q^{2n-1}(\zeta^{1/2} - \zeta^{-1/2})^2) \\ &= \prod_{n \geq 1} ((1-Q^{2n-1})^2 + 4Q^{2n-1} \sin^2(\pi\xi)). \end{aligned}$$

Further, $B_2(\log |u|/\log |q|) = B_2(1/2) = 1/12$, so the formula for the archimedean local height is

$$\hat{\lambda}_\infty(z) = -\frac{1}{24} \log |q| + \log f(S),$$

where now

$$f(S) = \prod_{n \geq 1} ((1-Q^{2n-1})^2 + Q^{2n-1}S^2) \quad \text{and} \quad S = 2 \sin(\pi\xi).$$

So we need to solve

$$f(S) = q^{1/24} \exp(-\hat{\lambda}_\infty(P)) = Q^{1/12} \exp(-H + \log d + \lambda)$$

for S . The derivative

$$f'(S) = 2Sf(S) \sum_{n \geq 1} \frac{Q^{2n-1}}{(1-Q^{2n-1})^2 + Q^{2n-1}S^2}$$

is rapidly convergent, so again we can use Newton iteration to solve $f(S) = c$. Further, $f'(S) > 0$ for $S > 0$, so there is only one solution.

Step 16. Using the value of S from Step 15, we compute the point

$$z = (1/\pi) \sin^{-1}(S/2) + \tau/2 \in \mathbb{C}/L,$$

use the modified Weierstrass function to get a possible rational x -coordinate $\mathcal{P}(z)$, and multiply by the denominator and take the closest integer a to $d^2\mathcal{P}(z)$ to get

the numerator of the hypothetical rational point. [PARI: As in Step (12), compute `pointell(e, e[15]*z)`.]

Step 17. In this step we check if the rational number a/d^2 obtained in Step 16 gives a rational point on E . The remarks made in Step 13 apply also to this step. [PARI: `ordell(a/d^2)` will return the y -coordinates of any points P in $E(\mathbb{Q})$ with $x(P) = a/d^2$, and will return `[]` if there are no such points.]

Step 18. This is the end of the section dealing with the case that $E(\mathbb{R})$ has two components.

Step 19. This is the end of the loop over λ 's in Λ_{bad} .

Step 20. This is the end of the loop over possible denominators d .

Step 21. If this step is reached, then the sought after point has denominator greater than d_{max} .

This completes our detailed description of the steps in the Canonical Height Search Algorithm used as a straight search method. We next describe the modification used to determine if the (possibly empty) set $E(\mathbb{Z}_S)$ contains any non-torsion points.

Step a. See the justification for Steps (1)–(8) above.

Step b. We will search for points whose canonical height is H_0/m^2 until this height is smaller than 0.1.

Step c. Compute the canonical height $H = H_0/m^2$ of the hypothetical point. Thus H will be the height of a point Q satisfying $\hat{h}(mQ) = H_0$.

Step d. Compute a number d_{max} with the property that if $P = (a/d^2, b/d^3) \in E(\mathbb{Q})$ has height $\hat{h}(P) = H$, then necessarily $1 \leq d \leq d_{\text{max}}$. This estimate depends on an explicit bound for the difference between the canonical height and the naive height. See Proposition 4.2 below for details.

Step e. Loop over the allowable d 's. From Step (d), these d 's must satisfy $1 \leq d \leq d_{\text{max}}$, and since we are looking for S -integral points, they must be products of primes in S .

Step f. See the justification for Steps (10)–(19) above. If $E(\mathbb{Q})$ has a point P with $x(P) = a/d^2$ and $\hat{h}(P) = H_0/m^2$, then the algorithm will find this point and terminate.

Step g. This is the end of the loop over d , the allowable denominators.

Step h. This is the end of the loop over the multipliers m .

Step i. Compute a number d_{max} with the property that if $P = (a/d^2, b/d^3) \in E(\mathbb{Q})$ has height $\hat{h}(P) < 0.1$, then necessarily $1 \leq d \leq d_{\text{max}}$ and $|a| \leq d_{\text{max}}^2$. As in Step (d) above, we refer the reader to Proposition 4.2 for details. [PARI: `ordell(a/d^2)` will return the y -coordinates of any points P in $E(\mathbb{Q})$ with $x(P) = a/d^2$, and will return `[]` if there are no such points.]

Step j. Perform an exhaustive or sieve-assisted search for all points $P \in E(\mathbb{Q})$ with $x(P) = a/d^2$ satisfying $1 \leq d \leq d_{\text{max}}$ and $|a| \leq d_{\text{max}}^2$. (In principle, it's only necessary to consider d 's in $\mathbb{Z} \cap \mathbb{Z}_S^*$, but in practice, it doesn't take long to perform a complete search.)

Step k. If the algorithm has not discovered a point in $E(\mathbb{Q})$, then one concludes that $E(\mathbb{Z}_S)$ contains no non-torsion points.

3. A MODIFIED ALGORITHM IF $c_6 < 0$ AND $|j| \gg 1$

In this section we briefly discuss a modified algorithm which can be used when the multiplicative period $|q|$ is very close to 1. This will occur when c_6 is negative and the j -invariant of E is large. Note that the periods ω_1 and ω_2 computed in Step 7 give an isomorphism

$$\phi : E(\mathbb{C}) \longrightarrow \mathbb{C}/(\omega_1\mathbb{Z} + \omega_2\mathbb{Z})$$

defined over \mathbb{R} . Our algorithm depends quite strongly on the fact that ϕ is defined over \mathbb{R} , since we use the inclusions $E(\mathbb{Q}) \subset E(\mathbb{R}) \subset E(\mathbb{C})$.

As usual, let $\tau = \omega_2/\omega_1$ and $q = e^{2\pi i\tau}$. If $|q|$ is small, we consider the lattice $-\omega_2\mathbb{Z} + \omega_1\mathbb{Z}$ with corresponding $\tau' = -1/\tau = -\omega_1/\omega_2$ and $q' = e^{2\pi i\tau'}$. Then we have a commutative diagram of isomorphisms

$$\begin{array}{ccc} E(\mathbb{C}) & \xrightarrow{\psi} & E'(\mathbb{C}) \\ \phi \downarrow & & \phi' \downarrow \\ \mathbb{C}/(\omega_1\mathbb{Z} + \omega_2\mathbb{Z}) & \xrightarrow{z \rightarrow iz} & \mathbb{C}/(-\omega_2\mathbb{Z} + \omega_1\mathbb{Z}) \end{array}$$

In this diagram, E'/\mathbb{Q} is the \mathbb{C}/\mathbb{R} twist of E , which means in particular that $c_4(E') = c_4(E)$ and $c_6(E') = -c_6(E)$. The vertical maps are defined over \mathbb{R} , but the horizontal maps satisfy $\psi' = -\psi$.

The idea now is to perform Steps 11 and 15 working on the curve E' whose q' has small magnitude. We will explain how to do Step 11 when $\Delta > 0$ (equivalently $\omega_2 \in i\mathbb{R}$), and leave the other similar cases for the reader. The correspondence between E and E' in this case is given by the formulas

$$\tau' = -1/\tau, \quad q' = e^{2\pi i\tau'} = e^{-2\pi i/\tau}, \quad z' = iz, \quad u' = e^{2\pi iz'/(-i\omega_2)} = e^{2\pi iz/\omega_2}.$$

Note that $u' \in \mathbb{R}$ and $u' > 0$, since $\omega_2 \in i\mathbb{R}$. Similarly, $q' > 0$. We also observe that

$$\hat{\lambda}_\infty(z; E) = \hat{\lambda}_\infty(z'; E'),$$

since normalized local heights are model independent. So, given a value H for the canonical height of a rational point, a choice of $\lambda \in \Lambda_{\text{bad}}$, and a choice of denominator d , we solve the equation

$$|q'|^{(1/2)B_2(\log u' / \log q')} |1 - u'| \cdot \left| \prod_{n \geq 1} (1 - q'^n u') (1 - q'^n u'^{-1}) \right| = \exp(-H + \log d + \lambda)$$

numerically for $u' \in \mathbb{R}$. By periodicity, we can restrict attention to u' satisfying $(q')^{1/2} \leq u' \leq (q')^{-1/2}$. Having solved for u' , we can recover z by the formula $z = (\omega_2/2\pi i) \log u'$, and then continue with Step 12 of the algorithm.

4. ERROR ESTIMATES AND FURTHER REMARKS

In this section we will give two estimates which are needed to use the Canonical Height Search Algorithm in practice, and then we will make some remarks.

First, Step (3) of the algorithm says to compute $L'(E, 1)$ using an explicit series, so we need to estimate the error in taking only a finite number of terms in this series. We will use the trivial estimate for the exponential integral, namely

$$E_1(x) = \int_x^\infty \frac{dt}{te^t} \leq \frac{1}{x} \int_x^\infty \frac{dt}{e^t} = \frac{1}{xe^x}.$$

Further, we have the Hasse-Weil estimate

$$|a_n| \leq d(n)\sqrt{n}$$

for the n^{th} coefficient of $L(E, s)$, where $d(n)$ is the number of divisors of n . We will use the estimate

$$(2) \quad d(n) \leq n^{1/\log \log n},$$

which holds for virtually all $n \geq 100$. (To be more accurate, we should probably replace the exponent with $1.06602/\log \log n$, which will deal with $n = 6983776800 = 2^5 \cdot 3^3 \cdot 5^2 \cdot 7 \cdot 11 \cdot 13 \cdot 17 \cdot 19$. However, many $a(n)$'s will be a lot smaller than $d(n)\sqrt{n}$, so in practice we use (2).) Hence if we compute $L'(E, 1)$ using m terms of the series, we can estimate the error by

$$\begin{aligned} 2 \sum_{n \geq m} \frac{|a_n|}{n} E_1\left(\frac{2\pi n}{\sqrt{N}}\right) &\leq 2 \sum_{n \geq m} \frac{d(n)}{\sqrt{n}} \cdot \frac{\sqrt{N}}{2\pi n} e^{-2\pi n/\sqrt{N}} \\ &\leq \frac{\sqrt{N}}{\pi} \cdot \frac{1}{m^{3/2-1/\log \log m}} \cdot \frac{1}{e^{2\pi m/\sqrt{N}} - 1}. \end{aligned}$$

Thus in order to compute $L'(E, 1)$ accurate to within 10^{-k} , one needs to take roughly $ke\sqrt{N}$ terms. This may be compared with the Heegner point method, which requires computing $O(kN)$ terms of a slightly simpler series. We record our result as a proposition.

Proposition 4.1. *Let E/\mathbb{Q} be a modular elliptic curve of conductor N whose functional equation has odd sign. Then for any $m \geq 100$,*

$$\left| L'(E, 1) - 2 \sum_{n=1}^m \frac{a_n}{n} E_1\left(\frac{2\pi n}{\sqrt{N}}\right) \right| \leq \frac{\sqrt{N}}{\pi m^{3/2-1/\log \log m} (e^{2\pi m/\sqrt{N}} - 1)}.$$

Next we consider the use of our height search algorithm to prove that a curve has no integral (or S -integral) points. In this case we do not continue until a point appears, so we need an a priori bound for the height of the x -coordinate of the hypothetical point. Since we know the equation of E and the canonical height $\hat{h}(P)$, we can use an estimate for the difference between $\hat{h}(P)$ and $h(x(P))$. The original estimates for this difference are due to Dem'janenko and Zimmer; we will use the author's estimate [17], and will only give the upper bound required for the height search algorithm.

Proposition 4.2 ([17, Theorem 1.1]). *Let E/\mathbb{Q} be an elliptic curve given by a Weierstrass equation*

$$E : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6,$$

with integral coefficients, let $P \in E(\mathbb{Q})$, and write $x(P) = a/d^2$. Then

$$\begin{aligned} \log \max\{|a|, d^2\} &= h(x(P)) \\ &\leq 2\hat{h}(P) + \frac{1}{12}h(j) + \frac{1}{6} \log |\Delta| + \frac{1}{6} \log^+ |j| + \log^+ |b_2/12| + 2.64, \end{aligned}$$

where j is the j -invariant of E , Δ is the discriminant of the given Weierstrass equation, and $b_2 = a_1^2 + 4 * a_2$. (If $b_2 = 0$, then $\log 2$ may be subtracted from the upper bound.)

Remark 4.1. Suppose that $\Delta > 0$, so $E(\mathbb{R})$ is disconnected. If the bounded component of $E(\mathbb{R})$ has a rational torsion point T , then there is no need to search for rational points on the bounded component. This is true because if $P \in E(\mathbb{Q})$ is on the bounded component, then $P + T$ is on the unbounded component and $\hat{h}(P + T) = \hat{h}(P)$, so the algorithm will find $P + T$. Of course, it's possible that P has a smaller denominator than $P + T$, in which case the algorithm would find P first.

Remark 4.2. One may wish to compute $w(E)$, the sign of the functional equation of E , without doing an L -series computation. The value of $w(E)$ is given by a product of local signs,

$$w(E) = w_\infty(E) \prod_p w_p(E),$$

and in many cases there are simple formulas for $w_p(E)$. For example,

$$\begin{aligned} w_\infty(E) &= -1, \\ w_p(E) &= 1 \quad \text{if } E \text{ has good reduction,} \\ w_p(E) &= -1 \quad \text{if } E \text{ has split multiplicative reduction,} \\ w_p(E) &= 1 \quad \text{if } E \text{ has non-split multiplicative reduction.} \end{aligned}$$

(See [2, §6].) Further, if E has multiplicative reduction, or equivalently if

$$\text{ord}_p(\Delta) = -\text{ord}_p(j) > 0,$$

then one can distinguish between split versus non-split reduction by computing the quadratic residue symbol [16, V §5]

$$w_p(E) = -\left(\frac{-c_4c_6}{p}\right).$$

Finally, if E has additive reduction at $p \geq 5$, we mention that Rohrlich [12, Proposition 2] has given a simple algorithm to compute $w_p(E)$.

Remark 4.3. In certain cases the Canonical Height Search Algorithm can be used for curves of rank 2. More precisely, suppose that $E(\mathbb{Q})$ has rank 2. Using the value of $L''(E/\mathbb{Q}, 1)$ and the conjecture of Birch and Swinnerton-Dyer, we compute a value for $R(E/\mathbb{Q})\#\text{III}(E/\mathbb{Q})$, where $R(E/\mathbb{Q})$ is the height regulator [15, VIII §9]. Suppose for the sake of discussion that $\text{III}(E/\mathbb{Q}) = 1$. This means that we can find the value of

$$R(E/\mathbb{Q}) = \det \begin{vmatrix} \langle P_1, P_1 \rangle & \langle P_1, P_2 \rangle \\ \langle P_1, P_2 \rangle & \langle P_2, P_2 \rangle \end{vmatrix},$$

where $2\langle P, Q \rangle = \hat{h}(P + Q) - \hat{h}(P) - \hat{h}(Q)$. Suppose further that by a (sieve assisted) brute-force search we have found one point $P_1 \in E(\mathbb{Q})$. This is not an unreasonable situation, since there are elliptic curves of rank 2 with one small generator and one large generator. To cite just one example shown to me by John Cremona, the curve

$$y^2 = x^3 - 673$$

has rank 2 with smallest generators

$$P_1 = (29, -154) \quad \text{and} \quad P_2 = \left(\frac{33989323537}{3814421121}, \frac{1384230292401340}{235582462854081} \right).$$

Here is how to adapt our algorithm to help in finding P_2 . We loop over denominators d and d' , where d is a possible x -denominator for P_2 and d' is a possible denominator for $P_1 + P_2$ or $P_1 - P_2$. (Note that one of $P_1 \pm P_2$ will have canonical height less than $\hat{h}(P_1) + \hat{h}(P_2)$.) We also loop over $\lambda, \lambda' \in \Lambda_{\text{bad}}$, where λ is a possible value for $\hat{\lambda}_{\text{bad}}(P_2)$ and λ' is a possible value for $\hat{\lambda}_{\text{bad}}(P_1 + P_2)$. Since we already know the point P_1 , we can compute the corresponding value $z_1 \in \mathbb{C}/L$. This leads to the following equation for the regulator, where $z \in \mathbb{C}/L$ is the still unknown value corresponding to P_2 :

$$R = \begin{vmatrix} \hat{h}(P_1) & \left(\begin{array}{l} \hat{\lambda}_{\infty}(z_1 + z) + \log d' + \lambda' \\ -\hat{h}(P_1) - \hat{\lambda}_{\infty}(z) - \log d - \lambda \end{array} \right) \\ \left(\begin{array}{l} \hat{\lambda}_{\infty}(z_1 + z) + \log d' + \lambda' \\ -\hat{h}(P_1) - \hat{\lambda}_{\infty}(z) - \log d - \lambda \end{array} \right) & \hat{\lambda}_{\infty}(z) - \log d - \lambda \end{vmatrix}$$

(This is the formula to test whether d' is an x -denominator for $P_1 + P_2$. There is a similar formula to test whether d' is an x -denominator for $P_1 - P_2$. Both formulas should be checked, since one of $P_1 \pm P_2$ may have a much smaller denominator than the other.) Since we are assuming that the value of R is known from the L -series computation, and since the desired value of z actually lies in \mathbb{R}/\mathbb{Z} or $\mathbb{R}/\mathbb{Z} + \frac{1}{2}\tau$, we can use numerical methods to solve the equation for z . This gives a hypothetical real approximation $\mathcal{P}(z)$ for $x(P_2)$, and we conclude by checking if $d^2\mathcal{P}(z)$ is close to an integer a satisfying $a/d^2 \in x(E(\mathbb{Q}))$.

How efficient is the method just described? Suppose that $x(P_2)$ actually equals a_2/d_2^2 , where $|a_2|$ and d_2^2 are each approximately equal to the multiplicative height $D = H(P_2)$. The “brute-force” method requires a loop $a \leq |a_2|$ and $d \leq d_2$, so has running time $O(D^{3/2})$, and the algorithm for rank 1 described above has running time $O(D^{1/2})$. In the rank 2 case, we need to loop over possible denominators $d, d' \ll d_2$, so the running time will be $O(D)$. This is still better than the brute-force method, but in practical terms it may not actually be faster than a sieve-assisted search.

Remark 4.4. The method of homogeneous spaces provides a powerful tool for searching for rational points on elliptic curves. Assuming that E has no rational torsion, it is generally only feasible to use homogeneous spaces of degree 4, as described in [6, §3.6]. The search for the homogeneous space(s) has expected runtime $O(\Delta_E^{1/2})$. Assuming that a (locally trivial) homogeneous space C is found, one then does an exhaustive (possibly sieve assisted) search for points in $C(\mathbb{Q})$. The expected running time is $O(D^{1/2})$, so about the same as the Canonical Height Search Algorithm. In practice, it is not clear a priori which will be faster for any given elliptic curve.

However, in principle it is possible to combine the two methods. Thus one uses the known value of $\hat{h}(P)$ to eliminate one of the loops in the exhaustive search on C , yielding a running time of $O(D^{1/4})$. Unfortunately, the method does not seem to be practical at this point, because of the need to check a large number of possible common divisors of certain numbers. For this reason we have not included the lengthy details of the combined height-homogeneous space method.

5. NUMERICAL EXAMPLES

In this section we give four numerical examples. The first two are from Cremona’s list [6] of curves of conductor up to 10^4 . At the time this paper was originally written, these two curves were expected to have rank 1, but a sieve-assisted search up to $h(x) \leq 7.5$ had failed to find any points. Since the writing of this paper, Cremona has improved both the sieve-assisted search and the homogeneous space method so that they are able to deal with these two examples. Further, Zagier has explained that since the conductors are comparatively small, the Heegner Point method will also find rational points quite rapidly. However, we will describe these two examples in some detail so as to illustrate how the Canonical Height Search Algorithm is used as a straight search algorithm.

Our third example is a curve E of large conductor $N \approx 10^7$ which has a generator of large height $H \approx 16$. This example is certainly at (if not beyond) the limits of feasible computability using any of the other search algorithms, but we will be able to verify quite rapidly that $E(\mathbb{Z}) = \emptyset$, and after a fairly lengthy search (36 hours), we will be able to find a generator for $E(\mathbb{Q})$. Our final example is a curve E of conductor greater than 10^8 . In this case we will not hunt for a generator for $E(\mathbb{Q})$, but will be content to prove that $E(\mathbb{Z}[1/2]) = \emptyset$.

Example 1. For this first example we will give a detailed description of each step of the Canonical Height Search Algorithm. Hopefully this will be useful for testing should the reader decide to implement the algorithm. To save space, we will generally only write the first few significant digits of real numbers. However, the actual calculations for this example were all performed with 28 digits of accuracy. For larger examples, it might be worthwhile to work with 50 or 100 digits.

We will consider the elliptic curve

$$E : y^2 + xy = x^3 - x^2 - 12396x - 1140144.$$

First we compute the associated quantities:

$$\begin{aligned} b_2 = -3, \quad b_4 = -24792, \quad b_6 = -4560576, \quad b_8 = -150240384, \\ c_4 = 595017, \quad c_6 = 987761979, \quad \Delta = -442714581230976, \\ j = -7802330770032219/16396836341888. \end{aligned}$$

The discriminant factors as $\Delta = -2^7 \cdot 3^3 \cdot 71^6$. (This shows, in particular, that the given Weierstrass equation is minimal.) Using Tate’s algorithm and Ogg’s formula, we compute the reduction types, local Tamagawa numbers, and the exponent of the conductor for each prime of bad reduction:

$p = 2$	Type I_7	$c_2 = 1$	$f_2 = 1$
$p = 3$	Type II	$c_3 = 1$	$f_3 = 3$
$p = 71$	Type I_6	$c_{71} = 6$	$f_{71} = 1$.

So the conductor of E is

$$N = 2^{f_2} \cdot 3^{f_3} \cdot 71^{f_{71}} = 3834,$$

and the (global) Tamagawa number is $c = c_2 c_3 c_{71} = 6$. Further, a search for torsion points yields

$$E(\mathbb{Q})_{\text{tors}} = \{O, (217, -2629), (217, 2412)\},$$

so $T = \#E(\mathbb{Q})_{\text{tors}} = 3$.

At this point we will also compute the periods of E . The lattice of E is generated by the two periods

$$\omega_1 \approx 0.21217 \quad \text{and} \quad \omega_2 \approx 0.10608 + 0.23494i.$$

Associated to these are the quantities

$$\tau = \omega_2/\omega_1 \approx 0.5 + 1.10736i \quad \text{and} \quad q = e^{2\pi i\tau} \approx -0.00095121.$$

The fact that $|q|$ is small means that the height series will converge rapidly, which is good. Further, we note that since $\Delta < 0$, the real locus has only one component and the real period Ω is given by

$$\Omega = \int_{E(\mathbb{R})} dx/(2y + a_1x + a_3) = \omega_1 \approx 0.21217.$$

Next we consider the computation of the L -series via the sum

$$L'(E, 1) \approx 2 \sum_{n=1}^m \frac{a_n}{n} E_1\left(\frac{2\pi n}{\sqrt{N}}\right).$$

Proposition 4.1 tells us that if we take $m = 650$, the finite sum will differ from $L'(E, 1)$ by less than 10^{-30} . In order to compute the first 650 coefficients of the L -series of E , we start by computing the a_p 's with p prime via the formula

$$a_p = p + 1 - \#E(\mathbb{F}_p).$$

For such small primes, one can compute the a_p 's by the brute force formula

$$a_p = \sum_{0 \leq x < p} \left(\frac{4x^3 + b_2x^2 + 2b_4x + b_6}{p} \right),$$

although in general for larger primes one should use a more efficient method for computing $\#E(\mathbb{F}_p)$, such as Shanks' baby step-giant step method. In any case, the a_p 's for our curve are

$$\begin{aligned} a_1 = 1, \quad a_2 = -1, \quad a_3 = 0, \quad a_5 = -3, \quad a_7 = -1, \quad a_{11} = 3, \\ a_{13} = 2, \quad a_{17} = 0, \quad \dots \quad a_{613} = 44, \quad a_{617} = 18, \quad a_{619} = -10, \\ a_{631} = -37, \quad a_{641} = -6, \quad a_{643} = 14, \quad a_{647} = 36. \end{aligned}$$

Using these and the standard recursions

$$a_{mn} = a_m a_n \quad \text{for } \gcd(m, n) = 1, \quad a_{p^k} = a_p a_{p^{k-1}} - p a_{p^{k-2}} \quad \text{for } k \geq 2,$$

we compute all a_n 's for $n \leq 650$.

Now that we know the a_n 's, we can compute the value of the L -series

$$\begin{aligned} L'(E, 1) &\approx \sum_{n=1}^{650} \frac{a_n}{n} E_1\left(\frac{2\pi n}{\sqrt{3834}}\right) \\ &\approx 2.164605251532673588502104879 \end{aligned}$$

Using this, we obtain the canonical height of the sought-for rational point on E :

$$H = \frac{L'(E, 1)T^2}{2\Omega c} \approx 7.651655428781607647931654194.$$

So far we have done nothing new; such computations have been performed by many people. Now we begin the heart of our algorithm. We need to compile a list of possible contributions to the height H coming from the three primes of bad

reduction. The curve E has Type I_7 reduction at $p = 2$, so looking at Table 1 we see that the possible contributions are

$$\Lambda_{\text{bad}}(2) = \left\{ \left(\frac{7}{12} - \frac{i}{2} + \frac{i^2}{14} \right) \log 2 : 0 \leq i \leq \frac{7}{2} \right\} \\ \approx \{0.40434, 0.10727, -0.09077, -0.18979\}.$$

The reduction at $p = 3$ is Type II , so $\Lambda_{\text{bad}}(3) = \{(3/12) \log 3\} \approx \{0.27465\}$ has only one element. Finally, E has Type I_6 reduction at $p = 71$, so

$$\Lambda_{\text{bad}}(71) = \left\{ \left(\frac{6}{12} - \frac{i}{2} + \frac{i^2}{12} \right) \log 71 : 0 \leq i \leq \frac{6}{2} \right\} \\ \approx \{2.14523, 0.35754, -0.71508, -1.07261\}.$$

Now we take all possible sums using one element from $\Lambda_{\text{bad}}(2)$, one element from $\Lambda_{\text{bad}}(3)$, and one element from $\Lambda_{\text{bad}}(71)$, to form our list Λ_{bad} of possible bad prime contributions to the height. For our curve E , the set Λ_{bad} will have 16 elements,

$$\Lambda_{\text{bad}} \approx \{2.81033, 1.03421, -0.03146, -0.38668, 2.51327, \\ 0.73715, -0.32852, -0.68374, 2.31522, 0.53911, -0.52656, \\ -0.88179, 2.21620, 0.44009, -0.62558, -0.98081\}.$$

We are now ready for the main loops. We start with $d = 1$ and take the first element $\lambda \approx 2.81033 \in \Lambda_{\text{bad}}$. We need to solve

$$S \prod_{n \geq 1} ((1 - q^n)^2 + q^n S^2) = |q|^{-1/12} \exp(\lambda + \log d - H) \approx 0.014101.$$

Using Newton’s method, we find that $S \approx 0.0140742$ and $z = (1/\pi) \sin^{-1}(S/2) \approx 0.00224$. Now the modified Weierstrass \wp function gives $\mathcal{P}(z) \approx 4427271.46956$. We next compute a potential numerator for $x(P)$,

$$a = \lfloor d^2 \mathcal{P}(z) \rfloor = 4427271.$$

The fact that $d^2 \mathcal{P}(z)$ is not close to an integer virtually ensures that this is not the correct a/d^2 . In any case, we can check to see if the quantity

$$4a^3 + b_2 a^2 d^2 + 2b_4 a d^4 + b_6 d^6 = 347110888577755405881$$

is a perfect square. Alas, it is not, so we continue on to the next element of Λ_{bad} and repeat the process. After 16 iterations, we have exhausted all of the elements of Λ_{bad} and we still have no point in $E(\mathbb{Q})$. This shows that the point we are searching for is not an integer point (i.e., does not have $d = 1$), so we increment to $d = 2$ and begin again.

Eventually the value of d will equal $d = 354$, and looping through the elements of Λ_{bad} we come to the value $\lambda \approx 1.03421$. Solving as above, we find that

$$S \approx 0.844039, \quad z \approx 0.138678, \quad \mathcal{P}(z) \approx 1157.598159851894410929171077.$$

Then $d^2 \mathcal{P}(z)$ is an integer to within 20 decimal places, definitely a good sign, and we set

$$a = \lfloor d^2 \mathcal{P}(z) \rfloor = 145065571.$$

Now

$$\begin{aligned} 4a^3 + b_2a^2d^2 + 2b_4ad^4 + b_6d^6 &= 12081205978054421909840896 \\ &= 3475802925664^2 \end{aligned}$$

is a perfect square, so we have found the rational point

$$\left(\frac{145065571}{354^2}, \frac{1712224856765}{354^3} \right) \in E(\mathbb{Q}).$$

This example was computed on a Power Macintosh 7100/80 using PARI 1.39 (with no great effort to streamline the computations). It took approximately 1 second to check each potential denominator d , and thus approximately 6 minutes to find the point. As mentioned earlier, Zagier has observed that the Heegner Point Algorithm provides an even quicker way to solve this example.

Example 2. We only briefly give the details for our second example, the curve

$$E : y^2 + xy = x^3 - x^2 - 34911x - 2501928$$

of conductor $N = 3879$. Using 1000 terms of the L -series gives

$$L'(E, 1) \approx 5.034164639731370882701288647,$$

and using this and the fact that $\#E(\mathbb{Q})_{\text{tors}} = 4$, we compute

$$H \approx 14.41346735625563972827987389.$$

Further,

$$q \approx 0.00000002097967707074283079352765088 \quad \text{and} \quad \#\Lambda_{\text{bad}} = 6,$$

so each denominator d can be checked fairly rapidly. After a number of hours, we find that $d = 125714$ yields the point

$$\left(\frac{8218827853779}{125714^2}, -\frac{22261338488996940783}{125714^3} \right) \in E(\mathbb{Q}).$$

Zagier [21] has used PARI to recompute Example 2 via Heegner points in 10 to 15 seconds. So for our next two examples we choose curves having conductors which are at or beyond the range at which Heegner point computations are practical.

Example 3. For our third example, we will consider the curve

$$E : y^2 + xy + y = x^3 - x^2 - 21x - 152,$$

which has conductor $N = 10069019$. We will first prove that $E(\mathbb{Z})$ is empty, a calculation which takes only a few minutes. We will then proceed to do a complete search, which takes considerably longer, but does eventually produce a generator.

The curve E has discriminant and j -invariant

$$\Delta = -10069019 \quad \text{and} \quad j = -\frac{979146657}{10069019} \approx -97.2435.$$

We begin by computing the L -series using 46000 terms of the series and 50 digits of precision:

$$L'(E, 1) = 31.047460928677357675133339778016130653212869682533.$$

(This took PARI a little less than 6 minutes on an 80MHz Power Macintosh.) According to Proposition 4.1, this value is accurate to at least 42 decimal places. Next we compute the torsion subgroup, Tamagawa number, and real period,

$$T = 1, \quad c = 1, \quad \Omega = 0.97249\dots,$$

and using these we find the canonical height

$$H = 15.962820584929157630512545697200916399133798839124.$$

Finally, we compute the multiplicative period

$$q = -0.0018219151314974123101127422317525825667588087666281$$

and observe that q is sufficiently small to permit the algorithm to operate quite efficiently. Further, there is only one bad prime, namely N , and since the discriminant $\Delta = -N$ is also prime, the set Λ_{bad} consists of the single value

$$\Lambda_{\text{bad}} = \left\{ \frac{1}{12} \log N \right\} = \{1.343747820156317559\dots\}.$$

We are now ready to illustrate how the Canonical Height Search Algorithm can be used to verify that $E(\mathbb{Z}) = \emptyset$. The first step is to apply the algorithm with denominator $d = 1$ and height H . This yields no points, which shows that $E(\mathbb{Z})$ contains no points with canonical height H . Continuing, we apply the algorithm with denominator $d = 1$ and heights $H/4, H/9, \dots, H/13^2$, and still find no (integral) points. (Computation time is very small, on the order of 4 seconds to check all 13 height values.) This proves that if $E(\mathbb{Z})$ is non-empty, then it must contain a point Q of canonical height less than $H/14^2$. Proposition 4.2 then tells us that such a point Q will satisfy

$$\log |x(Q)| \leq 9.521, \quad \text{so} \quad |x(Q)| \leq 13640.$$

This bound is sufficiently small that it is a simple matter to perform an exhaustive search and verify (in approximately 2.8 seconds) that there are no such points. So we have proven that $E(\mathbb{Z})$ is empty.

The most time-consuming part of the computation was the evaluation of $L'(E, 1)$. Having computed $L'(E, 1)$ (and thus H) once, we see that for any particular d , it takes us less than 7 seconds to show that $E(\mathbb{Q})$ has no points of the form $(a/d^2, b/d^3)$. Using Proposition 4.2, we know that there is a point in $P \in E(\mathbb{Q})$ with canonical height H whose denominator satisfies

$$(3) \quad d < e^{20.6418} \approx 9.22 \cdot 10^8.$$

Clearly we cannot check all d 's up to this bound, but there is nothing to stop us from trying small d 's, say $d < 10^6$, and hoping that the denominator of $x(P)$ is smaller than 10^{12} . Indeed, although it takes around 36 hours, the Canonical Height Search Algorithm eventually discovers the point

$$P = \left(\frac{72574710196444}{463995^2}, \frac{600637381549819353188}{463995^3} \right) \in E(\mathbb{Q}).$$

Computing $\hat{h}(P)$ directly, we find that our hypothetical height H differs from $\hat{h}(P)$ by less than 10^{-43} , just as expected. We also find that the Canonical Height Search Algorithm is able to identify P using only 23 digits of H , so the computation would have been faster if we'd worked with 28 digits of accuracy.

We should also check that P is actually a generator of $E(\mathbb{Q})$. To do this, we use the analytic parametrization $E(\mathbb{C}) \cong \mathbb{C}/L$ to compute analytically the points $(1/m)P$ for $m = 2, 3, 5, \dots, 23$ and to check if any of them have rational coordinates.

This takes under a minute using the PARI functions `zell` and `pointell`. (Note we only need to check prime values of m .) Finally, we suppose that $P = mQ$ for some $m \geq 29$. In this case $\hat{h}(Q) \leq \hat{h}(P)/29^2$, so Proposition 4.2 tells us that $x(Q) = a/d^2$ with

$$|a| \leq 12038 \quad \text{and} \quad 1 \leq d \leq 109.$$

It takes under 5 minutes to check by brute force that such a Q does not exist. This completes the verification that $E(\mathbb{Q}) = \mathbb{Z}P$.

Total computation time to prove that $E(\mathbb{Z}) = \emptyset$ was 6 minutes, and the time to find a generator for $E(\mathbb{Q})$ was approximately 36 hours. (The latter figure could certainly be considerably reduced if more care were taken in programming the algorithm.)

For comparison purposes, we expect that the Heegner Point Algorithm would take far longer, although we do not say that it is impossible to use the Heegner Point Algorithm for this problem. The difficulty in using Heegner points here is that in order to compute the Heegner point analytically, one must sum a series $\sum a_n q_1^n/n$ with $|q_1| \approx e^{-2\pi/N}$, where the a_n 's are the usual coefficients of $L(E/\mathbb{Q}, s)$. So to get (say) 28 digits of accuracy requires (at least) $5 \cdot 10^7$ terms. However, PARI takes around 0.22 seconds to compute each a_n when n gets around 10^7 , so even a rough estimate suggests that it would take hundreds of hours to find P via Heegner points. Finally, we mention that sieve-assisted and homogeneous space searches (even for integral points) are unlikely to be feasible on a curve with such a large conductor and generating point.

Example 4. For our final example we will look at the curve

$$y^2 + xy + y = x^3 - x^2 + 174x + 256$$

of conductor and discriminant

$$N = 377812871 \quad \text{and} \quad \Delta = -377812871.$$

It took 64 minutes to compute 250000 terms of the series for $L'(E, 1)$, yielding

$$L'(E, 1) = 30.03146162954767138480670957295473260515,$$

and Proposition 4.1 says that this is accurate to at least 37 decimal places. From this, it is an easy matter to compute

$$H = 14.21034299783875777356455736075754044495.$$

We can now use the Canonical Height Search Algorithm to search for points. Applying it with $d = 1$ and height H gives no point, so we conclude that $E(\mathbb{Z})$ contains no points of height H . Similarly applying the algorithm with $d = 1$ and H/m^2 for $m = 2, 3, \dots, 25$, we find (in 15 seconds) that there are no points in $E(\mathbb{Z})$ with height greater than $H/625$. Next, Proposition 4.2 tells us that any integer point $Q \in E(\mathbb{Z})$ with $\hat{h}(Q) < H/625$ will satisfy

$$|x(Q)| \leq 54926.$$

It takes 11 seconds to verify by a brute-force search that there are no such points, which completes the proof that $E(\mathbb{Z}) = \emptyset$.

We will conclude by explaining how the Canonical Height Search Algorithm can be used to prove that $E(\mathbb{Z}[1/2]) = \emptyset$. Suppose first that there is a point

$P = (a/d^2, b/d^3) \in E(\mathbb{Z}[1/2])$ with $\hat{h}(P) = H$. Proposition 4.2 tells us that

$$1 \leq d \leq 340008469,$$

so we see that $d = 2^e$ for some $0 \leq e \leq 28$. Applying the Canonical Height Search Algorithm with height H and denominator $d = 1, 2, 4, 8, \dots, 2^{28}$ takes 18 seconds and yields no points. Next we look for $P \in E(\mathbb{Z}[1/2])$ satisfying $\hat{h}(P) = H/4$. Such a P has denominator $d \leq 7996$, so we only need check $d = 2^e$ for $1 \leq e \leq 12$. Again we get no points (in 8 seconds). Continuing, we take $H/9$ with $1 \leq e \leq 10$ (7 seconds), $H/16$ with $1 \leq e \leq 9$ (6 seconds), \dots , $H/225$ with $1 \leq e \leq 7$ (5 seconds). Total time expended is a scant 100 seconds, and we now know that $E(\mathbb{Z}[1/2])$ contains no points with height greater than $H/225$. On the other hand, if $\hat{h}(P) \leq H/225$, then Proposition 4.2 says that

$$|a| \leq 59551 \quad \text{and} \quad d = 2^e \quad \text{with} \quad 0 \leq e \leq 7.$$

It takes about 7.5 minutes to check by brute force that $E(\mathbb{Q})$ contains no such points, which completes the verification that $E(\mathbb{Z}[1/2]) = \emptyset$.

ADDED IN PROOF

John Cremona has informed me that his mwrank program is now able to solve Example 3 in just a few seconds using homogeneous spaces.

REFERENCES

1. C. Batut, D. Bernardi, H. Cohen, M. Olivier, *PARI-GP*, a computer system for number theory, Version 1.39.
2. B.J. Birch, H.P.F. Swinnerton-Dyer, *Elliptic curves and modular functions*, Modular Functions of One Variable IV (B.J. Birch, W. Kuyk, eds.), Lecture Notes in Math. 476, Springer-Verlag, Berlin, 1975. MR **52**:5685
3. A. Bremner, *On the equation $Y^2 = X(X^2 + p)$* , Number Theory and Applications (R.A. Mollin, ed.), Kluwer Academic Publishers, 1989, pp. 3–22. MR **92h**:11047
4. A. Bremner, J.W.S. Cassels, *On the equation $Y^2 = X(X^2 + p)$* , Math. Comp. **42** (1984), 257–264. MR **85f**:11017
5. H. Cohen, *A Course in Computational Algebraic Number Theory*, Graduate Texts in Math., vol. 138, Springer Verlag, Berlin, 1993. MR **94i**:11105
6. J. Cremona, *Algorithms for Modular Elliptic Curves*, Cambridge University Press, Cambridge, 1992. MR **93m**:11053
7. ———, private communication, November 1995.
8. F. Diamond, *On deformation rings and Hecke rings*, Annals of Math. **144** (1996), 137–166. MR **97d**:11172
9. N. Elkies, *Heegner point computations*, Algorithmic Number Theory (L.M. Adelman, M.-D. Huang, eds.), ANTS-I, Lecture Notes in Computer Science, vol. 877, 1994, pp. 122–133. MR **96f**:11080
10. B. Gross and D. Zagier, *Heegner points and derivatives of L -series*, Invent. Math. **84** (1986), 225–320. MR **87j**:11057
11. V.A. Kolyvagin, *Euler systems*, The Grothendieck Festschrift, Vol. II, Birkhäuser, Boston, 1990, pp. 435–483. MR **92g**:11109
12. D. Rohrlich, *Variation of the root number in families of elliptic curves*, Compositio Math. **87** (1993), 119–151. MR **94d**:11045
13. K. Rubin, *p -adic L -functions and rational points on elliptic curves with complex multiplication*, Invent. Math. **107** (1992), 323–350. MR **92m**:11063
14. J.H. Silverman, *The Néron-Tate Height on Elliptic Curves*, Ph.D. thesis, Harvard, 1981.
15. ———, *The Arithmetic of Elliptic Curves*, Graduate Texts in Math., vol. 106, Springer-Verlag, Berlin and New York, 1986. MR **87g**:11070
16. ———, *Advanced Topics in the Arithmetic of Elliptic Curves*, Graduate Texts in Math., vol. 151, Springer-Verlag, Berlin and New York, 1994. MR **96b**:11074

17. ———, *The difference between the Weil height and the canonical height on elliptic curves*, Math. Comp. **55** (1990), 723–743. MR **91d**:11063
18. R. Taylor and A. Wiles, *Ring-theoretic properties of certain Hecke algebras*, Annals of Math. **141** (1995), 553–572. MR **96d**:11072
19. D. Ulmer, *A construction of local points on elliptic curves over modular curves*, International Math. Research Notes **7** (1995), 349–363. MR **97b**:11076
20. A. Wiles, *Modular elliptic curves and Fermat's last theorem*, Annals of Math. **141** (1995), 443–551. MR **96d**:11071
21. D. Zagier, private communication.

MATHEMATICS DEPARTMENT, BOX 1917, BROWN UNIVERSITY, PROVIDENCE, RI 02912 USA
E-mail address: `jhs@gauss.math.brown.edu`