AN ALMOST LINEAR TIME ALGORITHM TESTING WHETHER THE MARKOFF GRAPH MODULO p IS CONNECTED

COLBY AUSTIN BROWN

Abstract. The Markoff graphs modulo p were proven by Chen (2024) to be connected for all but finitely many primes, and Baragar (1991) conjectured that they are connected for all primes, equivalently that every solution to the Markoff equation modulo p lifts to a solution over $\mathbb{Z}$. In this paper, we provide an algorithmic realization of the process introduced by Bourgain, Gamburd, and Sarnak [arXiv:1607.01530] to test whether the Markoff graph modulo p is connected for arbitrary primes. Our algorithm runs in $o(p^{1+\epsilon})$ time for every $\epsilon > 0$. We demonstrate this algorithm by confirming that the Markoff graph modulo p is connected for all primes less than one million.

1. Introduction

The Markoff equation is

$$x^2 + y^2 + z^2 - xyz = 0. \tag{1}$$

It is often written as

$$x^2 + y^2 + z^2 - 3xyz = 0,$$

but the solutions to the latter are in bijection with the solutions to the former, reduced by a factor of 3. In this paper, "the Markoff equation" will always refer to (1).

The set of integer Markoff triples $\mathcal{M}(\mathbb{Z})$ is closed under taking the negation of exactly two coordinates; in this paper, we will only consider the nonnegative solutions, denoted $\mathcal{M} = \mathcal{M}(\mathbb{Z}_{\geq 0})$. If $(x, y, z) \in \mathcal{M}$, then it is called a Markoff triple. Each of x, y, and z is a Markoff number; the set of all positive Markoff numbers is $x(\mathcal{M})$.

Fixing the values of y and z, the Markoff equation is quadratic in x. The values of x which form a Markoff triple $(x, y, z) \in \mathcal{M}$ are given by the quadratic equation

$$x = \frac{1}{2}\left(yz \pm \sqrt{y^2z^2 - 4(y^2 + z^2)}\right). \tag{2}$$

These two solutions are mapped to each other via $x \mapsto yz - x$. The *Markoff Vieta involutions* are this map applied to one coordinate of a Markoff triple:

$$V_1 \colon (x, y, z) \mapsto (yz - x, y, z),$$
$$V_2 \colon (x, y, z) \mapsto (x, xz - y, z),$$
$$V_3 \colon (x, y, z) \mapsto (x, y, xy - z).$$

We represent the action of the Markoff Vieta involutions on the set of solutions as the graph $\mathcal{G}$ with vertex set $\mathcal{M}$ and an edge between v_1 and v_2 labeled V_i if $V_i(v_1) = v_2$. Since $(0, 0, 0)$ is fixed by all three involutions, it is an isolated node in $\mathcal{G}$ with three loops. By a result of Markoff, all other nodes are connected to $(3, 3, 3)$ by a unique sequence of V_i's [14]. We represent the *Markoff tree rooted at* $(3, 3, 3)$ as $\mathcal{G}^\times = \mathcal{G} \setminus \{(0, 0, 0)\}$. See Figure 1 for a depiction of $\mathcal{G}^\times$.

We may also consider the Markoff equation modulo a prime $p > 2$,

$$x^2 + y^2 + z^2 - xyz \equiv 0 \bmod p,$$

with solutions $\mathcal{M}(\mathbb{F}_p)$. The values $x(\mathcal{M}(\mathbb{F}_p))$ which appear as solutions is characterized as follows.

Submitted to Algor. Num. Th. Symp.

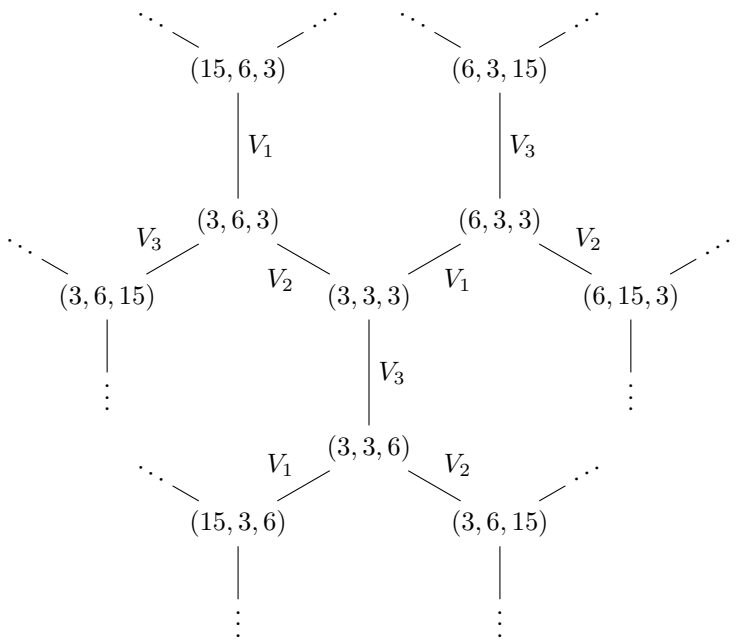

FIGURE 1. The Markoff tree $\mathcal{G}^\times$.

Lemma 1. *Let $x(\mathcal{M}(\mathbb{F}_p)^\times)$ denote the set of integers modulo p which appear in a non-trivial Markoff triple. Then,*

$$x(\mathcal{M}(\mathbb{F}_p)^\times) = \begin{cases} \mathbb{F}_p & -1 \text{ is a quadratic nonresidue modulo } p, \\ \mathbb{F}_p \setminus \{-2, 0, 2\} & \text{otherwise.} \end{cases}$$

Proof. The solutions $(x, y, z) \in \mathcal{M}(\mathbb{F}_p)$ are given by the quadratic equation (2). Without loss of generality, fix $z \in \mathbb{F}_p$. If $z = \pm 2$, then the discriminant of (2) is -16, and the square root lies in $\mathbb{F}_p$ exactly when -1 is a quadratic residue modulo p. Similarily, if $z = 0$, then the discriminant is $-4y^2$, and a solution exists when -1 is a quadratic residue or $y = 0$. But, if $y = z = 0$, then the only solution is the trivial one, $(0, 0, 0)$.

Now, fix $z \notin \{-2, 0, 2\}$. The map

$$y \mapsto y^2z^2 - 4(y^2 + z^2)$$

is 2-to-1 from $\mathbb{F}_p \setminus \{0\}$ to $\mathbb{F}_p$. Including $y = 0$, there are $(p+1)/2$ integers in the image of the map. But since there are only $(p-1)/2$ quadratic non-residues, there must be a y where the discriminant of (2) is a quadratic residue giving a x value for which $(x, y, z) \in \mathcal{M}(\mathbb{F}_p)$. □

As before, $\mathcal{M}(\mathbb{F}_p)$ is fixed by the Markoff Vieta involutions, and the only fixed point is $(0, 0, 0)$. We build a graph $\mathcal{G}_p^\times$ analogous to the Markoff tree with vertex set $\mathcal{M}(\mathbb{F}_p)^\times$ and edges given by the Markoff Vieta involutions. See Figure 2 for $\mathcal{G}_5^\times$ as an example.

Conjecture 1 (Baragar [1])**.** *For every prime p, the Markoff graph $\mathcal{G}_p^\times$ is connected and equivalent to $\mathcal{G}^\times$ with triples congruent modulo p identified.*

Conjecture 2. *For every prime p, the canonical projection $\mathcal{M} \to \mathcal{M}(\mathbb{F}_p)$ is onto.*

Lemma 2. *Conjectures 1 and 2 are equivalent.*

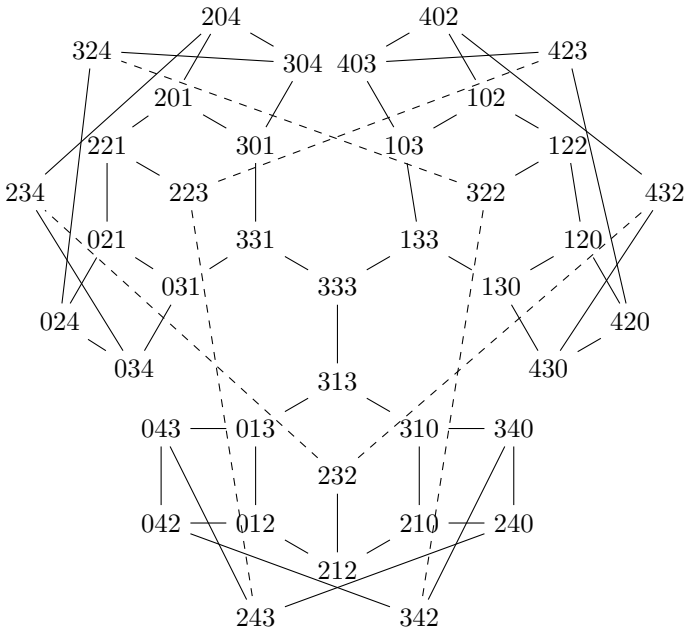

FIGURE 2. The Markoff graph $\mathcal{G}_5^\times$.

Proof. If $\mathcal{G}_p^\times$ is connected, then for any $v \in \mathcal{M}(\mathbb{F}_p)$, there is a sequence $V_{i_1}, \dots V_{i_n}$ forming a path from $(3,3,3)$ to v. Taking the vertices in this path modulo p,

$$(V_{i_1} \circ \cdots \circ V_{i_n})(3,3,3) \equiv v \bmod p.$$

Conversely, assume $\mathcal{M} \to \mathcal{M}(\mathbb{F}_p)$ is onto. For $v \in \mathcal{M}(\mathbb{F}_p)$, let $w \in \mathcal{M}$ such that $w \equiv v \bmod p$. Since $\mathcal{G}^\times$ is connected, there is a path S from $(3,3,3)$ to w. Now, S also forms a walk in $\mathcal{G}_p^\times$, connecting $(3,3,3)$ to v. And since every vertex is connected to $(3,3,3)$, they are all connected. □

Conjecture 2 is often called "the strong approximation conjecture".

Bourgain, Gamburd, and Sarnak developed a three-tiered approach in [3] showing that strong approximation holds for all primes p outside a zero density subset. This result was strengthened by Chen in [4] to strong approximation holding for all but a finite set of primes. Eddy, Fuchs, Litman, Martin, Tripeny, and Vanyo proved an upper bound of approximately $3.448 \cdot 10^{392}$ on p for which $\mathcal{G}_p^\times$ may be disconnected [7]. Their method proves that $\mathcal{G}_p^\times$ is connected for many primes beginning with $p = 1,327,363$, but is inconclusive for most primes less than 10^9. An important piece of their method is the notion of *maximal divisors*, whose connections to Markoff numbers were originally considered in [7]; we reintroduce them in Definition 2.

The strong approximation conjecture implies that solutions to the Markoff equation modulo a prime lift to Markoff triples over $\mathbb{Z}$. Bourgain, Gamburd, and Sarnak use the lifting property to show that almost all Markoff numbers are highly composite [3, Theorem 18]. The strong approximation conjecture is the first step towards proving the much larger conjecture that $\mathcal{G}_p^\times$ forms an expander family, first proposed in [3]. Recently, Fuchs, Lauter, Litman, and Tran explored applications of the Markoff graph being connected to cryptography [8].

In the course of studying the spectral gaps on Markoff graphs in [6], De Courcy-Ireland and Lee performed computations showing that $\mathcal{G}_p^\times$ is connected for all primes $p < 3,000$. However, their method requires calculating the adjacency matrix of $\mathcal{G}_p^\times$, which has on the order of $O(p^4)$ entries, and is therefore infeasible for determining connectivity for larger values of p. Our work aims to fill

in the gap between $p = 3,001$ and $p = 1,327,363$, and provide affirmative answers on the set of primes for which the results of [7] are inconclusive.

In this paper, we outline an algorithmic realization of the process introduced by Bourgain, Gamburd, and Sarnak, along with an implementation of that algorithm in Rust. In particular, we constructively prove the following theorem.

Theorem 1. *There exists an algorithm with runtime $o(p^{1+\epsilon})$ for every $\epsilon > 0$ returning an affirmative or inconclusive response to the query, is the graph $\mathcal{G}_p^\times$ connected?*

Conjecture 3. *The algorithm exhibited here returns an affirmative response for all primes.*

Our algorithm is significantly better than a naive, flood fill-like approach with an $O(p^2)$ runtime, there being $p^2 + O(p)$ solutions to the Markoff equation modulo p [6, Proposition 2.1]. We demonstrate our implementation by testing the connectivity of $\mathcal{G}_p^\times$ for all primes less than one million, and a random sample of primes less than one hundred million.

Theorem 2. *The graph $\mathcal{G}_p^\times$ is connected for all primes p less than one million.*

In the general case of the Markoff equation (1) having an integer on the right hand side, the solutions are still fixed by the Markoff Vieta involutions and rotation maps, but the corresponding graphs are no longer connected. A partial classification of these graphs over the integers was carried out in [9], and Bourgain, Gamburd, and Sarnak announced in [3] a forthcoming article extending their process to these more general graphs. Structures of the connected components for non-zero level sets over finite fields were studied in [6]. Extending the results of the present paper similarily is an interesting direction for future research.

In Section 2, we outline the arguments of [3] and [7], elaborating on the features we will utilize in our algorithm. In Section 3, we construct the data structures and procedures at the core of our algorithm. Finally, in Section 4, we present the algorithm in total, along with our computational results for all primes below one million, as well as auxillary data useful for short circuiting the computations and improving real world performance.

Acknowledgements. Many thanks to my advisor Elena Fuchs for her mentorship, encouragement, and suggestions throughout this paper. I also thank Matt Litman for helpful conversations, especially with regards to Section 2. Finally, thanks to Daniel Martin and Peter Sarnak for helpful feedback on an earlier draft. This paper is based on work supported by the National Science Foundation under grant DMS-2154624.

2. The structure of the Markoff graph modulo p

In this section, we review some of the structural properties of $\mathcal{G}_p^\times$. We begin our discussion with the orbits of the rotation maps, which will form useful walks along the graph.

Let $\mathcal{D}(n)$ be the set of positive divisors of n, and $\mathcal{D}(p \pm 1) = \mathcal{D}(p-1) \cup \mathcal{D}(p+1)$.

We will consider the group Γ generated by the *rotation maps*, which are compositions of a Markoff Vieta involution with a transposition. If $i \in \mathbb{Z}/3\mathbb{Z}$, then

$$\mathrm{rot}_i = \tau_{i+1,i+2} \circ V_{i+1}$$

is the rotation map fixing coordinate i. Writen explicitly,

$$\begin{aligned} rot_1 &\colon (a,b,c) \mapsto (a, c, ac-b), \\ rot_2 &\colon (a,b,c) \mapsto (ab-c, b, a), \\ rot_3 &\colon (a,b,c) \mapsto (b, bc-a, c). \end{aligned}$$

Lemma 3. *If $\mathcal{M}(\mathbb{F}_p)^\times$ is Γ-transitive, then $\mathcal{G}_p^\times$ is connected.*

Proof. Let $G = \langle V_1, V_2, V_3 \rangle$, and let the group S_3 act on $\mathcal{M}(\mathbb{F}_p)$ by permuting the coordinates of each Markoff triple. Since $V_i \circ \sigma = \sigma \circ V_{\sigma(i)}$, for all V_i and $\sigma \in S_3$, the Γ-action on $\mathcal{M}(\mathbb{F}_p)$ factors through $G \rtimes S_3$. In particular, every sequence of rotations is equivalent to a sequence of Markoff Vieta involutions followed by a permutation on the coordinates. Therefore, if $\gamma(x) = y$ for some $\gamma \in \Gamma$, then there is a path $V_{i_1}, \ldots, V_{i_n}$ from x to $\sigma(y)$ for some $\sigma \in S_3$.

Assuming $\mathcal{M}(\mathbb{F}_p)$ is Γ-transitive, there are paths P_1 from x to $(3,3,3)$ and P_2 from $(3,3,3)$ to $\sigma(y)$ for all $x, y \in \mathcal{G}_p^{\times}$. Then, P_1 followed by $\sigma^{-1}(P_2)$ is a path from x to y in $\mathcal{G}_p^{\times}$. □

Fix $(a,b,c) \in \mathcal{M}(\mathbb{F}_p)$. Denote the rotation map rot_1 restricted to acting on triples with fixed first coordinate a by

$$\begin{aligned} \mathrm{rot}'_a \begin{pmatrix} b \\ c \end{pmatrix} &= \begin{pmatrix} c \\ ac - b \end{pmatrix} \\ &= \begin{pmatrix} 0 & 1 \\ -1 & a \end{pmatrix} \begin{pmatrix} b \\ c \end{pmatrix}. \end{aligned}$$

Let $a \neq \pm 2$. If $a = \chi + \chi^{-1}$ for $\chi \in \mathbb{F}_{p^2}$, then the matrix diagonalizes as

$$\begin{pmatrix} 0 & 1 \\ -1 & \chi + \chi^{-1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ \chi & \chi^{-1} \end{pmatrix} \begin{pmatrix} \chi & 0 \\ 0 & \chi^{-1} \end{pmatrix} \begin{pmatrix} 1 & 1 \\ \chi & \chi^{-1} \end{pmatrix}^{-1}. \tag{3}$$

The size of the orbit of rot_1 acting on (a,b,c) is dependent on a only, and will be the same for any orbit with fixed first coordinate a. We therefore refer to the *order of a Markoff number a modulo p* as the size of any orbit with fixed first coordinate a. According to (3), this is equivalent to

$$\mathrm{ord}_p(a) = |\chi|$$

when $a \neq \pm 2$, where $|\cdot|$ is the multiplicative order in $\mathbb{F}_{p^2}^{\times}$.

The following lemma guarantees that the order of χ always divides $p \pm 1$, a fact we will use in our classification of Markoff triples. Equivalently, either $\chi \in \mathbb{F}_p^{\times}$, or $\chi \in \mathbb{F}_{p^2}^{\times}$ with norm 1 to $\mathbb{F}_p$.

Lemma 4. *Let $\chi \in \mathbb{F}_{p^2}^{\times}$. Then $(\chi + \chi^{-1}) \in \mathbb{F}_p$ if and only if $|\chi| \in \mathcal{D}(p \pm 1)$.*

Proof. If $\chi \in \mathbb{F}_{p^2}^{\times}$ with $|\chi| \in \mathcal{D}(p \pm 1)$, then

$$\begin{aligned} \left(\chi + \frac{1}{\chi}\right)^p &= \chi^p + \frac{1}{\chi^p} \\ &= \chi^{\mp 1} + \frac{1}{\chi^{\mp 1}} \\ &= \chi + \frac{1}{\chi}, \end{aligned}$$

and therefore $|\chi + \chi^{-1}| \in \mathcal{D}(p-1)$, equivalently $(\chi + \chi^{-1}) \in \mathbb{F}_p$.

Since $\mathbb{F}_{p^2}^{\times}$ is cyclic, there are $2p$ values $\chi \in \mathbb{F}_{p^2}^{\times}$ for which $|\chi|$ is divisible by $p-1$ or $p+1$. The map $\phi : \chi + \chi^{-1}$ is 2-to-1 on $\mathbb{F}_{p^2}^{\times}$, so the image of ϕ on the set $\{\chi : |\chi| \in \mathcal{D}(p \pm 1)\}$ is p values in $\mathbb{F}_p$. But, that is all of $\mathbb{F}_p$, so there must be no more values of $\chi \in \mathbb{F}_{p^2}^{\times}$ for which $\chi + \chi^{-1} \in \mathbb{F}_p$. □

The diagonalization (3) allows us to parameterize the Markoff triples in an orbit in terms of χ. Specifically, if $(a,b,c) \in \mathcal{M}(\mathbb{F}_p)$ and $a = \chi + \chi^{-1}$, then

$$\langle \mathrm{rot}'_a \rangle (b,c) = \left\{ k \left(\alpha \chi^{\ell} + \beta \frac{1}{\chi^{\ell}}, \alpha \chi^{\ell+1} + \beta \frac{1}{\chi^{\ell+1}} \right) : \ell \in \mathbb{Z} \right\},$$

where

$$k = \left(\chi - \frac{1}{\chi}\right)^{-1}, \ \alpha = \left(c - \frac{b}{\chi}\right), \text{ and } \beta = (\chi b - c).$$

A parameterization for arbitrary orbits is given in [7, Equation 3] as

$$\left\{\left(\chi+\frac{1}{\chi},\frac{\chi+\chi^{-1}}{\chi-\chi^{-1}}\left(r\chi^{\ell}+\frac{1}{r\chi^{\ell}}\right),\frac{\chi+\chi^{-1}}{\chi-\chi^{-1}}\left(r\chi^{\ell+1}+\frac{1}{r\chi^{\ell+1}}\right)\right):\ell\in\mathbb{Z}\right\}, \tag{4}$$

where any values $r,\chi\in\mathbb{F}_{p^2}^{\times}\setminus\{\pm1\}$ gives a set of triples solving the Markoff equation over $\mathbb{F}_{p^2}$ fixed by rot_1. For fixed $\chi\in\mathbb{F}_{p^2}^{\times}$, the orbits given by $r_1,r_2\in\mathbb{F}_{p^2}^{\times}$ according to (4) will be the same exactly when $r_1\langle\chi\rangle=r_2\langle\chi\rangle$. We therefore seek one representative r for each coset of $\langle\chi\rangle$ giving

$$\frac{r+r^{-1}}{\chi-\chi^{-1}}\in\mathbb{F}_p,$$

so that our expression (4) gives solutions to the Markoff equation in $\mathbb{F}_p$. The solutions are characterized by Lemma 5.

Lemma 5. *Let $\chi,r\in\mathbb{F}_{p^2}^{\times}$ and $\chi\neq\pm1$. Let $k\in\mathbb{F}_p$ be a quadratic nonresidue.*

(a) If $\chi\in\mathbb{F}_p^{\times}$, then $\chi-\chi^{-1}\in\mathbb{F}_p$.
(b) If $|\chi|\in\mathcal{D}(p+1)$, then $\chi-\chi^{-1}\in\sqrt{k}\mathbb{F}_p$.
(c) If $|r|\in\mathcal{D}(2(p+1))\setminus\mathcal{D}(p+1)$, then $r+r^{-1}\in\sqrt{k}\mathbb{F}_p$.

Proof. The proof of (a) is identical to the proof of Lemma 4 where $|\chi|\in\mathcal{D}(p-1)$. So, assume $|\chi|\in\mathcal{D}(p+1)$. Then,

$$\left(\chi-\frac{1}{\chi}\right)^{p}=\frac{1}{\chi}-\chi,$$

therefore $|\chi-\chi^{-1}|\notin\mathcal{D}(p-1)$. However,

$$\begin{aligned}\left(\chi-\frac{1}{\chi}\right)^{2p}&=\left(\chi^2+\frac{1}{\chi^2}-2\right)^{p}\\&=\left(\frac{1}{\chi^2}+\chi^2-2\right),\end{aligned}$$

so $(\chi-\chi^{-1})^2$ has multiplicative order dividing $p-1$. Finally, let $|r|\in\mathcal{D}(2(p+1))$. Then,

$$\begin{aligned}\left(r+\frac{1}{r}\right)^{2p}&=\left(r^2+\frac{1}{r^2}+2\right)^{p}\\&=\left(\frac{1}{r^2}+r^2+2\right),\end{aligned}$$

therefore $(r+r^{-1})^2$ has multiplicative order dividing $p-1$. □

Corollary 1. *Let $\chi,r\in\mathbb{F}_{p^2}^{\times}$ with $|\chi|\in\mathcal{D}(p+1)$, $|r|\in\mathcal{D}(2(p+1))\setminus\mathcal{D}(p+1)$, and $\chi\neq\pm1$. Then, $\chi-\chi^{-1}=b_1\sqrt{k}$ and $r+r^{-1}=b_2\sqrt{k}$ for $b_1,b_2,k\in\mathbb{F}_p$ and k a quadratic nonresidue. With these choices of χ and r, the orbit (4) contains triples with all coordinates in $\mathbb{F}_p$.*

In the spirit of [3], we say that an element $a\in\mathbb{Z}/p\mathbb{Z}$ is *parabolic* if $a=\pm2$. Otherwise, it is *hyperbolic* if $\mathrm{ord}_p(a)$ divides $p-1$, or *elliptic* if $\mathrm{ord}_p(a)$ divides $p+1$.

Remark 1. The orbits of triples with fixed parabolic coordinates have sizes $\mathrm{ord}_p(a)=p$ or $2p$. If a is parabolic and $a=\chi+\chi^{-1}$, then $\chi=\pm1$. The matrix in (3) is not invertible, and $|\chi|$ is 1 or 2.

Definition 1. The *order of the triple* $(a,b,c)\in\mathcal{M}(\mathbb{F}_p)$ is

$$\mathrm{ord}_p((a,b,c))=\max(\mathrm{ord}_p(a),\mathrm{ord}_p(b),\mathrm{ord}_p(c)).$$

If the order of a triple is $p\pm1$, then it is called *a triple of maximal order.*

Bourgain, Gamburd, and Sarnak showed in [3] that $\mathcal{G}_p^\times$ is connected for every prime p outside a zero density subset of the primes. They did so by first observing that all triples of maximal order and parabolic triples belong to the same connected component $\mathcal{C}_p$. (An explicit proof that triples with parabolic coordinates are connected to triples of maximal order is given in [2, Lemma 3.3].) Then, they iteratively lower the bound on the orders of triples which must be connected to $\mathcal{C}_p$. In the first step, the *endgame*, all triples of order at least $p^{\delta+1/2}$ are shown to be connected to $\mathcal{C}_p$ for some δ dependent on p. In the *middle game*, they show that there is an $\epsilon > 0$ dependent only on p for which $\mathrm{ord}_p(a) > p^\epsilon$ implies that every orbit of rot'_a contains a triple of order larger than $\mathrm{ord}_p(a)$. By successively walking from triple to triple of higher order, eventually a triple of order at least $p^{\delta+1/2}$ is reached. Finally, in the *opening*, they bound the size of $|\mathcal{G}_p \setminus \mathcal{C}_p|$ from below, and show that for primes p with $p \pm 1$ not having too many factors, there can not be enough unaccounted for triples to constitute a connected component disconnected from $\mathcal{C}_p$. Specifically, Bourgain, Gamburd, and Sarnak proved in [3] a lower bound of $(\log p)^{1/3}$ on the size of any connected component. This bound was improved by Konyagin, Makarychev, Shparlinski, and Vyugin [12] to $(\log p)^{7/9}$, and again improved by Chen [4] to the size being divisible by p.

The bounds on the endgame were made explicit in [7], where the following inequality was proved.

Lemma 6 ([7, Proposition 6.1])**.** *A Markoff triple of order $d \in \mathcal{D}(p \pm 1)$ is connected to $\mathcal{C}_p$ provided*

$$d > \frac{8\sqrt{p}(p \pm 1)\tau(p \pm 1)}{\phi(p \pm 1)}. \tag{5}$$

We will call the right hand side of (5) for each choice of sign the *endgame breakpoints* and denote them $B_{(p,\pm)}$.

We now wish to bound from below the orders of triples which are part of the middle game. We will consider an orbit of size t, and give an upper bound on the number of triples in that orbit with order less than t. If the number of such triples is less than t, then the orbit being considered must contain a triple of order larger than t. A suitable bound is given in [7] and relies on an approximation of Corvaja and Zannier [5].

Lemma 7 ([7, Lemma 2.1])**.** *If $\chi \in \mathbb{F}_{p^2}^\times$ has order $t > 2$, then the number of values n between 0 and $t-1$ for which*

$$\mathrm{ord}_p\left(\frac{\chi + \chi^{-1}}{\chi - \chi^{-1}}\right)\left(sr^n + (sr^n)^{-1}\right)$$

divides d is at most

$$\frac{3}{2}\max\left(\sqrt[3]{6td}, \frac{4td}{p}\right). \tag{6}$$

We now introduce the notion of maximal divisors, coined in [7].

Definition 2 ([7, Definition 1.2])**.** Let n be a positive integer, and let $x \in \mathbb{R}$. A positive divisor d of n is *maximal with to respect to x* if $d \le x$ and there is no other positive divisor d' of n such that $d' \le x$ and $d \mid d'$. The set of maximal divisors of n with respect to x is denoted $\mathfrak{M}_x(n)$.

Fix a divisor t of $p \pm 1$. If we sum (6) over all maximal divisors of $p-1$ and $p+1$ with respect to t, then we have an upper bound on the number of triples with order less than t in each orbit of order t. Explicitly, if

$$t > \sum_{d \in \mathfrak{M}_t(p \pm 1)} \frac{3}{2}\max\left(\sqrt[3]{6td}, \frac{4td}{p}\right), \tag{7}$$

then every orbit of order t is guaranteed to have a triple of order larger than t. (In fact, this is still double counting elements of order dividing more than one maximal divisor; sharpening this bound is a direction for future work.)

Let L_p be the smallest $d \in \mathcal{D}(p \pm 1)$ such that every $d' \ge d$, with $d' \in \mathcal{D}(p \pm 1)$, satisfies (7). If it exists, then we call L_p the *middle game breakpoint*.

Lemma 8. *Every triple of order at least L_p is connected to $\mathcal{C}_p$.*

Proof. Let $v_1 \in \mathcal{M}(\mathbb{F}_p)^\times$ with order $d > L_p$. Then, v_1 is connected to some triple v_2 of order $d' \gtrsim d$. This argument may be repeated until a triple of order $p \pm 1$ is found, at which point we have reached $\mathcal{C}_p$. □

Algorithm 1 shows a procedure for calculating L_p.

Algorithm 1: Finding the middle game breakpoint L_p.

Input: prime p.
Result: L_p, either a real number or **None**.
$L_p \leftarrow$ **None**;
for t **in** $\mathcal{D}(p \pm 1) \setminus \{2\}$ *in ascending order* **do**
  $\mathsf{sum} \leftarrow 0$;
  for d **in** $\mathcal{D}$, $d < t$ **do**
    $\mathsf{sum} \leftarrow \mathsf{sum} + \max(\sqrt[3]{6dt}, 4dt/p)$;
  end
  if $t \geq \mathsf{sum}$ **and** L_p **is None then**
    $L_p \leftarrow t$;
  end
  else if $t < \mathsf{sum}$ **then**
    $L_p \leftarrow$ **None**;
  end
end
return L_p;

In [4], Chen showed that the size of any connected component disconnected from $\mathcal{C}_p$ is divisible by p. This result was used in [7] to show the following.

Lemma 9 ([7, Lemma 2.2])**.** *If $p > 3$, then $|\mathcal{G}_p \setminus \mathcal{C}_p|$ is divisible by $4p$.*

Lemma 6 and Corollary 8 give two bounds on the orders of triples in $\mathcal{C}_p$. Our goal is now to count the number of triples which we have not yet shown are connected to $\mathcal{C}_p$. If the number of triples thus counted is less than $4p$, then $\mathcal{G}_p$ is connected by Lemmas 3 and 9.

Definition 3. Let $\phi : \mathcal{M}(\mathbb{F}_p) \to \{0, 1\}$ be the boolean function

$$\phi(a) = \begin{cases} 1 & a \text{ is hyperbolic and } \operatorname{ord}_p(a) < \min(L_p, B_{(p,-)}), \\ 1 & a \text{ is elliptic and } \operatorname{ord}_p(a) < \min(L_p, B_{(p,+)}), \\ 0 & \text{otherwise,} \end{cases}$$

where L_p and $B_{(p,\pm)}$ are the middle game and endgame breakpoints, respectively. A coordinate a has *small order* if $\phi(a) = 1$. The set of coordinates with small order is denoted $\mathcal{S}$.

Definition 4. Let φ be a boolean function $\mathcal{M}(\mathbb{F}_p) \to \{0, 1\}$. A triple $v \in \mathcal{M}(\mathbb{F}_p)$ is a *bad triple* if $\varphi(v) = 0$. The set of bad triples is denoted $\mathcal{B}$.

We now wish to construct a boolean function φ with two competing mandates. First, φ must be as sensitive as possible; that is, we wish to minimize the number of triples v connected to $\mathcal{C}_p$ but for which $\varphi(v) = 0$. Second, we would like φ to be easy to compute, in the sense that a computer can quickly identify bad triples with respect to φ. To this end, we have two strategies:

Strategy 1. For each $(a, b) \in \mathcal{S} \times \mathcal{S}$, determine the zero, one, or two values of c for which $(a, b, c) \in \mathcal{M}(\mathbb{F}_p)$ using (2). Any triple found this way for which a, b, and c all have small order is a bad triple.

Strategy 2. For each $a \in \mathcal{S}$ and $a = \chi + \chi^{-1}$, and for each coset of $\langle\chi\rangle$, pick a representative $r \in \mathbb{F}_{p^2}^\times/\langle\chi\rangle$ and calculate the entries in the orbit of rot'_a given by (4). Note that $r \in \mathbb{F}_p$ if a is hyperbolic, and $r \in \mathbb{F}_{p^2}^\times$ with order dividing $2(p+1)$ if a is elliptic; see Corollary 1. Check a set number of triples in the orbit[1]; if they are all of small order, then every triple in the orbit is bad. Otherwise, none of the triples in the orbit are bad.

For every $a \in \mathbb{F}_p$, we run one of the two above strategies to identify all bad triples with first coordinate a. We choose the strategy requiring fewer checks, i.e., we choose Strategy 1 when $|\mathcal{S} \times \mathcal{S}| < (p \pm 1)/|\chi|$, and Strategy 2 otherwise.

Remark 2. Strategy 2 is more sensitive, since a triple may have all coordinates with small order and yet still be "good". It is possible that, for some prime p, choosing to only run Strategy 2 could affirmatively show that $\mathcal{G}_p$ is connected while sometimes running Strategy 1 is inconclusive.

Strategy 2 has a second advantage; it will only ever test Markoff triples. Strategy 1, on the other hand, must filter through (a, b) pairs for which no value of c gives a Markoff triple. The runtime cost of performing this filter is similar to the weakening of the estimations in [7] caused by assuming each pair (a, b) has two values of c making $(a, b, c) \in \mathcal{M}(\mathbb{F}_p)$. Nevertheless, when the middle game breakpoint or $|\chi|$ is small, Strategy 1 may still be faster.

3. Computational Machinery

In this section, we describe the data structures and procedures needed to implement the strategies described in Section 2. We begin with the *factor trie*[2], encoding a canonical poset on $\mathcal{D}(p \pm 1)$. This poset will allow us to recursively generate only those elements $a \in \mathbb{F}_p$ satisfying $\mathrm{ord}_p(a)$ less than some upper bound (or a being parabolic), without having to filter the entirety of $\mathbb{F}_p$ by order.

Associate a word $w(n)$ to each $n = p_1^{t_1} \cdots p_n^{t_n}$ with primes in increasing order given by

$$w(n) = \underbrace{p_1 \dots p_1}_{t_1} \cdots \underbrace{p_n \dots p_n}_{t_n}. \tag{8}$$

Let $\mathcal{D}(n)$ be the set of positive divisors of n. Define a partial ordering $\prec$ on $\mathcal{D}(p \pm 1)$ as $n \prec m$ if the word $w(n)$ is a prefix for the word $w(m)$. If $n \prec m$ and there is no ℓ such that $n \prec \ell \prec m$, then we write $n \prec\cdot\, m$.

Define the factor trie $T_n = (V, E)$ as the graph with vertex set $V = \{w(d) : d \in \mathcal{D}(n)\}$ and directed edge set $E = \{(w(\ell), w(m)) : \ell \prec\cdot\, m\}$. When the context is clear, we will identify $\mathcal{D}(n)$ and $\prec\cdot$ with the vertex and edge sets directly. An example factor trie for $n = 60$ is shown in Figure 3.

Associate to each vertex of T_n a finite cyclic group via the covariant functor $\mathcal{F}_n : T_n \to \mathbf{Ab}$ defined by $\mathcal{F}_n : w(d) \mapsto \mathbb{Z}/d\mathbb{Z}$ and taking the directed edge $\ell \prec\cdot\, m$ to the inclusion map $\mathbb{Z}/\ell\mathbb{Z} \hookrightarrow \mathbb{Z}/m\mathbb{Z}$. Note that there are no intermediate subgroups between these groups.

The tries $T_{p\pm1}$ will form the call graphs for our recursive procedure, with vertex $d \notin \{1, 2\}$ generating the elements $\{x \in \mathbb{Z}/p\mathbb{Z} : \mathrm{ord}_p(x) = d\}$; these are exactly the generators of $\mathcal{F}_{p\pm1}(d)$. (If $d \in \{1, 2\}$, then $x = \pm 2$ and is parabolic, see Remark 1.) If $d_1 \prec d_2$, then $|\mathcal{F}_{p\pm1}(d_1)| < |\mathcal{F}_{p\pm1}(d_2)|$, so our x values are generated with increasing $\mathrm{ord}_p(x)$ moving down the call graph induced by $T_{p\pm1}$, and we can stop the procedure when we have reached the desire bound.

We could represent elements of the group $\mathcal{F}_{p\pm1}(d)$ as integers modulo d. However, it will be more convenient to represent them instead as integer arrays based on the unique decomposition of these groups into a direct sum of prime power cyclic groups. The multiplicative group operations can then

[1]See Figure 5 for the maximum number of triples checked per orbit. We cap the number of triples checked to prevent our program hanging on a prime for which an exceptional number of checks must be performed. In practice, we found no such exceptional prime out of all primes less than one million, or among those randomly sampled less than 110,000,000.

[2]A *trie* is a tree where the vertex set is the set of unique paths from the root. Some authors use this as the definition for a tree; others define a tree as a connected acyclic graph. We use the term *trie* to emphasize the word structure of the graph.

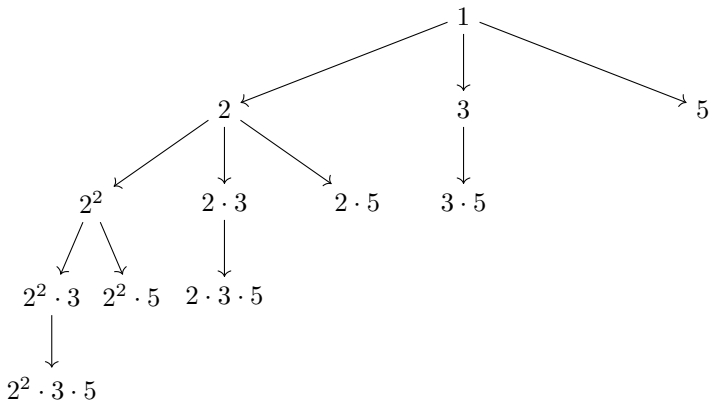

FIGURE 3. The factor trie for $n = 60$.

be performed by the CPU using only componentwise addition and subtraction, which are much faster than direct multiplication and reciprocal taking. One drawback is that addition and subtraction are not easily represented, so a conversion back to $\mathbb{F}_p$ is necessary before any ring operations. We now describe this representation.

Let $p-1 = p_1^{t_1} \cdots p_n^{t_n}$ and $p+1 = q_1^{s_1} \cdots q_m^{s_m}$, and assume $p_1 = q_1 = 2$.

Fix, once and for all, a $\mathbb{Z}$-basis of $\mathbb{F}_p^\times$ of the form $\{g_i\}_{i=1}^n$, where $|g_i| = p_i^{t_i}$. Making this choice is a special case of [13, Algorithm 6.1], which is our preferred method. This choice induces a group isomorphism

$$\iota_{p-1} : \mathbb{F}_p^\times \xrightarrow{\sim} \bigoplus_{i=1}^n \mathbb{Z}/p_i^{t_i}\mathbb{Z}$$

given by

$$\prod_{i=1}^n g_i^{r_i} \mapsto (r_1, \ldots, r_n). \tag{9}$$

An identical construction exists for the elliptic coordinates by replacing $\mathbb{F}_p^\times$ with the subgroup $E \subset \mathbb{F}_{p^2}^\times$ with order $p+1$. Taken together, we have the commutative diagram

$$\begin{array}{ccccc} & & \mathbb{Z}/p\mathbb{Z} & & \\ & \overset{\psi_{p-1}}{\nearrow} & \uparrow & \overset{\psi_{p+1}}{\nwarrow} & \\ \mathbb{F}_p^\times & \xhookrightarrow{\iota_{p-1}} & \left(\bigoplus_{i=1}^n \mathbb{Z}/p_i^{t_i}\mathbb{Z}\right) \amalg \left(\bigoplus_{i=1}^m \mathbb{Z}/q_i^{s_i}\mathbb{Z}\right) & \xhookleftarrow{\iota_{p+1}} & E \end{array}$$

where $\psi_{p\pm1}(\chi) = \chi + \chi^{-1}$ (see Lemma 4). To the vertex $d = p_1^{d_1} \cdots p_n^{d_n} \in \mathcal{D}(p-1)$ we have now associated both the subgroup of $\mathbb{F}_p^\times$ of order d and the group

$$\mathcal{F}_{p-1}(d) = \mathbb{Z}/d\mathbb{Z} \cong \bigoplus_{i=1}^n \mathbb{Z}/p_i^{d_i}\mathbb{Z}$$

containing integer arrays given by (9), and similarly for the divisors of $p+1$.

Remark 3. Since $\psi_{p\pm1}(\chi) = \psi_{p\pm1}(\chi^{-1})$, the maps ψ_{p-1} and ψ_{p+1} are 2-to-1 when restricted to $\mathbb{F}_p^\times \setminus \{\pm1\}$ and $E \setminus \{\pm1\}$, respectively. Since $\psi_{p-1}(\pm1) = \psi_{p+1}(\pm1)$, the induced (dotted) map in the commutative diagram is 2-to-1.

We now describe the recursive procedure for generating a value $a \in \mathcal{S}$ (see Definition 3) with order $\mathrm{ord}_p(a) = d$. We begin with the integer array $(0, 0, \ldots, 0)$ associated to $w(1)$, and propagate

this array down the factor trie recursively, with each node $w(d)$ yielding[3] all the coordinates $a \in \mathbb{F}_p$ with $\mathrm{ord}_p(a) = d$, mapping the integer array to a via the inverse map of (9).

Remark 4. Parabolic triples are connected to $\mathcal{C}_p$ by definition, so we do not want to generate χ values for which $|\chi| = d \in \{1, 2\}$ (see Remark 1).

In light of Remark 3, we would like to yield only one representative from each $\{\chi, \chi^{-1}\}$ pair. The following definition and lemma provides a characterization for making this choice canonical.

Definition 5. Let $p \pm 1 = p_1^{t_1} \cdots p_n^{t_n}$ with $p_1 = 2$, and let $(r_1, \ldots, r_n) \in \oplus_{i=1}^n \mathbb{Z}/p_i^{t_i}\mathbb{Z}$ be an integer array according to (9) not representing a $\chi \in \mathbb{F}_{p^2}^\times$ with order 1 or 2. Let

$$k = \min\{j : r_j \neq 0 \text{ and } (p_j, r_j) \neq (2, 2^{t_1-1})\}.$$

The integer array is *in the lower half* if $r_k \leq p_k^{t_k}/2$.

Lemma 10. *Let $\chi \in \mathbb{F}_{p^2}^\times$ such that $|\chi| \notin \{1, 2\}$ and $|\chi| \in \mathcal{D}(p \pm 1)$. Then, exactly one of $\iota_{p\pm1}(\chi)$ or $\iota_{p\pm1}(\chi^{-1})$ is in the lower half.*

Proof. Consider a hyperbolic $a \in \mathbb{F}_p$ with $a = \chi + \chi^{-1}$ and $\chi \in \mathbb{F}_p^\times$. (The elliptic case is handled identically.) Let $\chi = g_1^{r_1} \cdots g_n^{r_n}$. Since $\iota_{p-1}(\chi^{-1}) = (a_1, \ldots, a_n)$ where

$$a_i = (p_i^{t_i} - r_i) \bmod p_i^{t_i},$$

if $|\chi| \notin \{1, 2\}$, then at least one of r_k and $p_k^{t_k} - r_k$ is less than or equal to $p_k^{t_k}/2$. Therefore, at least one of $\iota_{p-1}(\chi)$ and $\iota_{p-1}(\chi^{-1})$ is in the lower half.

Now assume that both χ and χ^{-1} are in the lower half. For each j such that $p_j \neq 2$, it must be that $r_j = 0$. But then our array is either $(0, \ldots, 0)$ or $(2^{t_1-1}, \ldots, 0)$, rearranging if necessary so $p_1 = 2$, and these arrays correspond to $|\chi| = 1$ and $|\chi| = 2$, respectively. □

We will only yield integer arrays in the lower half, thus generating each desired element of $\mathbb{F}_p$ exactly once. We do this by setting upper bounds on the entries in the arrays yielded by node $w(d)$ for $d = p_1^{d_1} \cdots p_m^{d_m}$ according to

$$r_j \leq \ell_{d,j} = \begin{cases} p_j^{t_j}/2 & j = \min\{j : p_j^{d_j} \neq 2\}, \\ p_j^{t_j} & \text{otherwise.} \end{cases} \tag{10}$$

We now present our algorithm for yielding Markoff coordinates $a \in \mathcal{S}$ in Algorithm 2. Each value of χ yielded corresponds to a Markoff coordinate $\chi + \chi^{-1}$.

Besides generating all $a \in \mathcal{S}$, we can use the factor trie to generate cosets of $\langle\chi\rangle$ for $\chi \in \mathbb{F}_{p^2}^\times$, as in Strategy 2. For this we need one representative for each coset of $\langle\chi\rangle$. That is, for $|\chi| \in \mathcal{D}(p \pm 1)$, we need one value of order d for each divisor d of $(p \pm 1)/|\chi|$. Let $|\chi| = p_1^{s_1} \cdots p_n^{s_n}$. We can repurpose Algorithm 2 by instead interpreting the yielded values as r in (4), and replacing the upper bounds in (10) with

$$\ell_{d,j} = \begin{cases} 0 & d_j \leq s_j, \\ p_j^{t_j}/2 & d_j > s_j \text{ and } j = \min\{j : p_j^{d_j} \neq 2\}, \\ p_j^{t_j} & \text{otherwise.} \end{cases} \tag{11}$$

[3]A `yield` statement is similar to a `return` statement in that it passes data back to the caller of the function; it differs in that execution continues in the callee, allowing for more than one value to be yielded.

Algorithm 2: Generating χ corresponding to desired Markoff coordinates $\chi + \chi^{-1} \in \mathcal{S}$.

Input: Factor trie $T_{p\pm1}$ with $p \pm 1 = p_1^{t_1} \cdots p_n^{t_n}$,
Limits $\ell_1, \ldots, \ell_n$ as in (10),
Node $w(d)$ with $d = p_1^{d_1} \cdots p_m^{d_m}$,
Array of integers $r_1, \ldots, r_m$.
Base case: `propagate(`$w(1)$`,` $0, \ldots, 0$`)`
Result: Stream of χ values.
fn `propagate(`$w(d)$`,` $r_1, \ldots, r_m$`)`

```
for i in 0 to p_m do
    r'_m ← r_m + p_m^{t_m − d_m};
    if r'_m > ℓ_{d,m} then
        break;
    end
    if d_m < t_m then
        propagate(w(p_1^{d_1} ··· p_m^{d_m+1}), r_1, ..., r'_m);
    end
    if i = 0 then
        /* Yielding when i = 0 would not respect the trie structure, i.e.,
           there would be multiple call paths to yield the same value.    */
        continue;
    end
    if 2 < d < min(L_p, B_{(p,±)}) then
        yield g_1^{r_1} ··· g_{m−1}^{r_{m−1}} · g_m^{r'_m};
    end
    for j in m + 1 to n do
        if d_j < t_j and w(d · p_j) or there is a d' ≻ d such that d' < min(L_p, B_{(p,±)}) then
            propagate(w(d · p_j), r_1, ..., r'_m);
        end
    end
end
```

4. Our Algorithm & Results

In this section, we combine the process described in Section 2 with the machinery constructed in Section 3 to produce a comprehensive algorithm for proving that $\mathcal{G}_p^{\times}$ is connected. The code described in this section can be found at `https://github.com/colbyaustinbrown/libbgs`, with executable code found in the `examples` directory. We used the Rust programming language, version 1.78, for our `libbgs` library defining the data structures described in Section 3. Rust is a strongly, statically typed imperative language prioritizing speed, memory safety, and concurrency. We now present our algorithm for testing whether $\mathcal{G}_p^{\times}$ is connected.

Algorithm 3 (Testing $\mathcal{G}_p$ for connectivity)**.**

(1) Calculate the prime factorizations for $p \pm 1$, and construct the factor tries $T_{p\pm1}$. Also, fix $\mathbb{Z}$-bases of $\mathbb{F}_p^\times$ and E for (9), and also a value m with order $2^t || p^2 - 1$ (see Corollary 1) using [13, Algorithm 6.1]. Note that the choice of basis requires $O(\log p)$ choices via randomly sampling elements of $\mathbb{F}_{p^2}$, with each choice being made within an expected 2 samples.
(2) Use (5) and Algorithm 1 to determine the endgame and middle game breakpoints.
(3) Generate all coordinates with small order $a \in \mathbb{F}_p$ according to Definition 3 via Algorithm 2, and proceed to step 4 for each a.
(4) If $(p \pm 1)/\operatorname{ord}_p(a) < |\mathcal{S} \times \mathcal{S}|$, go to step 4(b); otherwise, go to step 4(a).
   (a) Perform Strategy 1: Use Algorithm 2 to generate all $b \in \mathbb{F}_p$ with small orders. For each b value, calculate the 0, 1, or 2 values of c for which $(a, b, c) \in \mathcal{G}_p$. Any c value with small order gives a bad triple (a, b, c).
   (b) Perform Strategy 2: Generate r values given by Algorithm 2 modified with (11), and use them to calculate up to (fixed) n values b from (4). (If a is elliptic, the values r must be multiplied by the value m chosen in Step 1 before being plugged in to (4), see Corollary 1.) If every b has order smaller than L, then every triple in the orbit is bad.
(5) Count the number of bad triples found. If the result is less than $4p$, then $\mathcal{G}_p$ is connected. Otherwise, the algorithm is inconclusive.

Using Algorithm 3, we confirmed that $\mathcal{G}_p$ is connected for all $p < 1,000,000$, confirming Theorem 2. We also ran our algorithm for a random sample of 1,000 primes $p < 110,000,000$, and $\mathcal{G}_p^\times$ was connected for each of these primes, too. We did not find any $\mathcal{G}_p$ with at least p bad triples, with two exceptions: $p = 7,558,541$ and $p = 96,840,901$, for which the number of bad triples found was 9,716,411 and 103,370,751, respectively. The criteria for connectivity of [4] was nearly always sufficient in our tests, although we did rely on the improved bounds of [7] that any $|\mathcal{G}_p^\times \setminus \mathcal{C}_p| \geq 4p$. The number of bad triples per prime is shown in Figure 4. Table 1 shows the prime factorizations of $p \pm 1$ and bad triple counts for a sample of primes less than one million. Generally, the number of bad triples is inversely correlated with the number of prime divisors (with multiplicity) of $p \pm 1$.

TABLE 1. For $p < 1,000,000$, the 10 primes with the largest ratio $|\mathcal{B}|/p$, and the 5 largest primes with $\mathcal{B} = \emptyset$.

	Hyperbolic		Elliptic	
p	$p-1$	Bad Triples	$p+1$	Bad Triples
825,287	$2 \cdot 7 \cdot 11 \cdot 23 \cdot 233$	277,287	$2^3 \cdot 3 \cdot 137 \cdot 251$	320,209
916,879	$2 \cdot 3 \cdot 17 \cdot 89 \cdot 101$	251,391	$2^4 \cdot 5 \cdot 73 \cdot 157$	425,410
804,203	$2 \cdot 7 \cdot 17 \cdot 31 \cdot 109$	295,979	$2^2 \cdot 3^2 \cdot 89 \cdot 251$	286,714
936,259	$2 \cdot 3 \cdot 17 \cdot 67 \cdot 137$	307,155	$2^2 \cdot 5 \cdot 13^2 \cdot 277$	362,722
734,803	$2 \cdot 3 \cdot 29 \cdot 41 \cdot 103$	171,268	$2^2 \cdot 7^2 \cdot 23 \cdot 163$	351,854
550,811	$2 \cdot 5 \cdot 13 \cdot 19 \cdot 223$	211,593	$2^2 \cdot 3 \cdot 197 \cdot 233$	168,181
858,701	$2^2 \cdot 5^2 \cdot 31 \cdot 277$	279,547	$2 \cdot 3 \cdot 13 \cdot 101 \cdot 109$	304,704
843,229	$2^2 \cdot 3^2 \cdot 59 \cdot 397$	320,154	$2 \cdot 5 \cdot 37 \cdot 43 \cdot 53$	250,655
995,677	$2^2 \cdot 3 \cdot 11 \cdot 19 \cdot 397$	0	$2 \cdot 497,839$	0
995,987	$2 \cdot 497,993$	0	$2^2 \cdot 3 \cdot 7 \cdot 71 \cdot 167$	0
996,783	$2^2 \cdot 3 \cdot 5 \cdot 37 \cdot 449$	0	$2 \cdot 498,391$	0
997,583	$2 \cdot 498,791$	0	$2^4 \cdot 3 \cdot 7 \cdot 7,969$	0
999,958	$2 \cdot 499,979$	0	$2^3 \cdot 3 \cdot 5 \cdot 13 \cdot 641$	0

In the worst case, steps 1 and 2 require on the order of $|\mathcal{D}(p \pm 1)| + |\mathcal{D}(p+1)|$ computations. Step 4 executes at most p times, and each iteration is bounded above by the runtime of step 4(b),

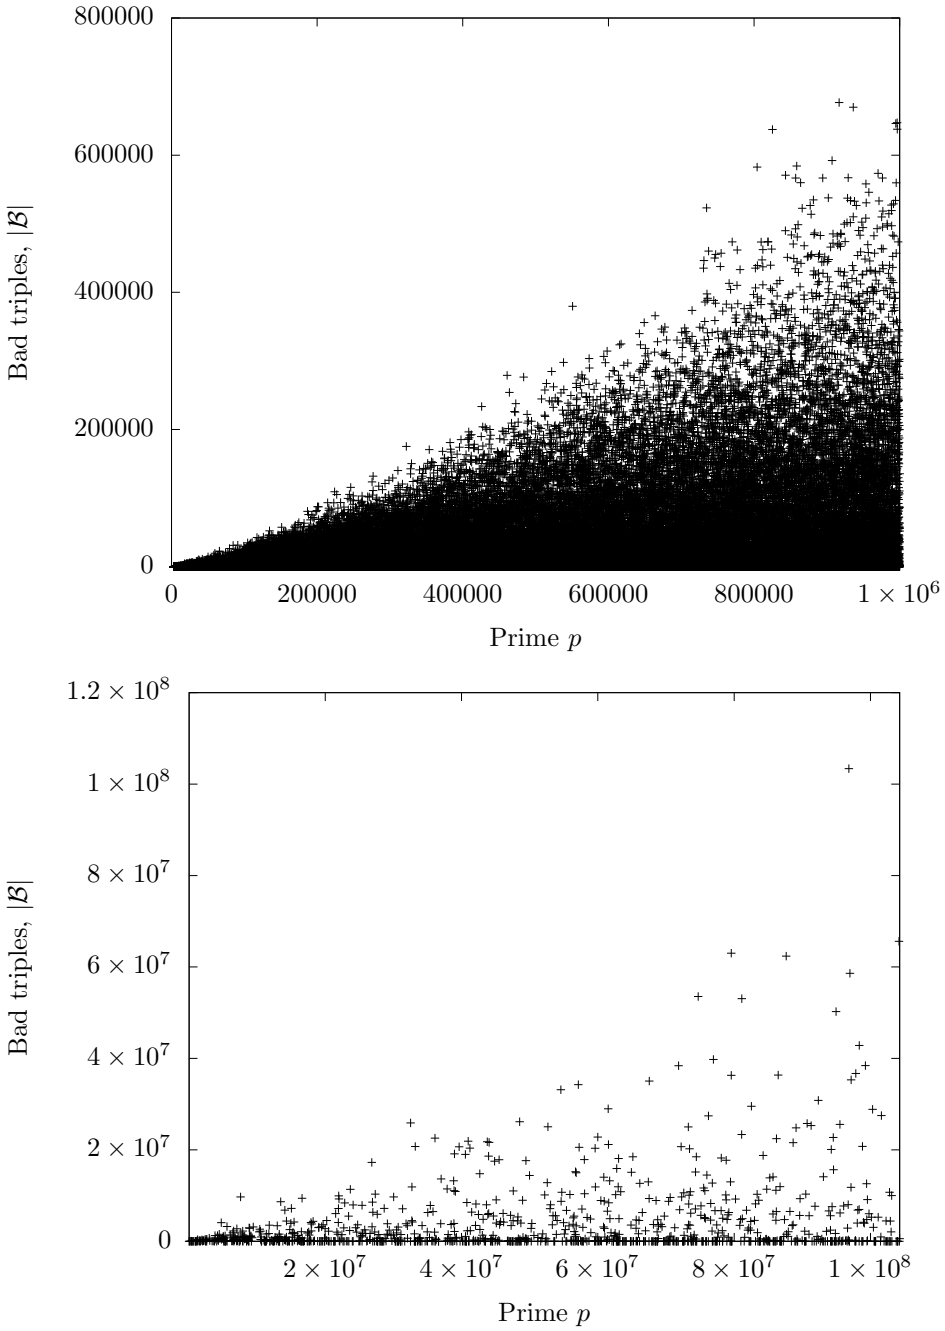

FIGURE 4. Total number of bad triples for all primes less than $1,000,000$ (top) and random primes less than $110,000,000$ (bottom).

which checks at most $\mathcal{D}(p \pm 1)$ orbits. The check on each orbit terminates after finding any triple not in $\mathcal{B}$; the largest number of such checks per prime is shown in Figure 5.

For the cap given in step 4(b), we used $n = 60$. In our execution of Algorithm 3, step 4(b) was terminated after 60 checks for only one prime, $p = 119,659$, and the algorithm still affirmatively

determined $\mathcal{G}_p^{\times}$ is connected. We note that this cap is necessary to guarentee an almost linear asymptotic runtime at the expense of a possibly inconclusive response by the algorithm; however, the lack of growth in the number of required checks over the large range $p < 110,000,000$ (as shown in Figure 5) provides evidence that this check does not reduce the algorithm's power. We use the following result.

Theorem 3 ([10, theorem 315])**.** *For every $\epsilon > 0$, the divisor function $\mathcal{D}(n) \in o(n^{\epsilon})$.*

Therefore, our algorithm runs in $o(p^{1+\epsilon})$ time, in the worst case. The runtimes for all the primes we tested are shown in Figure 6. As emphasized in the figures, the almost linear bound on the runtime is a worst-case upper bound, and our algorithm performs well below this upper limit for a large number of primes.

The actual time to run Algorithm 3 depends heavily on both the endgame and middle game breakpoints. A plot of a the breakpoints for our random primes beneath 110 million, calculated using equation (5) and Algorithm 1, is shown in Figure 7. The first prime for which Algorithm 1 returns a middle game breakpoint is $p = 1,328,247$.

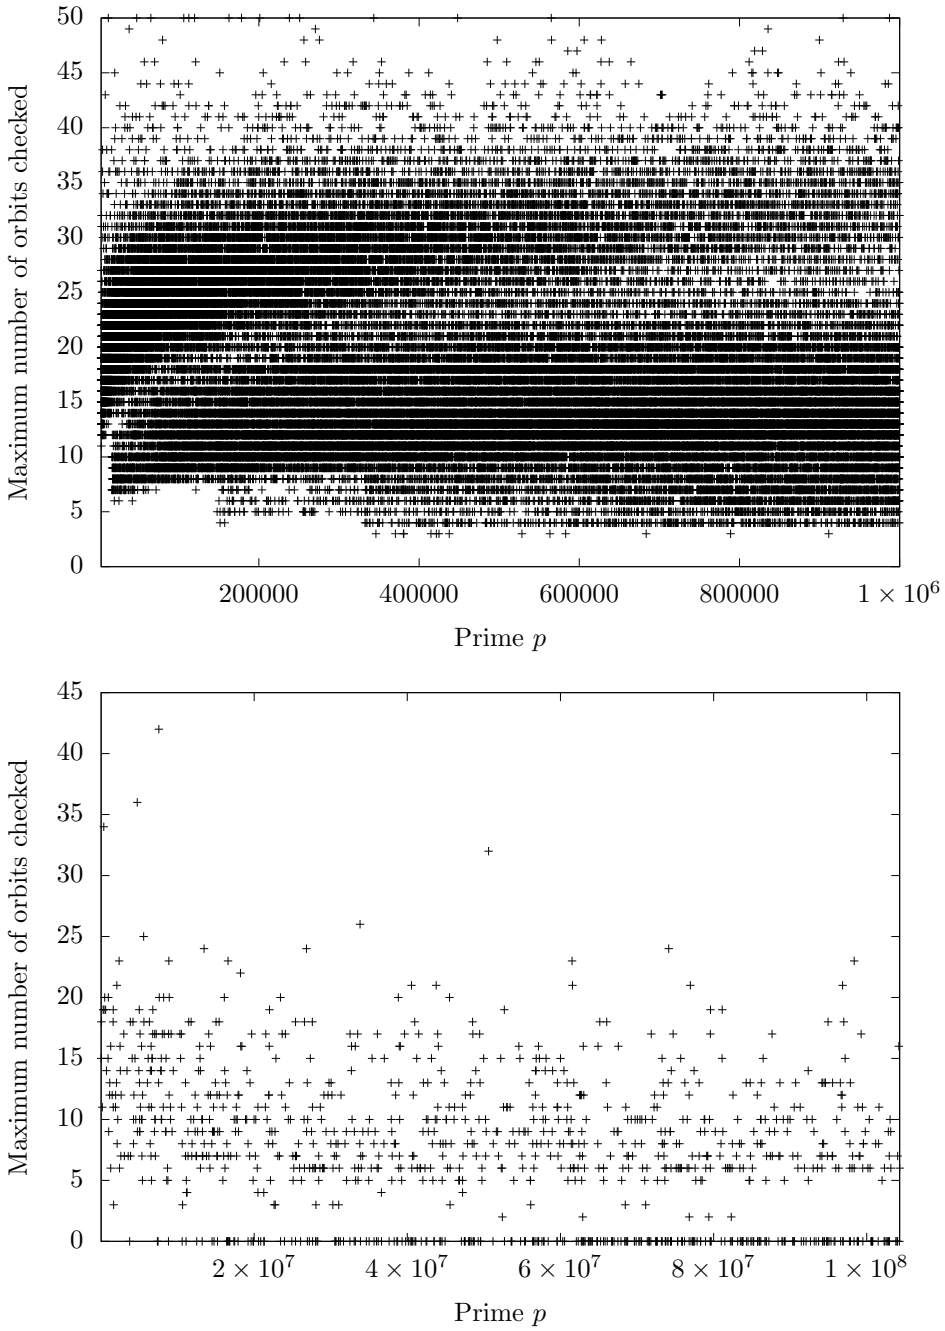

FIGURE 5. Maximum number of orbits checked during any iteration of Step 4(b) for all primes less than $1,000,000$ (top) and random primes less than $110,000,000$ (bottom).

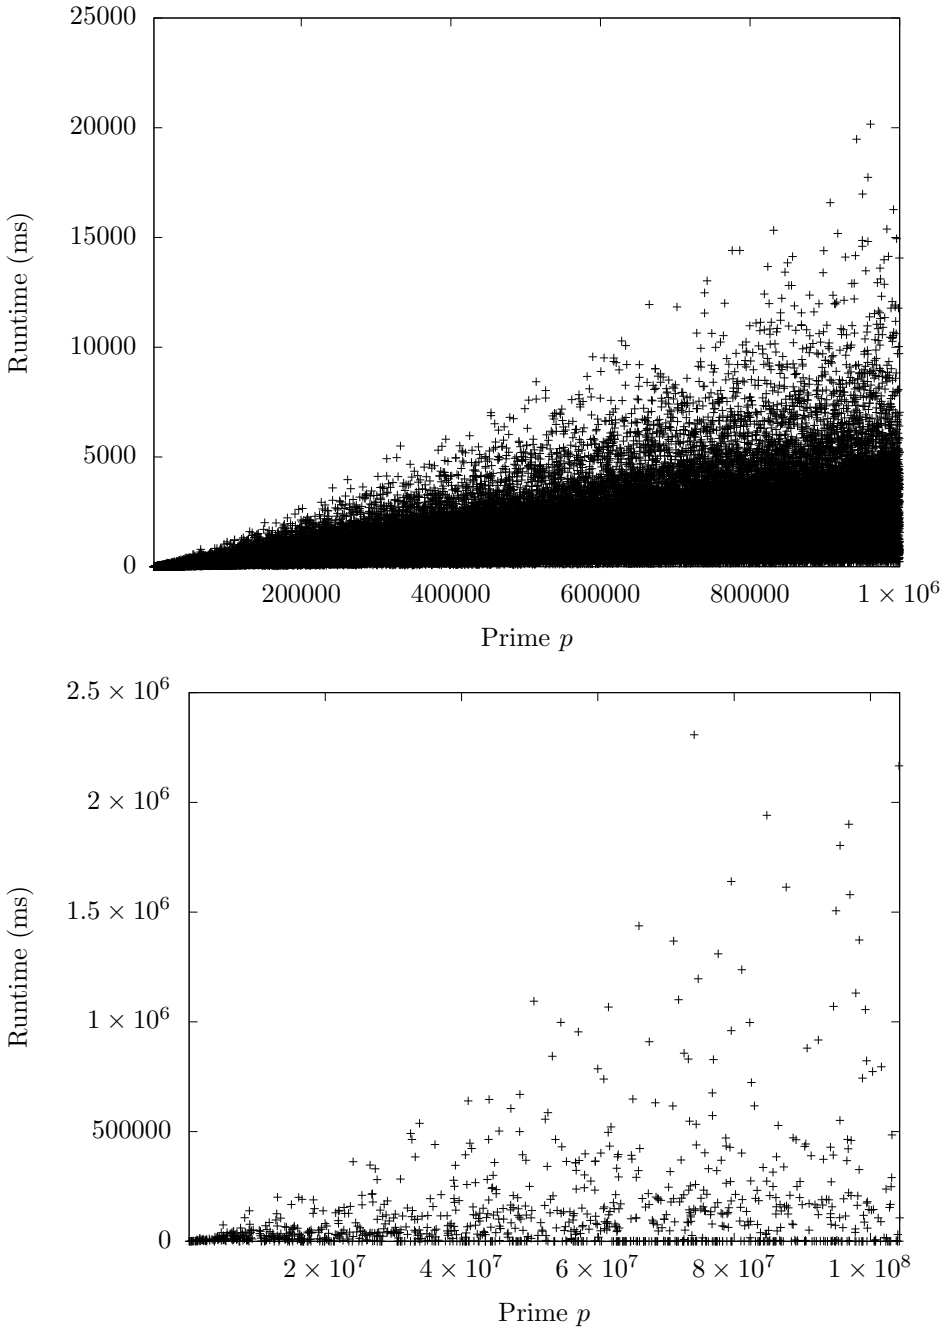

FIGURE 6. Runtimes for Algorithm 3 on all primes less than $1,000,000$ (top) and our random sample of primes less than $110,000,000$ (bottom). Measured on an 11th Gen Intel Core i5-11320H processor.

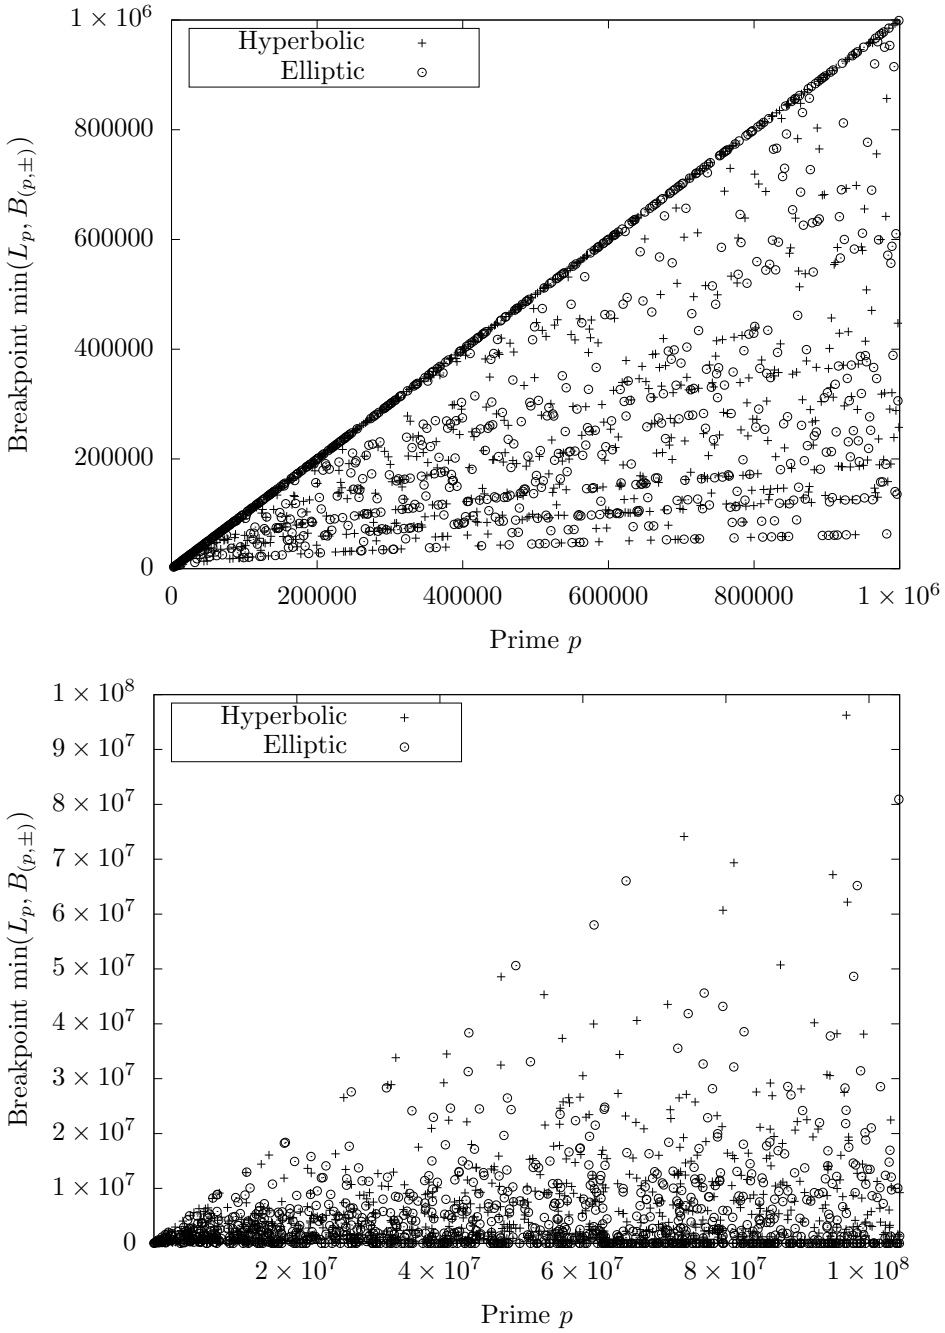

FIGURE 7. Breakpoints for all primes less than $1,000,000$ (top) and random primes less than $110,000,000$ (bottom).

REFERENCES

[1] Arthur Baragar. "The Markoff equation and equations of Hurwitz". Brown University, 1991.
[2] Elisa Bellah, Siran Chen, Elena Fuchs, and Lynnelle Ye. "Bounding Lifts of Markoff Triples mod p". 2023. arXiv: 2311.11468 [math.NT].

[3] Jean Bourgain, Alexander Gamburd, and Peter Sarnak. "Markoff Surfaces and Strong Approximation: 1". 2016. arXiv: 1607.01530 [math.NT].

[4] William Y. Chen. "Nonabelian level structures, Nielsen equivalence, and Markoff triples". In: *Annals of Mathematics* 199.1 (2024). ISSN: 0003-486X.

[5] Pietro Corvaja and Umberto Zannier. "Greatest common divisors of $u - 1$, $v - 1$ in positive characteristic and rational points on curves over finite fields". In: *Journal of the European Mathematical Society* 15.5 (2013), pp. 1927–1942. DOI: 10.4171/JEMS/409.

[6] Matthew De Courcy-Ireland and Seugjae Lee. "Experiments with the Markoff Surface". In: *Experimental mathematics* 31.3 (2022), pp. 814–829.

[7] Jillian Eddy, Elena Fuchs, Matthew Litman, Daniel Martin, Nico Tripeny, and Devin Vanyo. "Connectivity of Markoff Mod-P Graphs and Maximal Divisors". 2023. arXiv: 2308.07579 [math.NT].

[8] E. Fuchs, K. Lauter, M. Litman, and A. Tran. "A cryptographic hash function from Markoff triples". In: *Mathematical Cryptology* (2022).

[9] Amit Ghosh and Peter Sarnak. "Integral points on Markoff type cubic surfaces". In: *Inventiones mathematicae* 229.2 (2022), pp. 689–749. ISSN: 0020-9910.

[10] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, 1975.

[11] Steve Klabnik and Carol Nichols. *The Rust Programming Language*. 2022. URL: https://doc.rust-lang.org/book (visited on 11/29/2023).

[12] Sergei V. Konyagin, Sergey V. Makarychev, Igor E. Shparlinski, and Ilya V. Vyugin. "On the structure of graphs of Markoff triples". In: *The Quarterly Journal of Mathematics* 71 (2020).

[13] Frank Lübeck. "Standard generators of finite fields and their cyclic subgroups". In: *Journal of Symbolic Computation* 117.51-67 (2023).

[14] A. Markoff. "Sur les formes quadratiques binaires indéfinies". fe. In: *Mathematische annalen* 15.3-4 (1879).

[15] Henry S. Warner Jr. *Hacker's Delight, Second Edition*. Addison-wesley Professional, 2012.