THE FROBENIUS NUMBER OF A SHIFTED SQUARE SEQUENCE: A COMPLETE FORMULA, THE GENUS, AND A SHARP BOUND AT $a \equiv 7 \pmod 8$

KOREA SUPERINTELLIGENCE LABS

Abstract. For an integer $a \geq 2$ let S_a be the numerical semigroup generated by the shifted squares $a,\ a+1^2,\ a+2^2,\ a+3^2, \ldots$, and let $\mathrm{F}(a)$ be its Frobenius number. Liu and Xin determined $\mathrm{F}(a)$ under a hypothesis — the existence of an $r < a$ with $\iota(r) = 4$, $\iota(a+r) \geq 3$, $\iota(2a+r) \geq 2$, where $\iota(n)$ is the least number of positive squares summing to n — and conjectured that the hypothesis holds for every $a > 30$; their table records three values of a that neither of their two theorems reaches. We give a formula for $\mathrm{F}(a)$ that is valid for *every* $a \geq 2$ with no hypothesis and no case left over: writing $n = ta + r$ with $0 \leq r < a$, membership in S_a is decided by at most four evaluations of ι, the gap set lies in $[0, 4a)$, and $\mathrm{F}(a)$ is read off the highest non-empty of four rungs, the top two of which are the theorems of Liu and Xin. The same description gives the genus $|\mathbb{N} \setminus S_a|$ as an explicit count of at most $4a$ points, a quantity neither source computes. We also prove $\mathrm{F}(a) \geq 4a - 8$ for $a \equiv 7 \pmod 8$, $a \geq 15$, sharpening the constant 48 of Song for that class by a factor of six, with equality already at $a = 15$; and we prove Liu and Xin's Theorem 2.8 and their Conjecture 2.10, the latter by a single congruence class rather than a case analysis, which in particular covers $a = 73$ and $a = 77$. All statements are machine-checked in Lean 4.

1. Introduction

A *numerical semigroup* is a submonoid $S \subseteq \mathbb{N}$ with finite complement; its *Frobenius number* $\mathrm{F}(S)$ is the largest integer not in S, and its *genus* is $|\mathbb{N} \setminus S|$. Determining $\mathrm{F}(S)$ from a generating set is the Frobenius problem [6]. For a sequence of generators given by a polynomial shift the problem has a uniform flavour, and the case of squares is the one this paper is about. For an integer $a \geq 1$ put

$$A(a) \;=\; \big(a,\ a+1^2,\ a+2^2,\ a+3^2,\ \ldots\big), \qquad S_a \;=\; \big\langle\, a+k^2 \;:\; k \geq 0 \,\big\rangle \;\subseteq\; \mathbb{N}.$$

Since a and $a+1$ are coprime generators, $\mathbb{N} \setminus S_a$ is finite for every $a \geq 2$; we write $\mathrm{F}(a)$ for the Frobenius number of S_a and $\mathrm{g}(a)$ for its genus. At $a = 1$ the generator $1 + 0^2 = 1$ already gives $S_1 = \mathbb{N}$, so there is no Frobenius number and $a \geq 2$ is assumed throughout.

Throughout, $\iota(n)$ denotes the least number of *positive* squares whose sum is n (so $\iota(0) = 0$, and $\iota(n) \leq 4$ for all n by Lagrange's four-square theorem); it is sequence A002828 of [5].

Shifted square sequences entered the Frobenius literature with Einstein, Lichtblau, Strzebonski and Wagon [2], who treated the *finite* sequences $(a, a+1^2, \ldots, a+M^2)$: their Theorem 5 gives, for $M \leq 7$ and a beyond an explicit bound, a formula $\frac{1}{M^2}(a^2 + c_{M,j}a) - d_{M,j}$ with $j = a \bmod M^2$, and their Problem 4 of Section 20.2 asks for one valid for every M. The infinite sequence $A(a)$ does not appear there.

Liu and Xin [3] answer that problem for the finite sequences in their Section 4; their Section 2, which is what concerns us, treats the infinite sequence $A(a)$. We cite [3] throughout by the numbering of its arXiv version v4. For the squares their Theorem 2.8 states: if some r with $1 \leq r \leq a-1$ satisfies

$$\iota(r) = 4, \qquad \iota(a+r) \geq 3, \qquad \iota(2a+r) \geq 2, \tag{1}$$

then $\mathrm{F}(a) = 3a + \max\{r\}$, the maximum over all such r. Their Theorem 2.9 covers a further list of a at which (1) has no solution, and their Conjecture 2.10 asserts that (1) does have a solution whenever $a > 30$. Their Table 1, in the appendix, tabulates $\mathrm{F}(a)$ for $2 \leq a \leq 42$ together with a

column recording which of the two theorems applies; at $a = 2, 3, 6$ that column reads `no`, and the paper gives no formula at those three values.

Song [8] takes up the same sequence. The first version of that preprint opens by calling the computation of $\mathrm{F}(a)$ "still an open problem", citing [2] and [3], and solves the case $a = (2k)^2$. The current version v3, retitled *The Frobenius problem for shifted square sequences*, proves Conjecture 2.10 by a case analysis on $a \bmod 8$: the classes $a \equiv 0, 4$ receive an exact formula (its Theorem 2.3), the remaining six classes a uniform bound $\mathrm{F}(a) \geq 4a - C(c)$ for large $a \equiv c$ (its Theorem 2.4), and the finitely many $a > 30$ below those thresholds a table. The six constants recorded there are

$$C(1) = 66, \quad C(2) = 43, \quad C(3) = 44, \quad C(5) = 70, \quad C(6) = 47, \quad C(7) = 48.$$

So the literature leaves four things standing. There is no formula for $\mathrm{F}(a)$ at $a = 2, 3, 6$, and more generally none at any a where (1) fails; the genus of S_a is computed nowhere, although Liu and Xin, with Ye and Yin, compute the genus of a neighbouring family [4]; the size of the deficiency $4a - \mathrm{F}(a)$ is bounded only class by class and, in several classes, not tightly; and Conjecture 2.10, while proved in [8], is proved by an argument whose finite part appears not to reach every a (Remark 5.6).

Results. Everything below rests on one identity, proved in Section 3 directly from the definition of S_a: a sum of j generators is ja plus a sum of j squares, zeros allowed, because $a = a + 0^2$ is itself a generator, so

$$n \in S_a \iff \exists\, j \geq 0 : \; ja \leq n \;\text{ and }\; \iota(n - ja) \leq j.$$

This is the membership form of the Apéry-set lemma of [3] (their Lemma 2.7); we prove it here from scratch. Two consequences do all the work: everything $\geq 4a$ lies in S_a, so the gap set is contained in $[0, 4a)$; and writing $n = ta + r$ with $0 \leq r < a$, the condition $ja \leq n$ is $j \leq t$, so at most four evaluations of ι decide whether n is a gap.

A complete formula (Section 4). For $0 \leq t \leq 3$ let $R_t(a)$ be the largest $r < a$, if any, satisfying the t-th of the four conditions listed in (2). Then for every $a \geq 2$

$$\mathrm{F}(a) \;=\; t^* a + R_{t^*}(a), \qquad t^* = \max\{t \leq 3 : R_t(a) \text{ exists}\},$$

with $R_0(a) = a - 1$ always, so the right-hand side is defined for every $a \geq 2$ (Theorem 4.1). Rung 3 is Liu and Xin's Theorem 2.8 and rung 2 is their Theorem 2.9; rungs 1 and 0 appear in neither source, and they are exactly what is needed at $a = 2, 3, 6$. The same section records $a - 1 \leq \mathrm{F}(a) \leq 4a - 1$ and a decision procedure that computes $\mathrm{F}(a)$ from $\Theta(a)$ evaluations of ι.

The genus (Section 7). $\mathrm{g}(a)$ is the number of $n < 4a$ passing the same four-evaluation test (Theorem 7.1); for instance $\mathrm{g}(31) = 74$ and $\mathrm{g}(250) = 651$.

A sharp bound at $a \equiv 7 \pmod 8$ (Section 6). For $a \equiv 7 \pmod 8$ with $a \geq 15$ one has $\mathrm{F}(a) \geq 4a - 8$ (Theorem 6.3), a sixfold sharpening of the constant $C(7) = 48$ of [8], with a two-line proof and with equality already at $a = 15$, where $\mathrm{F}(15) = 52$.

The two results of [3] (Section 5). We prove their Theorem 2.8 (Theorem 5.1) and their Conjecture 2.10 (Theorem 5.4). The conjecture is also proved in [8]; the proof here is a single congruence class instead of six cases and leaves no finite range to check by hand, and in particular it covers $a = 73$ and $a = 77$, which the case analysis of [8] appears to miss.

Every theorem, proposition, corollary and lemma below has been formally verified in Lean 4 against *Mathlib*; Section 9 says exactly what rests on exhaustive computation. The remarks are commentary; the two numerical measurements that were computed outside the formal development appear only there, and are labelled as such where they stand.

2. Preliminaries

Definition 2.1. For $a \in \mathbb{N}$ let $S_a \subseteq \mathbb{N}$ be the additive submonoid generated by $\{a + k^2 : k \in \mathbb{N}\}$. A natural number f is *the Frobenius number* of a submonoid $T \subseteq \mathbb{N}$ if $f \notin T$ and every $n > f$ lies in T; such an f is unique when it exists, and we write $\mathrm{F}(a)$ for the Frobenius number of S_a.

Uniqueness is immediate from the definition, and it is what licenses the phrase "the" Frobenius number in every statement below: two candidates $f < g$ would put g both inside and outside T. The existence of $\mathrm{F}(a)$ for $a \geq 2$ is not assumed anywhere; it comes out of Theorem 4.1.

Definition 2.2. For $n \in \mathbb{N}$,

$$\iota(n) \;=\; \min\bigl\{\, m : n = x_1^2 + \cdots + x_m^2 \text{ with all } x_i \geq 1 \,\bigr\},$$

the least number of positive squares summing to n.

The infimum is over a non-empty set by Lagrange's four-square theorem, which also gives $\iota(n) \leq 4$ for every n. It is convenient to allow zero summands: write $\mathrm{Sq}_{\leq m}(n)$ for the assertion that n is a sum of at most m squares, zeros permitted. Deleting the zero entries of such a representation only shortens it, so $\iota(n) \leq m$ if and only if $\mathrm{Sq}_{\leq m}(n)$, and the two least-lengths agree.

Proposition 2.3. *$\mathrm{Sq}_{\leq m}(n)$ is decidable, uniformly in m and n, by the recursion that peels one square $k^2 \leq n$ with $0 \leq k \leq \lfloor\sqrt{n}\rfloor$ at a time; consequently ι is computable, and $\iota(n)$ is determined by one positive and at most four negative instances of $\mathrm{Sq}_{\leq m}$.*

The last clause is what makes ι usable inside an exhaustive check: to certify $\iota(n) = m$ one exhibits a representation of length m and refutes lengths $0, \ldots, m-1$, and by Lagrange the case $m = 4$ needs no positive instance at all.

Remark 2.4. Liu and Xin write $g(A)$ for the Frobenius number of the semigroup generated by a sequence A. We reserve g for the genus, as is usual for numerical semigroups, and write F for the Frobenius number.

3. The membership criterion

Theorem 3.1. *For all $a, n \in \mathbb{N}$,*

$$n \in S_a \iff \exists\, j \in \mathbb{N} : \quad ja \leq n \;\text{ and }\; \iota(n - ja) \leq j.$$

Equivalently, $n \notin S_a$ if and only if $\iota(n - ja) > j$ for every j with $ja \leq n$.

Proof sketch. For $\Rightarrow$, induct over the closure defining S_a: a single generator $a + k^2$ is the case $j = 1$, the zero element is $j = 0$, and if n_1 uses j_1 generators and n_2 uses j_2 then $n_1 + n_2$ uses $j_1 + j_2$, the square parts concatenating. For $\Leftarrow$, pad: a representation of $n - ja$ by $\ell \leq j$ squares becomes a representation of n as a sum of j generators by appending $j - \ell$ copies of the generator $a = a + 0^2$. The passage between ι and $\mathrm{Sq}_{\leq j}$ is the equivalence noted in Section 2; the fact that the infimum defining ι is attained uses Lagrange. □

Corollary 3.2. *If $4a \leq n$ then $n \in S_a$. Hence every gap of S_a is smaller than $4a$.*

Proof. Take $j = 4$ in Theorem 3.1 and use $\iota(n - 4a) \leq 4$. □

Write $n = ta + r$ with $0 \leq r < a$. Then $ja \leq n$ holds exactly for $j \leq t$, and by Corollary 3.2 only $t \leq 3$ matters. So membership is decided by at most four evaluations of ι, and the gap test has four instances, one per value of t:

(2)

t	$ta + r$ is a gap of $S_a \iff$
3	$\iota(r) = 4 \;\wedge\; \iota(a + r) \geq 3 \;\wedge\; \iota(2a + r) \geq 2$
2	$\iota(r) \geq 3 \;\wedge\; \iota(a + r) \geq 2$
1	$\iota(r) \geq 2$
0	$r \geq 1$

(In rung $t = 3$ the condition $\iota(3a + r) \geq 1$ coming from $j = 0$ is automatic, and $\iota(r) > 3$ is $\iota(r) = 4$ because $\iota \leq 4$; rungs $t = 2, 1, 0$ simplify in the same way.)

Proposition 3.3. *Let $a \geq 2$. Then the largest $n < 4a$ satisfying the rung of* (2) *indexed by $t = \lfloor n/a \rfloor$ exists, and it is the Frobenius number* $\mathrm{F}(a)$*. In particular* $\mathrm{F}(a)$ *exists, and it is computed by* $\Theta(a)$ *evaluations of* ι*.*

Proof sketch. The test in (2) decides non-membership below $4a$, by Theorem 3.1 and Corollary 3.2. The search is non-empty because $a-1$ passes it: the only j with $ja \leq a-1$ is $j = 0$, and $\iota(a-1) \geq 1$. The largest n passing the test is a gap, and every larger n is either $\geq 4a$, hence in S_a by Corollary 3.2, or fails the test, hence lies in S_a as well. □

4. The complete formula

For $0 \leq t \leq 3$ let

$$R_t(a) \;=\; \max\{\, r < a \;:\; r \text{ satisfies rung } t \text{ of } (2)\,\},$$

undefined when the set is empty. Rung 0 is satisfied by $r = a - 1$ for every $a \geq 2$, so $R_0(a) = a - 1$ always exists. Note also that $ta + R_t(a) < (t+1)a$ whenever $R_t(a)$ exists, so the candidate values increase with t: taking the highest non-empty rung is the same as taking the largest value that any non-empty rung produces.

Theorem 4.1. *For every $a \geq 2$, the Frobenius number of S_a exists and equals*

$$\mathrm{F}(a) \;=\; \begin{cases} 3a + R_3(a) & \text{if } R_3(a) \text{ exists,} \\ 2a + R_2(a) & \text{else, if } R_2(a) \text{ exists,} \\ a + R_1(a) & \text{else, if } R_1(a) \text{ exists,} \\ a - 1 & \text{otherwise.} \end{cases}$$

There is no hypothesis on a and no case left over.

Proof sketch. Each of the four cases is a separate statement, proved from Theorem 3.1 in the same shape. Take the second case for illustration. That $2a + R_2(a)$ is a gap is (2) at $t = 2$. That every $n > 2a + R_2(a)$ lies in S_a splits: if $n \geq 3a$ then, since rung 3 is empty, writing $n = 3a + s$ with $s < a$ (or invoking Corollary 3.2 when $n \geq 4a$) and negating rung 3 of (2) produces a $j \in \{1, 2, 3\}$ with $\iota(n - ja) \leq j$; if $2a \leq n < 3a$ then $n = 2a + s$ with $s > R_2(a)$, and maximality of $R_2(a)$ negates rung 2 at s, producing $j \in \{1, 2\}$. The first case is the same argument one rung up, the third and fourth one and two rungs down. Assembling: if $R_3(a)$ exists the first case applies; otherwise rung 3 is empty and the second applies if $R_2(a)$ exists; and so on, with the fourth case always available. □

The top two cases are known. The first is Theorem 5.1 below, which is Theorem 2.8 of [3]. The second is Theorem 2.9 of [3]: their statement lists three conditions on r, and the proof of Theorem 4.1 shows that once rung 3 is known to be empty those three reduce to the pair $\iota(r) \geq 3 \wedge \iota(a + r) \geq 2$ of rung 2, since $\iota(r) = 4 \wedge \iota(a + r) \geq 3$ would then force $\iota(2a + r) = 1$. The bottom two cases are in neither source. They are not decorative: rungs 3 and 2 are both empty at $a = 2, 3, 6$ — precisely the three values marked `no` in Table 1 of [3] — where rung 1 gives $\mathrm{F}(3) = 3 + 2 = 5$ and $\mathrm{F}(6) = 6 + 5 = 11$, and rung 0 gives $\mathrm{F}(2) = 1$.

Corollary 4.2. *For every $a \geq 2$, $a - 1 \leq \mathrm{F}(a) \leq 4a - 1$.*

Proof. The upper bound is Corollary 3.2. For the lower bound, $a - 1$ is a gap, as in the proof of Proposition 3.3. □

Remark 4.3. Theorem 4.1 replaces, as a way of computing $\mathrm{F}(a)$, the Apéry-set certificates of Section 8: those cost $\Theta(a^2\sqrt{a})$ per value and ship a certificate of length $\Theta(a)$, whereas the ladder needs $\Theta(a)$ evaluations of ι — in practice far fewer, since $R_3(a)$ is usually close to $a-1$ — and no certificate at all.

5. Theorem 2.8 and Conjecture 2.10 of Liu and Xin

Theorem 5.1. ([3, Theorem 2.8]) *Let $a \geq 2$ and let $r < a$ satisfy $\iota(r) = 4$, $\iota(a+r) \geq 3$, $\iota(2a+r) \geq 2$. If no s with $r < s < a$ satisfies the same three conditions, then $\mathrm{F}(a) = 3a + r$.*

Proof sketch. This is the top case of Theorem 4.1, and it is short enough to give in full. That $3a+r$ is a gap: by Theorem 3.1 one must refute $\iota(3a + r - ja) \leq j$ for $j = 0, 1, 2, 3$, and these four are $\iota(3a+r) \geq 1$, $\iota(2a+r) \geq 2$, $\iota(a+r) \geq 3$, $\iota(r) \geq 4$, which are the hypotheses. That every $n > 3a+r$ is in S_a: if $n \geq 4a$ use Corollary 3.2; otherwise $n = 3a + s$ with $r < s < a$, maximality says one of the three conditions fails at s, and each failure names a $j \in \{1, 2, 3\}$ with $\iota(n - ja) \leq j$. □

Liu and Xin conjecture that the hypothesis of Theorem 5.1 is never vacuous beyond $a = 30$:

Conjecture 5.2 ([3, Conjecture 2.10])**.** For every $a > 30$ there exists r with $1 \leq r \leq a-1$ satisfying (1).

We prove it. The three conditions of (1) are lower bounds on ι, and lower bounds need only the *easy* directions of the classical two- and three-square theorems — the congruence obstructions. Neither the three-square theorem of Legendre nor the two-square theorem of Fermat is needed, and neither is used.

Lemma 5.3. (1) *If $n \equiv 7 \pmod 8$ then $\iota(n) = 4$.*

(2) *If $n \equiv 3 \pmod 4$, or $n \equiv 6 \pmod 8$, or $n \equiv 12 \pmod{16}$, or $n \equiv 24 \pmod{32}$, or $n \equiv 3, 6 \pmod 9$, then $\iota(n) \geq 3$.*

(3) *If $n \bmod 360$ is not one of the 36 squares modulo 360, then n is not a perfect square, and $\iota(n) \geq 2$.*

Proof sketch. (1) Squares are $0, 1, 4$ modulo 8, and no three of those sum to 7 modulo 8; this is an exhaustive check of the $8^3 = 512$ triples of residues. With $\iota(n) \leq 4$ from Lagrange this gives $\iota(n) = 4$. This is the easy half of Legendre's three-square theorem.

(2) The last two cases: if $3 \mid x^2 + y^2$ then $3 \mid x$ and $3 \mid y$, so $9 \mid x^2 + y^2$; as an exhaustive check this is the $9^2 = 81$ pairs of residues modulo 9. The first four cases are the 2-adic descent. If $x^2 + y^2$ is even then $x \equiv y \pmod 2$, and writing $x = y + 2d$ gives $(y + 2d)^2 + y^2 = 2\big((y + d)^2 + d^2\big)$, so a sum of two squares halves to a sum of two squares. Iterating, a representation of $2^e m$ with m odd descends to one of m, and no $m \equiv 3 \pmod 4$ is a sum of two squares (an exhaustive check of the 16 pairs of residues modulo 4). The four congruences listed are the cases $e = 0, 1, 2, 3$ of "odd part $\equiv 3 \pmod 4$".

(3) The 36 square residues modulo 360 are computed once, by an exhaustive check of the 360 residues; a number congruent to none of them is not a square, and a positive non-square needs at least two squares. □

Theorem 5.4. *Conjecture 5.2 holds: for every $a \geq 31$ there is an r with $1 \leq r < a$ satisfying (1). Consequently, for every $a \geq 31$,*

$$\mathrm{F}(a) \;=\; 3a + \max\{\, r < a \;:\; \iota(r) = 4,\ \iota(a + r) \geq 3,\ \iota(2a + r) \geq 2 \,\}.$$

Proof sketch. Call r *certified* for a if $r \equiv 7 \pmod 8$, $a + r$ matches one of the six congruences of Lemma 5.3(2), and $2a + r$ fails the square-residue test of Lemma 5.3(3). By Lemma 5.3 a certified r satisfies (1), so Theorem 5.1 applies to the largest such r.

Being certified depends on a only through $a \bmod 1440$: the moduli involved are $4, 8, 16, 32$ and 9, all dividing 1440, and $360 \mid 2 \cdot 1440$, which handles the doubled argument $2a + r$. So the whole conjecture reduces to a finite check. *The finite search that was run is this: for each of the 1440 residues $b < 1440$ and each of the eight candidates $r \in \{7, 15, 23, 31, 39, 47, 55, 63\}$ in turn, take the first certified candidate; verify that one exists, that it is < 64, and — when $31 \leq b \leq 63$ — that it is $< b$.* All 1440 residues pass. Given $a \geq 31$, let r be the first certified candidate for $a \bmod 1440$; it is certified for a and $r < 64$, so $r < a$ when $a \geq 64$, while for $31 \leq a \leq 63$ the third clause of the check gives $r < a$ directly. No finite range is left over. □

Remark 5.5. The threshold is exactly where the construction stops. At $a = 30$ the first certified candidate is $r = 39$, which is not < 30; and $a = 30$ is in fact the largest a at which (1) has no solution at all (Section 8). The certification is not vacuous either: at $a = 30$ the candidate $r = 7$ is rejected because $a + r = 37 = 1^2 + 6^2$.

Remark 5.6. Conjecture 5.2 is also proved in [8], by a case analysis on $a \bmod 8$ using the same three ingredients, plus an enumeration of $x^2 - y^2 \in \{24, 48\}$ in the cases $a \equiv 1, 5 \pmod 8$. In that argument the uniform bounds for those two classes start at $a \geq 81$ and $a \geq 85$ respectively, while its Table 1, which is to cover the remaining $a > 30$, lists $a = 33, 41, 49, 57, 65$ for $a \equiv 1$ and $a = 37, 45, 53, 61, 69$ for $a \equiv 5$; the values $a = 73$ and $a = 77$ fall in neither, so the finite part of that proof appears not to reach them. Both are covered by Theorem 5.4, which needs no finite table.

6. Closed forms in three residue classes

Theorem 5.1 turns a single certified r into a lower bound for $\mathrm{F}(a)$, since $\mathrm{F}(a) = 3a + R_3(a) \geq 3a + r$; combined with Corollary 4.2 this gives two-sided bounds, and when the certified r is $a - 1$ it gives an exact value.

Theorem 6.1. *If* $8 \mid a$ *and* $a \geq 8$ *then* $\mathrm{F}(a) = 4a - 1$ *exactly.*

Proof. Take $r = a - 1$. Then r, $a + r = 2a - 1$ and $2a + r = 3a - 1$ are all $\equiv 7 \pmod 8$, so Lemma 5.3(1) gives $\iota(r) = \iota(a + r) = \iota(2a + r) = 4$ and (1) holds. Since $r = a - 1$ is the largest index available, $R_3(a) = a - 1$ and Theorem 5.1 gives $\mathrm{F}(a) = 3a + (a - 1) = 4a - 1$. □

Theorem 6.2. *If* $a \equiv 4 \pmod 8$ *and* $a \geq 12$ *then* $4a - 5 \leq \mathrm{F}(a) \leq 4a - 1$.

Proof. Take $r = a - 5$. Then $r \equiv 7 \pmod 8$, so $\iota(r) = 4$; $a + r = 2a - 5 \equiv 3 \pmod 4$, so $\iota(a + r) \geq 3$ by Lemma 5.3(2); and $2a + r = 3a - 5 \equiv 7 \pmod 8$ is not a square modulo 8, so $\iota(2a + r) \geq 2$. Hence $R_3(a) \geq a - 5$ and $\mathrm{F}(a) \geq 3a + (a - 5)$. The upper bound is Corollary 4.2. □

Theorem 6.1 is the first case of Theorem 2.3 of [8], and Theorem 6.2 is the lower half of the second case of the lemma preceding it (Lemma 2.2 there, which also asserts $R_3(a) \leq a - 4$ in that class). That upper half is out of reach of the congruence obstructions used here: $R_3(a) \leq a - 4$ needs $\iota(a - 1) \leq 3$ for $a - 1 \equiv 3 \pmod 8$, which is the hard direction of Legendre's three-square theorem. The next result is in neither source: for $a \equiv 7 \pmod 8$ the constant recorded in [8] is $C(7) = 48$, and the true constant is 8.

Theorem 6.3. *If* $a \equiv 7 \pmod 8$ *and* $a \geq 15$ *then* $4a - 8 \leq \mathrm{F}(a) \leq 4a - 1$.

Proof. Take $r = a - 8$. Then $r \equiv 7 \pmod 8$, so $\iota(r) = 4$ by Lemma 5.3(1). Next $a + r = 2a - 8 = 2(a - 4)$ with $a - 4 \equiv 3 \pmod 8$, so the odd part of $a + r$ is $\equiv 3 \pmod 4$; concretely $a + r \equiv 6 \pmod 8$, and Lemma 5.3(2) gives $\iota(a + r) \geq 3$. Finally $2a + r = 3a - 8 \equiv 5 \pmod 8$ is not a square modulo 8, so $\iota(2a + r) \geq 2$. Hence $R_3(a) \geq a - 8$ and $\mathrm{F}(a) \geq 3a + (a - 8)$; the upper bound is Corollary 4.2. □

Remark 6.4. The constant 8 cannot be lowered: $\mathrm{F}(15) = 52 = 4 \cdot 15 - 8$, a value that is machine-checked independently in Section 8, so the bound is attained. Outside the formal development we computed $4a - \mathrm{F}(a)$ for all $a \leq 2 \cdot 10^5$, evaluating ι from the classical characterisation of Fermat, Legendre and Lagrange over a smallest-prime-factor sieve: over the class $a \equiv 7 \pmod 8$ its maximum is exactly 8, so Theorem 6.3 is sharp on the whole of that range, and over all $a \leq 2 \cdot 10^5$ its maximum is 47, attained first at $a = 6982$ with $a \equiv 6 \pmod 8$. The per-class maxima are, for $a \bmod 8 = 0, 1, \ldots, 7$: $1, 34, 43, 31, 5, 30, 47, 8$. These are measurements, not theorems; the only one of them inside the formal development is the value at $a = 15$.

7. The genus

Neither [3] nor [8] computes the genus of S_a, although the genus of a neighbouring family of semigroups is computed in [4]. The description of the gap set in Section 3 gives it immediately.

Theorem 7.1. *For every $a \geq 1$,*

$$\mathrm{g}(a) \;=\; \bigl|\mathbb{N} \setminus S_a\bigr| \;=\; \#\bigl\{\, n < 4a \;:\; n \text{ satisfies the rung of (2) indexed by } t = \lfloor n/a \rfloor \,\bigr\},$$

a count of at most $4a$ points, each decided by at most four evaluations of ι.

Proof. By Corollary 3.2 the gap set is contained in $\{0, \ldots, 4a-1\}$, and on that range (2) decides membership (Theorem 3.1). So the gap set is the displayed finite set, and its cardinality as a set of naturals is the cardinality of that finite set. □

Theorem 7.2. $\mathrm{g}(2) = 1$, $\mathrm{g}(3) = 3$, $\mathrm{g}(6) = 8$, $\mathrm{g}(31) = 74$ *and* $\mathrm{g}(250) = 651$.

Proof. Evaluation of the count of Theorem 7.1 at the five values of a. □

Remark 7.3. Outside the formal development, the same count gives

$$\mathrm{g}(a) \;=\; 1,\, 3,\, 6,\, 7,\, 8,\, 12,\, 16,\, 16,\, 17,\, 20,\, 24,\, 24,\, 28,\, 32,\, 36,\, 36,\, 37,\, 38,\, 44,\, 43 \qquad (a = 2, \ldots, 21),$$

of which the entries at $a = 2, 3, 6$ are the machine-checked ones of Theorem 7.2. This sequence returns no match in [5], nor does the deficiency sequence $4a - \mathrm{F}(a)$.

8. Values, cross-checks and controls

The formula of Section 4 is not the only route to $\mathrm{F}(a)$ available here, and the two routes were run against each other. The second route is the classical one. For a numerical semigroup S containing a, the Apéry set is $\mathrm{Ap}(S, a) = \{w(c) : c < a\}$ with $w(c) = \min\{s \in S : s \equiv c \pmod a\}$, and Brauer–Shockley gives $\mathrm{F}(S) = \max \mathrm{Ap}(S, a) - a$ [1, 7].

Proposition 8.1. *Let $a \geq 2$ and $K \in \mathbb{N}$ with $4a \leq K^2$, and let $w : \mathbb{N} \to \mathbb{N}$ satisfy $w(0) = 0$, the residue condition $w(c) \equiv c \pmod a$ for $c < a$ with each $w(c)$ exhibited as an explicit sum of generators, and the Bellman inequalities $w\bigl((c + k^2) \bmod a\bigr) \leq w(c) + a + k^2$ for all $c < a$ and all $k \leq K$. Then $\mathrm{F}(a) = \max_{c<a} w(c) - a$. The cap $K = \lfloor \sqrt{6a} \rfloor + 1$ satisfies $4a \leq K^2$.*

The only step here that is not textbook is why a finite cap on the shift index suffices, and the hypothesis $4a \leq K^2$ is exactly what makes it work: a generator $a + k^2 < 5a$ forces $k^2 < 4a \leq K^2$, so below $5a$ the semigroup generated by the capped set $\{a + k^2 : k \leq K\}$ agrees with S_a; and Lagrange puts $4a + c$ in the capped semigroup for every $c < a$, so no class minimum exceeds $5a$.

Theorem 8.2. *For every a with $2 \leq a \leq 250$ the Frobenius number $\mathrm{F}(a)$ has been determined from an explicit certificate as in Proposition 8.1. In particular*

$$\mathrm{F}(a) \;=\; 1,\, 5,\, 11,\, 13,\, 11,\, 20,\, 31,\, 24,\, 27,\, 29,\, 43,\, 37,\, 49,\, 52,\, 63,\, 58,\, 69,\, 53,\, 75,\, 61 \qquad (a = 2, \ldots, 21),$$

and $\mathrm{F}(31) = 121$, $\mathrm{F}(32) = 127$, $\mathrm{F}(250) = 997$.

The certificates were produced by a Dijkstra shortest-path computation over $\mathbb{Z}/a$ with edge weights $a + k^2$ and self-checked before being emitted; each was then re-checked from the definition of S_a. For $2 \leq a \leq 61$ these values are sequence A357995 of [5] and agree with it term by term; for $2 \leq a \leq 42$ they agree with Table 1 of [3]. Beyond $a = 61$ they are not tabulated in any source we found, but they are determined by Theorem 5.1, which applies at every a in that range: they are an extension of a known computation, not a new result.

Theorem 8.3. *At each of the 235 values of $a \leq 250$ at which (1) has a solution, the value $3a + \max\{r\}$ predicted by Theorem 5.1 equals the Frobenius number certified in Theorem 8.2.*

Theorem 8.4. *For $1 \le a \le 250$, the condition* (1) *has no solution exactly at the fifteen values*

$$a \in \{1, 2, 3, 4, 5, 6, 7, 9, 10, 11, 13, 19, 21, 22, 30\}.$$

In particular $a = 30$ is the last such value in this range, and $R_3(31) = 28$.

The fifteen values are not new: Liu and Xin list $a = 4, 5, 7, 9, 10, 11, 13, 19, 21, 22, 30$ as the values where their Theorem 2.9 is needed and mark $a = 2, 3, 6$ as needing neither theorem, which is the same list for $a \ge 2$ (the value $a = 1$ is vacuous, there being no r with $1 \le r \le 0$). Their check runs to $a = 42$; the range above is $a \le 250$. Theorem 5.4 supersedes the question of what happens beyond: no $a \ge 31$ can join the list.

Independently of both routes, the ladder of Theorem 4.1 was evaluated directly.

Theorem 8.5. $\mathrm{F}(2) = 1$, $\mathrm{F}(3) = 5$, $\mathrm{F}(6) = 11$, $\mathrm{F}(31) = 121$, $\mathrm{F}(32) = 127$, $\mathrm{F}(250) = 997$, *by evaluation of the formula of Theorem 4.1; moreover at $a = 2, 3, 6$ the rungs 3 and 2 are both empty, and rung 1 is empty at $a = 2$ but not at $a = 3$ or $a = 6$.*

These six values agree with the Apéry certificates of Theorem 8.2, which were obtained by a completely independent argument — no Apéry vector and no Bellman inequality enters Theorem 4.1, and no evaluation of ι enters Proposition 8.1.

Finally, the controls. A formalisation can pass every check and prove nothing, so the following were stated and proved as well.

Proposition 8.6.
(1) F *is exact: at $a = 32$ neither 126 nor 128 is the Frobenius number, and likewise at $a = 31$ and $a = 250$ for the neighbours of the stated values.*
(2) *The hypothesis $4a \le K^2$ of Proposition 8.1 is load-bearing. The Apéry vector of $\langle 5, 6\rangle$ passes the residue condition* and *the Bellman inequalities at $K = 1$, and its maximum is $19 + 5$; without the hypothesis Proposition 8.1 would give $\mathrm{F}(5) = 19$, whereas $\mathrm{F}(5) = 13$.*
(3) *Neither half of each Boolean test is vacuous: the residue check and the Bellman check of Proposition 8.1 each fail on an explicit input; the congruence certificate of Lemma 5.3(2) fails at $25 = 3^2 + 4^2$ and succeeds at 21; the square-residue test of Lemma 5.3(3) succeeds at 49 and fails at 47.*
(4) *Lemma 5.3(1) does not prove $\iota = 4$ everywhere: $\iota(4) = 1$ while $\iota(7) = 4$.*
(5) *The hypothesis $a \ge 2$ is not decorative: $S_1 = \mathbb{N}$, so S_1 has no Frobenius number.*
(6) *The threshold in Theorem 5.4 is where the construction stops, not a convenience: see Remark 5.5.*

9. Verification

Every theorem, proposition, corollary and lemma above has been formally verified in Lean 4 (toolchain `v4.32.0`) against *Mathlib*, and the development is free of `sorry`; repository reference to be added at submission. The mathematical content is kernel-checked: Theorems 3.1, 4.1, 5.1, 5.4, 6.1, 6.2, 6.3, 7.1, Corollaries 3.2 and 4.2, Lemma 5.3 and Propositions 2.3, 3.3, 8.1 depend on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`, and the 1440-residue check inside Theorem 5.4 runs in the kernel and depends on no axioms at all. What rests on exhaustive computation is exactly the following, and nothing else. Inside Lemma 5.3: the 512 residue triples modulo 8, the 81 pairs modulo 9, the 16 pairs modulo 4, and the 360 residues modulo 360. Inside Theorem 5.4: the 1440 residue classes modulo 1440, each against at most eight candidates $r \le 63$ — this is the whole finite content of the proof, and it leaves no range of a to be checked separately. Inside Section 8: one certificate per a for Theorem 8.2, discharged in the kernel for $2 \le a \le 80$ and by compiled evaluation (Lean's `native_decide`) for $81 \le a \le 250$; one evaluation per a for Theorem 8.3; a single evaluation over the whole range $1 \le a \le 250$ for Theorem 8.4; and one evaluation per value for Theorems 7.2 and 8.5 — all by compiled evaluation. The theorems those evaluations instantiate — Theorems 4.1, 5.1 and 7.1, and Proposition 8.1 — are themselves kernel-checked, so compiled evaluation is used only to supply numerical inputs, never to establish a general statement. Every

numerical claim was first produced by an independent implementation in Python — the ladder was reimplemented from the membership criterion and reproduced all 41 entries ($2 \le a \le 42$) of Table 1 of [3] and their "which theorem applies" column — and the formal proof was then written against it. The two numerical measurements of Remarks 6.4 and 7.3 lie outside the formal development and are labelled as such where they appear; the remaining remarks are either commentary on the literature or, in the case of Remark 5.5, verified statements about the construction.

References

[1] A. Brauer and J. E. Shockley, *On a problem of Frobenius*, J. Reine Angew. Math. **211** (1962), 215–220.

[2] D. Einstein, D. Lichtblau, A. Strzebonski and S. Wagon, *Frobenius numbers by lattice point enumeration*, Integers **7**(1) (2007), #A15, 63 pp.

[3] F. Liu and G. Xin, *On Frobenius Numbers of Shifted Power Sequences*, Ann. Comb. (2026), published online 9 April 2026, doi:10.1007/s00026-026-00811-3; arXiv:2210.02722. Numbering here follows the arXiv version v4 (23 July 2025).

[4] F. Liu, G. Xin, S. Ye and J. Yin, *On the Frobenius number and genus of a collection of semigroups generalizing repunit numerical semigroups*, Semigroup Forum **110** (2025), 357–383; arXiv:2306.03459.

[5] OEIS Foundation Inc., *The On-Line Encyclopedia of Integer Sequences*, `https://oeis.org`; entries A357995 (Frobenius numbers of $A(a)$, 60 terms, $a = 2, \ldots, 61$; M. Marcus and B. Dobbelaere, 2022) and A002828 (ι).

[6] J. L. Ramírez Alfonsín, *The Diophantine Frobenius Problem*, Oxford Lecture Series in Mathematics and its Applications, vol. 30, Oxford University Press, 2005.

[7] J. C. Rosales and P. A. García-Sánchez, *Numerical Semigroups*, Developments in Mathematics, vol. 20, Springer, 2009.

[8] K. Song, *The Frobenius problem for shifted square sequences*, arXiv:2605.25542v3 (9 July 2026). Version v1 (25 May 2026) carried the title *The Frobenius problem for shifted square sequences starting with a multiple of* 4.