THE MAXIMAL HEIGHT OF A DIVISOR OF $x^{p^2q^2}-1$: A CONJECTURE OF RYAN, WARD AND WARD BEYOND ITS VERIFIED RANGE, AND TWO EXPLICIT FORMULAS

KOREA SUPERINTELLIGENCE LABS

Abstract. For a positive integer n let $B(n)$ be the largest height of a polynomial in $\mathbb{Z}[x]$ dividing x^n-1, the height of a polynomial being its largest coefficient in absolute value. Pomerance and Ryan determined $B(p^k)$ and $B(pq)$, and Kaplan determined $B(p^2q)$; the first case in which both exponents exceed 1 is the subject of Conjecture 3.3 of Ryan, Ward and Ward, which asserts that for primes $p<q$ the value $B(p^2q^2)$ is the larger of $H(\Phi_p\Phi_q\Phi_{p^2q}\Phi_{pq^2})$ and $H(\Phi_p\Phi_q\Phi_{p^2}\Phi_{q^2})$. Their paper verifies it for $2\le p<q<60$ and states no explicit formula for either height; Wang, who in 2015 proved the source's other two conjectures about two-prime n, wrote that it is not at all clear what the formula for $B(p^2q^2)$ should be; the conjecture is absent from the list of settled cases in Sanna's 2022 survey; and the database behind the original verification is no longer reachable. We verify the conjecture at 107 pairs outside $2\le p<q<60$, namely for $p\in\{2,3,5,7,11,13\}$ and q running over the primes from 61 up to 241, 199, 113, 113, 97 and 97 respectively, each with the exact values of $B(p^2q^2)$ and of both heights compared, and at four further pairs described below. We then give two explicit formulas for $B(p^2q^2)$ at a fixed p, each verified on a finite range and offered as a conjecture: $B(4q^2)=4$ for every prime $3\le q\le 241$, and $B(9q^2)=12$, 11 or 9 according as $q\equiv 5$, $q\equiv 7$, or $q\equiv 1,2,4,8 \pmod 9$, for every prime $5\le q\le 199$; on the same two ranges we also give the explicit formula for the height of the conjecture's own first product, which is the formula the source says it lacks. We record negative controls, in particular that $B(p^2q^2)$ is not a function of q mod p^2 in general, and we formulate the refinement that it is one once $q>p^2$. For that refinement we then certify two residue classes at $p=11$:

$$B(11^2\cdot 137^2)=B(11^2\cdot 379^2)=161,\qquad B(11^2\cdot 131^2)=B(11^2\cdot 373^2)=121=p^2,$$

the four primes involved all exceeding $p^2=121$ and the two members of each class differing by exactly $2p^2$. The first common value lies strictly above p^2, so that agreement is not the one the known lower bound $B(p^aq^b)\ge\min\{p^a,q^b\}$ forces; and the two classes carry different values, so on this range $B(11^2q^2)$ is neither eventually constant in q nor equal to $\min\{p^2,q^2\}$, which is what makes the refinement a statement with content. The largest instance is $n=11^2\cdot 379^2=17\,380\,561$, an exhaustive maximum over 512 products of degree up to n. A strictly finer invariant is certified as well: the eight sub-maxima obtained by fixing which of Φ_1,Φ_p,Φ_{p^2} occur agree across the class 16 mod 121 and separate the two classes. Every theorem and lemma below, and the propositions that record the controls, are machine-checked in Lean 4; the conjectures, the remarks labelled as computations outside the formal development, and one comparison of two columns of a machine-checked table are not.

1. Introduction

The objects. For $f\in\mathbb{Z}[x]$ write $H(f)$ for the *height* of f, "the largest coefficient of f in absolute value" [9, §1], and let Φ_n denote the nth cyclotomic polynomial, so that

$$x^n-1=\prod_{d\mid n}\Phi_d(x). \tag{1}$$

Pomerance and Ryan [8] introduced

$$B(n)=\max\{H(f): f\mid x^n-1 \text{ and } f\in\mathbb{Z}[x]\} \tag{2}$$

— the sequence [7] — and proved $B(p^k)=1$ and $B(pq)=\min\{p,q\}$ for distinct primes p,q [8, Proposition 2.2 and Lemma 2.1], together with the maximal order of B, namely $\limsup_{n\to\infty}\log\log B(n)/(\log n/\log\log n)=\log 3$ [8, (1.5)]. Kaplan [3] added $B(p^2q)=\min\{p^2,q\}$

for distinct primes p and q — equivalently $B(pq^2) = \min\{p, q^2\}$, which [9, §3] cites as [3, Theorem 6] and describes as reached "by means of a thorough case-by-case analysis" computing the height of every divisor of $x^{pq^2} - 1$ — together with upper bounds for general n that do not depend on its largest prime factor. From the latter, [9, Theorem 3.5] records — in its own words, "a rephrasing of a special case" of a theorem of Kaplan — that for fixed p, a and b the quantity $B(p^a q^b)$ takes only finitely many values as q runs over the primes. For the state of the subject we refer to Sanna's survey [10], which also records the general upper bounds of Bzdęga and of Zhang and the work of Decker and Moree [2] determining the coefficient sets of all 64 divisors of $x^{p^2q} - 1$.

By (1) and the irreducibility of the Φ_d over $\mathbb{Z}$, the divisors of $x^n - 1$ are, up to sign, exactly the $2^{\tau(n)}$ products $\prod_{d \in S} \Phi_d$ with S a set of divisors of n, so (2) is a maximum over a finite set. This is the reduction on which every computation of $B(n)$ rests; [9] states it in one sentence, and we prove it in Lemma 2.1.

The question. The formulas above cover $n = p^k$, $n = pq$, $n = pq^2$ and $n = p^2q$; Wang [11] later added a closed formula for $B(pq^3)$ and, for $p < q$, closed formulas for $B(pq^4)$ and $B(pq^5)$. In every one of these cases at most one of the two exponents exceeds 1. The first case in which both do is $n = p^2q^2$, where $\tau(n) = 9$ and (2) is a maximum over $2^9 = 512$ products. Ryan, Ward and Ward state, verbatim:

> **Conjecture 3.3.** Let $p < q$ be primes. Then $B(p^2q^2)$ is the larger of $H(\Phi_p(x)\Phi_q(x)\Phi_{p^2q}(x)\Phi_{pq^2}(x))$ and $H(\Phi_p(x)\Phi_q(x)\Phi_{p^2}(x)\Phi_{q^2}(x))$.

and, after the two worked examples we reproduce in Remark 2.5:

> In addition to not having a proof for this conjecture, we also lack an explicit formula for the height of the polynomial. The conjecture has been checked for the primes indicated in Table 1.

The row of that table reads "$p^2q^2 \mid 2 \le p < q < 60 \mid 463 \mid$ Conjecture 3.3". So the published verification range is $2 \le p < q < 60$; we return in Remark 2.3 to the count 463, which does not match that range, and in Remark 2.4 to the database [1] in which the underlying data were deposited.

Two later sources locate the conjecture. Wang [11], who proved the two companion conjectures of [9] about $B(pq^b)$ — their Conjectures 3.4 and 3.6 — and gave the closed formulas just mentioned, writes in his final section, verbatim: "It is clear that the same propositions used in this section will be useful in studying $B(p^aq^b)$, however it is not at all clear what the resulting formula should be, even for $B(pq^4)$ if $p > q$ or $B(p^2q^2)$." Sanna's survey [10], after noting that Ryan, Ward and Ward "made several conjectures on $B(n)$, for n having two, three, or four prime factors, based on extensive numerical computations", continues "Some of these conjectures were proved by Wang (2015)" and then enumerates precisely the $B(pq^b)$ statements: that $p \mid B(pq^b)$, that $B(2q^b) = 2$, that $B(pq^b) > p$ for $b \ge 3$, and that $q \equiv \pm r \pmod p$ with $b \le 5$ forces $B(pq^b) = B(pr^b)$. The p^2q^2 conjecture is not among them.

What is settled here. Nothing below proves Conjecture 3.3, and we do not claim a proof of any of the formulas we state; apart from the two lemmas of Section 2, which are ordinary proofs, every result below is an exhaustive finite computation, machine-checked, over an explicitly named set of pairs. What is new is the range, and the shape that the data on two lines turn out to have.

(1) **The conjecture outside its verified range.** Theorem 3.1 certifies Conjecture 3.3 at 138 pairs (p, q), of which 107 lie outside $2 \le p < q < 60$: the whole line $p = 2$ to $q = 241$, the whole line $p = 3$ to $q = 199$, the primes $61 \le q \le 113$ at $p = 5$ and $p = 7$, and the primes $61 \le q \le 97$ at $p = 11$ and $p = 13$. The largest instance certified is $n = 13^2 \cdot 97^2 = 1\,590\,121$. Each pair carries the exact values of $B(p^2q^2)$ and of both heights the conjecture compares, and each carries a certificate that the nine coefficient arrays used are the nine cyclotomic polynomials (Section 2).

(2) **Two explicit formulas.** Theorems 4.1 and 4.2: $B(4q^2)=4$ for every prime $3\le q\le 241$, and

$$B(9q^2)=\begin{cases}12, & q\equiv 5 \pmod 9,\\ 11, & q\equiv 7 \pmod 9,\\ 9, & q\equiv 1,2,4,8 \pmod 9,\end{cases}$$

for every prime $5\le q\le 199$. Both are formulas for $B(p^2q^2)$ at a fixed p, which no source we have found supplies at any fixed p. The value 9 is $\min\{3^2,q^2\}$, so the second formula also says that the lower bound $B(p^2q^2)\ge p^2$ of [9] is sharp on four of the six residue classes and misses by 3 and by 2 on the other two. We offer both as conjectures (Conjectures 4.5 and 4.6).

(3) **The missing height formula, on those two lines.** Theorem 5.1 supplies the explicit formula for the first of the two products of Conjecture 3.3 on the same two ranges, and records that on the whole line $p=3$ that first product is the one attaining the maximum.

(4) **Controls, and a refinement.** Proposition 6.1 certifies that neither branch of the conjecture's maximum suffices alone, that the lower bound $B(p^2q^2)\ge p^2$ of [9] is not an equality, that the certificate used throughout is not vacuous, and — delimiting item (2) — that $B(p^2q^2)$ is *not* a function of $q \bmod p^2$ in general. Section 7 formulates the refinement that it is one as soon as $q>p^2$, and reports the tests available.

(5) **Two residue classes at $p=11$, certified.** Section 8 puts two instances of that refinement on the same footing as item (1): $B(11^2\cdot 137^2)=B(11^2\cdot 379^2)=161$ and $B(11^2\cdot 131^2)=B(11^2\cdot 373^2)=121$, with Conjecture 3.3 holding at all four pairs. The four primes exceed $p^2=121$ and the members of each class differ by exactly $2\cdot 121$. One common value lies above p^2 and the other is exactly p^2, so the two classes together (Theorem 8.3) show that the residue of q modulo p^2 is doing real work. The largest instance is $n=11^2\cdot 379^2=17\,380\,561$. A finer invariant, the eight sub-maxima cut out by the inclusion pattern of Φ_1,Φ_p,Φ_{p^2}, is certified too (Theorem 8.4).

(6) **A way to check one exhaustive maximum in pieces.** Section 9 isolates two elementary identities about the depth-first walk over a power set that carries a running maximum: it splits at depth three into eight independent subtrees, and its seed enters only through a maximum. Together they let one evaluation of $B(p^2q^2)$ be replaced by eight that can be run separately and combined by arithmetic afterwards. That is what made the largest instance above checkable at all, and it applies to any exhaustive search of this shape.

Sections 2–7 are written so that the two things trusted in each statement are visible: which finite set was enumerated, and which part of the argument is a proof. Section 10 says what is machine-checked and what is not.

Changes from version 1. The first version of this paper contained Sections 2–7 in the form in which they stand here: every numbered statement of that version keeps its number and its wording, and nothing in it is withdrawn. What is added is the following.

- Section 8, the two residue classes at $p=11$ (Theorems 8.1–8.4 and Proposition 8.5). In the first version Conjecture 7.1 was supported only by computations outside the formal development; it now has four certified instances, and Theorem 8.3 states in certified form that the two classes carry different values, which is what makes the conjecture non-vacuous.
- Section 9, the splitting identities, and with them the certification of the largest instance in the paper, $n=11^2\cdot 379^2$, which is 10.9 times the largest instance of Theorem 3.1.
- Remark 8.7, which carries out the test at $p=13$ that Remark 7.2 only priced, and corrects that price; the body of Remark 7.2 is left as the first version had it. Both are computations outside the formal development.

- Section 10 is rewritten to cover the new material, and the abstract and this introduction are rewritten. The bibliography is unchanged except in [4], where a note records where the statement used in Remark 8.8 was read.

2. Preliminaries

All polynomials have integer coefficients, and $H(f)$ is as in Section 1. We write $\tau(n)$ for the number of divisors of n and, for primes $p<q$,

$$H_1(p,q)=H\big(\Phi_p\Phi_q\Phi_{p^2q}\Phi_{pq^2}\big),\qquad H_2(p,q)=H\big(\Phi_p\Phi_q\Phi_{p^2}\Phi_{q^2}\big), \tag{3}$$

the two heights of Conjecture 3.3, so that the conjecture reads $B(p^2q^2)=\max\{H_1(p,q),H_2(p,q)\}$. The nine divisors of p^2q^2 are

$$1,\quad p,\quad q,\quad p^2,\quad q^2,\quad pq,\quad p^2q,\quad pq^2,\quad p^2q^2,$$

so (2) is a maximum over $2^9=512$ products, each of degree at most $n=p^2q^2$.

The second product. The second product of (3) is flat and its height is known outright. Indeed $\Phi_p(x)\Phi_{p^2}(x)=1+x+\cdots+x^{p^2-1}$ — this is Lemma 3.2 of [9], $\prod_{k=1}^{a}\Phi_{p^k}(x)=\sum_{c=0}^{p^a-1}x^c$ — whence

$$\Phi_p\Phi_q\Phi_{p^2}\Phi_{q^2}=\Big(\sum_{c=0}^{p^2-1}x^c\Big)\Big(\sum_{c=0}^{q^2-1}x^c\Big),\qquad H_2(p,q)=\min\{p^2,q^2\}=p^2 \tag{4}$$

for $p<q$, the coefficients of the product being $1,2,\dots,p^2,\dots,p^2,\dots,2,1$. That polynomial is the witness for the lower bound

$$B(p^aq^b)\ \geq\ \min\{p^a,q^b\} \tag{5}$$

of [9, Theorem 3.1], which we refer to as (5) from here on to avoid a clash with the numbering of the present paper. So Conjecture 3.3 is equivalent to

$$B(p^2q^2)=\max\big\{p^2,\,H_1(p,q)\big\}\qquad (p<q), \tag{6}$$

and its content is an upper bound: no product of cyclotomic factors of $x^{p^2q^2}-1$ beats both the flat product (4) and $\Phi_p\Phi_q\Phi_{p^2q}\Phi_{pq^2}$. Every table below records H_2 as computed, and on the 42 pairs of Section 3 with $p\geq 5$ and $q\geq 61$ the recorded value is p^2, in agreement with (4).

Why a power set is the whole story.

Lemma 2.1. *Let $n\geq 1$ and $f\in\mathbb{Z}[x]$. Then $f\mid x^n-1$ if and only if f is associated in $\mathbb{Z}[x]$ to $\prod_{d\in S}\Phi_d$ for some set S of divisors of n. Consequently every divisor of x^n-1 has the same height as one of the $2^{\tau(n)}$ subset products, every subset product is a divisor, and*

$$B(n)=\max\Big\{H\Big(\prod_{d\in S}\Phi_d\Big):S\subseteq\{d:d\mid n\}\Big\}.$$

Proof. Each Φ_d is irreducible over $\mathbb{Z}$, hence prime in the unique factorisation domain $\mathbb{Z}[x]$, and x^n-1 is their product by (1). A divisor of a product of primes is a unit times a sub-product: by induction on the index set, either the first prime divides f, and one cancels it, or f is relatively prime to it and one passes to the remaining product. The units of $\mathbb{Z}[x]$ are ±1 and $H(-f)=H(f)$, so associated polynomials have equal heights; the converse inclusion is that a sub-product of a factorisation divides the whole. □

This is the reduction [9, §2] asserts in a sentence — "observe that any f that would give a maximal height is a product of cyclotomic polynomials since $x^n-1=\prod_{d\mid n}\Phi_d(x)$" — and it is what makes every computation below a statement about $B(n)$ rather than about a list of polynomials. It also reconciles the definition (2), in which f ranges over all integer divisors, with the variant used in [10], in which f ranges over the monic ones.

The certificate. The computations work with coefficient arrays, not with a symbolic cyclotomic polynomial: for each pair (p,q) a table $(f_d)_{d\mid p^2q^2}$ of nine integer arrays is built from $\Phi_1 = x-1$, $\Phi_p = 1+\cdots+x^{p-1}$, $\Phi_{pq} = (x^{pq}-1)(x-1)/\big((x^p-1)(x^q-1)\big)$ (for an explicit formula for Φ_{pq} see [4]) and repeated use of $\Phi_{mp}(x) = \Phi_m(x^p)$ for $p \mid m$. None of those identities is trusted. Instead each table is submitted to the *tower certificate*

$$\prod_{d\mid m} f_d = x^m - 1 \qquad \text{for every one of the nine } m \mid p^2q^2, \tag{7}$$

and the next lemma says that (7) pins the table completely.

Lemma 2.2. *Let $n \geq 1$ and let $(f_d)_{d\geq 1}$ be a family of polynomials in $\mathbb{Z}[x]$ with $\prod_{d\mid m} f_d = x^m - 1$ for every m dividing n. Then $f_d = \Phi_d$ for every d dividing n.*

Proof. Strong induction on $d \mid n$. Split $\prod_{e\mid d} f_e = f_d \cdot \prod_e f_e$ over the proper divisors e of d, and split $\prod_{e\mid d}\Phi_e = x^d - 1$ the same way. The two tails agree by the induction hypothesis, and the tail $\prod_e \Phi_e$ is monic, hence non-zero; cancel it in the domain $\mathbb{Z}[x]$. □

So a passing (7) certifies that the nine arrays are the nine cyclotomic polynomials, and no construction identity has to be believed. For each pair (p,q) the computation therefore records a five-field verdict

$$\big(\text{tower},\ \text{conj},\ B(p^2q^2),\ H_1(p,q),\ H_2(p,q)\big), \tag{8}$$

where *tower* is the outcome of (7), the three numbers are computed by enumerating all 512 subset products, and *conj* is the Boolean $B(p^2q^2) = \max\{H_1, H_2\}$, that is, Conjecture 3.3 at (p,q).

Remark 2.3 (the source's count). The "Data Points" cell of the p^2q^2 row of Table 1 of [9] reads 463, while the range in the same row, $2 \leq p < q < 60$, contains the 17 primes below 60 and hence exactly $\binom{17}{2} = 136$ pairs. The two are not consistent with each other. Other rows of the same table do agree with their ranges — the $2q^b$ rows for $b = 3$ and $b = 4$ read 24 and 20, exactly the numbers of primes q with $2 < q < 100$ and with $2 < q < 75$ — so the column is not counting something else throughout. We read the *range* as the source's claim, since that is what verifying a conjecture about pairs of primes means, and we reproduce all 136 of its pairs (Remark 2.5).

Remark 2.4 (the source's database). Table 1 of [9] points at [1] for the data behind every conjecture in the paper. Checked on 3 September 2026, that address redirects to its `https` form and returns HTTP 500, and the Wayback Machine holds no snapshot of the page, so the published range $2 \leq p < q < 60$ is the only surviving statement of what was checked for Conjecture 3.3. This bears on how the results below should be read. [9, §2] reports having computed $B(n)$ "for almost 300000 distinct n" with "4 or fewer prime factors" and n as large as 56 796 482, so the lost database may well have contained some of the values recomputed here. What is new below is therefore *coverage outside the source's published verification range*, not the assertion that no one has ever computed these numbers.

Remark 2.5 (a computation outside the formal development). Before any of the statements below was attempted, an independent implementation in C of the definitions (2) and (1) — build the nine factors, check (7), enumerate the $2^{\tau(n)}$ subset products — was checked against everything published that we could check it against. It reproduces the four numerical assertions printed in [9]: $B(3^25^2) = 12$, attained by $\Phi_3\Phi_5\Phi_{45}\Phi_{75}$ and not by $\Phi_3\Phi_5\Phi_9\Phi_{25}$; $B(5^211^2) = 25$, attained by $\Phi_5\Phi_{11}\Phi_{25}\Phi_{121}$ and not by the other product; $B(7^283^2) = 64$; and $B(3\cdot 31\cdot 1009) = 599$. In both of the first two the product the source names is among those attaining the maximum — at $(3,5)$ it is the only one, at $(5,11)$ one of six — and the product the source excludes is not, so the agreement is with the factorisations printed there and not merely with the numbers. It confirms Conjecture 3.3 at all 136 pairs of the source's range $2 \leq p < q < 60$, at a cost of 452.5 CPU-seconds. And it agrees with the independently computed table of $B(n)$ in [7] — a different algorithm by a different author,

which factors $x^n - 1$ outright — at all 655 values with $\tau(n) \le 12$ in the range $1 \le n \le 719$ covered there, among them the seven values $B(36)$, $B(100)$, $B(196)$, $B(225)$, $B(441)$, $B(484)$, $B(676)$ of the shape p^2q^2. There were no disagreements in any of these checks. These are measurements, not theorems; they are reported because a paper about a published function should begin by reproducing its published values.

3. The conjecture outside its verified range

Define the following set of pairs of primes:

$$\begin{aligned}\mathcal{K} = {} & \{(2,q) : 3 \le q \le 241\} \ \cup \ \{(3,q) : 5 \le q \le 199\} \\ & \cup \ \{(p,q) : p \in \{5,7\},\ 61 \le q \le 113\} \ \cup \ \{(p,q) : p \in \{11,13\},\ 61 \le q \le 97\},\end{aligned} \tag{9}$$

q always running over primes. These are $52+44+26+16 = 138$ pairs, of which the 107 with $q \ge 61$ lie outside the range $2 \le p < q < 60$ verified in [9]. Three further pairs inside that range, $(5,7)$, $(5,11)$ and $(7,11)$, are certified in Section 6 and used there, so that 141 distinct pairs are certified somewhere below, 107 of them outside the source's range.

Theorem 3.1. *For every pair $(p,q) \in \mathcal{K}$,*

$$B(p^2q^2) = \max\bigl\{H(\Phi_p\Phi_q\Phi_{p^2q}\Phi_{pq^2}),\ H(\Phi_p\Phi_q\Phi_{p^2}\Phi_{q^2})\bigr\},$$

that is, Conjecture 3.3 of [9] *holds at (p,q). Moreover the individual values $B(p^2q^2)$, $H_1(p,q)$ and $H_2(p,q)$ are as certified in the tables underlying this statement.*

Proof sketch. The whole statement rests on exhaustive computation, and on nothing else. For each of the 138 pairs the nine coefficient arrays are constructed, the tower certificate (7) is evaluated — nine polynomial identities, one per divisor of p^2q^2 — and all $2^9 = 512$ subset products are formed by a depth-first walk over the power set, their heights compared, and the verdict (8) recorded. The computation is exact integer arithmetic throughout; there is no estimate anywhere in it. The pairs are grouped into 14 blocks, and per block one evaluation establishes that every recorded verdict is the verdict the pair actually produces. Everything read off the resulting tables afterwards — that the tower and conjecture fields are true in every row — is a check on numerals, carried out by the proof kernel and not by the evaluator. Lemma 2.2 converts a true tower field into the statement that the nine arrays are the nine cyclotomic polynomials, and Lemma 2.1 converts the maximum over the 512 products into $B(p^2q^2)$. The largest instance is $(p,q) = (13,97)$, $n = 1\,590\,121$; the subset products there have degree up to n. □

Proposition 3.2. *At $(p,q) = (13,79)$ one has*

$$H_1(13,79) = 157 < 169 = 13^2 = H_2(13,79) = B(13^2 \cdot 79^2),$$

and $(13,79)$ is the only pair of $\mathcal{K}$ at which $H_1(p,q) < B(p^2q^2)$.

Proof. Read off the certified verdicts of Theorem 3.1, comparing the B and H_1 entries of all 138 rows. □

So outside the range verified in [9] the maximum of Conjecture 3.3 is still sometimes attained strictly by its second product alone — rarely, but not never. Inside that range the same happens at $(5,11)$, one of the two examples [9] prints in support of the conjecture (Proposition 6.1(1)). At the remaining 137 pairs of $\mathcal{K}$ the first product attains the maximum, alone or jointly.

Remark 3.3 (a computation outside the formal development). The C implementation of Remark 2.5 confirms Conjecture 3.3 at 386 pairs in all, with no failure: the whole line $p = 2$ for the 108 primes $3 \le q \le 599$, the line $p = 3$ for the 76 primes $5 \le q \le 397$, $p = 5$ for the 54 primes $7 \le q \le 269$, $p = 7$ for the 34 primes $11 \le q \le 163$, $p = 11$ for the 31 primes $13 \le q \le 151$ together with $q = 271$ and $q = 373$, $p = 13$ for the 26 primes $17 \le q \le 131$, and the 55 remaining pairs of the source's range with $17 \le p$. Of these 386 pairs, 250 lie outside $2 \le p < q < 60$. The largest is $(11,373)$, where

$n = 16\,834\,609$ and Φ_n has degree $\varphi(n) = 15\,263\,160$. The whole sweep cost about 0.88 CPU-hours. Only the 138 pairs of (9) are covered by Theorem 3.1; the rest are reported as evidence.

4. Two explicit formulas

The source's complaint that it lacks "an explicit formula for the height of the polynomial" has a companion: in the literature we have searched, no formula for $B(p^2q^2)$ itself appears at any fixed p, in contrast with $B(pq) = \min\{p, q\}$, $B(p^2q) = \min\{p^2, q\}$ and Wang's $B(2q^b) = 2$ for every prime $q > 2$ and $B(3q^b) = 3 \cdot 2^{\lfloor (b-1)/2 \rfloor}$ for every prime $q > 3$ [11, Theorem 1.5], none of which reaches $n = p^2q^2$. The two smallest lines produce one each.

Theorem 4.1. *For every prime q with $3 \le q \le 241$,*

$$B(4q^2) = 4.$$

Theorem 4.2. *For every prime q with $5 \le q \le 199$,*

$$B(9q^2) = \begin{cases} 12, & q \equiv 5 \pmod 9, \\ 11, & q \equiv 7 \pmod 9, \\ 9, & q \equiv 1, 2, 4, 8 \pmod 9. \end{cases}$$

Proof sketch for both. The certified tables of Theorem 3.1 at $p = 2$ and $p = 3$ carry the value $B(p^2q^2)$ in every row: 52 rows for the primes $3 \le q \le 241$ and 44 rows for the primes $5 \le q \le 199$. The two statements are then the assertion that a column of a finite table of numerals is constant, respectively that it agrees with the function $q \mapsto 12, 11, 9$ of q mod 9; both are checks on numerals, performed by the proof kernel, requiring no further computation with polynomials. The finite search behind them is the one described in the proof of Theorem 3.1: for each of those 96 primes, all 512 subset products of the nine certified cyclotomic factors of $x^{p^2q^2}-1$ were formed and their heights compared. Among the 44 primes of Theorem 4.2 the six residue classes modulo 9 are represented 8, 8, 7, 8, 6 and 7 times in the order $1, 2, 4, 5, 7, 8$, so no class is carried by a single witness. □

Theorem 4.1 is an upper bound in substance: the lower bound $B(4q^2) \ge \min\{4, q^2\} = 4$ is (5), and what is verified is that none of the 512 divisors built from the nine cyclotomic factors of $x^{4q^2}-1$ has a coefficient exceeding 4 in absolute value. Likewise the value $9 = \min\{3^2, q^2\}$ of Theorem 4.2 is exactly the value given by (5), so on the four classes $q \equiv 1, 2, 4, 8 \pmod 9$ that bound is sharp, while on $q \equiv 5$ and $q \equiv 7$ it misses by 3 and by 2 respectively. Written in terms of $p = 3$ the three values are p^2, $p^2 + p - 1$ and $p^2 + p$; we mention the shape without suggesting it persists, since $p = 2$ attains only p^2 and the values computed at $p = 5$ include 27 and 55, which are neither.

Remark 4.3 (not a ± phenomenon). Theorem 1.8 of [11] says that for odd primes $p < q < r$ with $q \equiv \pm r \pmod p$ and $b \le 5$ one has $B(pq^b) = B(pr^b)$. Theorem 4.2 is not an instance of that pattern with p replaced by p^2: the classes 5 and $-5 \equiv 4$ modulo 9 carry the different values 12 and 9.

Remark 4.4 (a computation outside the formal development). The C sweep of Remark 3.3 extends both formulas without exception: $B(4q^2) = 4$ for all 108 primes $3 \le q \le 599$, and the value table of Theorem 4.2 for all 76 primes $5 \le q \le 397$, whose distribution among the classes $1, 2, 4, 5, 7, 8$ modulo 9 is $12, 14, 12, 14, 13, 11$.

Neither theorem is a proof, and we state the general assertions as what they are.

Conjecture 4.5. $B(4q^2) = 4$ for every odd prime q.

Conjecture 4.6. For every prime $q \ge 5$, $B(9q^2) = 12$ if $q \equiv 5 \pmod 9$, $B(9q^2) = 11$ if $q \equiv 7 \pmod 9$, and $B(9q^2) = 9$ otherwise.

A proof of either would, as far as we can tell, be the first explicit formula for $B(p^2q^2)$ at a fixed p. For $p = 2$ the shape of the problem is unusually favourable and we record it, since it is where we would start. With $A(x) = \Phi_q(x) = 1 + x + \cdots + x^{q-1}$, the nine cyclotomic factors of $x^{4q^2} - 1$ are

$$x - 1, \quad x + 1, \quad x^2 + 1, \quad A(x), \quad A(-x), \quad A(-x^2), \quad A(x^q), \quad A(-x^q), \quad A(-x^{2q}),$$

by $\Phi_{2m}(x) = \Phi_m(-x)$ for odd $m > 1$ and $\Phi_{2m}(x^2) = \Phi_{4m}(x)$. So each of the 512 divisors is a product of at most three of $x\pm1$, x^2+1 with a subset of six substituted copies of a single $\{0,1\}$-polynomial, and two collapses are immediate: $A(x)A(-x) = 1+x^2+\cdots+x^{2q-2}$ and $A(x)A(x^q) = 1+x+\cdots+x^{q^2-1}$, both flat. A case analysis in the style of the one by which Kaplan [3] settled $B(pq^2)$ looks feasible.

5. The height of the conjecture's first product

By (4) the second product of Conjecture 3.3 has height p^2 outright, so "the polynomial" whose height [9] says it cannot express is the first, $\Phi_p\Phi_q\Phi_{p^2q}\Phi_{pq^2}$. On the two lines of Section 4 we can express it.

Theorem 5.1. (1) *For every prime q with $3 \le q \le 241$,*

$$H\big(\Phi_2\Phi_q\Phi_{4q}\Phi_{2q^2}\big) = 4 \qquad \text{and} \qquad H\big(\Phi_2\Phi_q\Phi_4\Phi_{q^2}\big) = 4,$$

so on the line $p = 2$ both products of Conjecture 3.3 attain its maximum.

(2) *For every prime q with $5 \le q \le 199$,*

$$H\big(\Phi_3\Phi_q\Phi_{9q}\Phi_{3q^2}\big) = \begin{cases} 12, & q \equiv 5 \pmod 9, \\ 11, & q \equiv 7 \pmod 9, \\ 9, & q \equiv 1,2,4,8 \pmod 9, \end{cases}$$

and moreover $B(9q^2) = H\big(\Phi_3\Phi_q\Phi_{9q}\Phi_{3q^2}\big)$ for every one of those 44 primes: on the whole line $p = 3$ the maximum of Conjecture 3.3 falls on its first product, which is to say that this product never has height below the value 9 given by (5).

Proof sketch. Every quantity here is a column of the same two certified tables that carry Theorems 4.1 and 4.2; the verdict (8) records H_1 and H_2 alongside B. The four assertions are, respectively, that the H_1 column at $p = 2$ is constant at 4, that the H_2 column at $p = 2$ is constant at 4, that the H_1 column at $p = 3$ agrees with the same function of q mod 9 as before, and that the B and H_1 columns at $p = 3$ agree row by row. All four are checks on numerals over tables already certified; no new enumeration of subset products is performed for this theorem, and the underlying finite search is that of Theorem 3.1 at those 96 primes. □

Remark 5.2 (a computation outside the formal development). Printing, for each pair, the sets S at which the maximum of Lemma 2.1 is attained explains the shape of Theorem 5.1(2). Over the 44 primes of that statement, whenever $B(9q^2) > 9$ — the 14 primes with $q \equiv 5$ or $7 \pmod 9$ — the maximum is attained at exactly one set, namely $\{\Phi_3, \Phi_q, \Phi_{9q}, \Phi_{3q^2}\}$, the first product of Conjecture 3.3; at the other 30 primes, where $B(9q^2) = 9$, there are 22 or 23 maximising sets, among them both of the products the conjecture compares. These counts come from the C implementation and are not part of the formal development.

6. Controls

A family of assertions produced by one program should be delimited from both sides.

Proposition 6.1. (1) (Neither branch alone.) $H_1(5,11) = 21 < 25 = B(5^2 11^2)$ *and* $H_2(3,5) = 9 < 12 = B(3^2 5^2)$. *So neither product of Conjecture 3.3 is $B(p^2q^2)$ by itself; both branches of the maximum are needed.*

(2) (The lower bound is not an equality.) $B(3^2 5^2) = 12 > 9 = 3^2$, *so* (5) *is not sharp in general.*

(3) (The certificate is not vacuous.) *The tower certificate* (7) *evaluates to true on the correct table at* $(3,5)$*, and to false both on the table obtained by adding* 1 *to one coefficient of one factor and on the table obtained by replacing* Φ_{p^2} *by* Φ_p*.*
(4) (The hypothesis $p<q$ is a normalisation.) *The verdict* (8) *is unchanged by exchanging the two arguments, at* $(3,5)$ *and at* $(5,11)$*.*
(5) (No general law modulo p^2.) $7\equiv 107 \pmod{25}$ *while* $B(5^2 7^2)=45\neq 49=B(5^2 107^2)$*, and* $11\equiv 109 \pmod{49}$ *while* $B(7^2 11^2)=69\neq 92=B(7^2 109^2)$*. So* $B(p^2q^2)$ *is not a function of* p *and* $q \bmod p^2$*, and Theorem 4.2 is a statement about* $p=3$ *rather than an instance of a general law.*

Proof sketch. Items (1), (2) and (5) are read off certified verdicts (8) at the six pairs $(3,5)$, $(5,7)$, $(5,11)$, $(5,107)$, $(7,11)$, $(7,109)$, produced by the same enumeration of 512 subset products as everywhere above; the pairs of (1) are the two examples printed in [9]. Items (3) and (4) are direct evaluations of (7) and of the verdict on the modified or swapped inputs. □

Proposition 6.2. *At each of the eleven pairs*

$$(2,3),\ (2,97),\ (3,5),\ (3,7),\ (3,101),\ (5,3),\ (5,7),\ (5,11),\ (7,11),\ (11,5),\ (11,13)$$

the same power-set enumeration, run over the four divisors of pq *and over the six divisors of* p^2q *and with the tower certificate passing in every case, returns* $B(pq)=\min\{p,q\}$ *and* $B(p^2q)=\min\{p^2,q\}$*. These are the theorems of Pomerance and Ryan* [8] *and of Kaplan* [3]*; they are recomputed here as a control on the engine, not offered as results.*

Proof sketch. For each pair the four, respectively six, coefficient arrays are built by the same code as in Section 2, (7) is checked over the divisors of pq respectively p^2q, and all 2^4 respectively 2^6 subset products are enumerated. A machine that returned a wrong value here could not be believed on p^2q^2; the list includes pairs with $p>q$ so that both sides of each minimum are exercised. □

7. A refinement, and what remains

The finiteness statement [9, Theorem 3.5] quoted in Section 1 says that for fixed p,a,b the list of values of $B(p^aq^b)$ is finite as q varies. Theorem 4.2 exhibits such a list, of length three, indexed by the residue of q modulo $p^2=9$; Proposition 6.1(5) shows that no such indexing can hold for all q at $p=5$ or $p=7$. In every violation in the computed range (Remark 3.3), the offending q is smaller than p^2: $q=7,13,17$ at $p=5$, $q=11$ at $p=7$, $q=29$ at $p=11$. That suggests the following.

Conjecture 7.1. Let p be a prime. For primes $q,q'>p^2$ with $q\equiv q' \pmod{p^2}$,

$$B(p^2q^2)=B(p^2q'^2).$$

Remark 7.2 (a computation outside the formal development). Over the 386 pairs of Remark 3.3 the refinement has no counterexample, and every residue class modulo p^2 all of whose computed members exceed p^2 was constant. The tests the computed range allows are: at $p=5$, the classes 7 mod 25 ($q=107,157,257$, all giving 49), 13 mod 25 ($q=113,163,263$, all giving 50), 17 mod 25 ($q=67,167$, both 50) and 22 mod 25 ($q=47,97,197$, all 40), together with fourteen further classes each of which was likewise constant; at $p=7$, the classes 4 mod 49 ($q=53,151$, both 63) and 10 mod 49 ($q=59,157$, both 98); and at $p=11$ the class 10 mod 121, whose two computed members $q=131$ and $q=373$ both exceed 121 and both give $B(11^2q^2)=121$ — the decisive test available, since $n=11^2\cdot 373^2=16\,834\,609$ and the pair cost 302 CPU-seconds. Consistently with the conjecture, the class 29 mod 121 is *not* constant ($B(11^2 29^2)=367$, $B(11^2 271^2)=331$) and its smaller member 29 is below 121. The next test of the same kind, $q=233$ against $q'=571$ at $p=13$, has $n=13^2\cdot 571^2=55\,100\,929$, which we priced at about 2 300 CPU-seconds and 2.2 gigabytes; it lies outside the sweep reported here.

What remains open is everything general: Conjecture 3.3 itself, for which we have only enlarged the tested range; the two formulas of Section 4, for which the $p = 2$ case has the concrete attack sketched there; the value of $H_1(p, q)$ at any $p \geq 5$, where the data show no pattern we can state; and Conjecture 7.1, which is at present a statement about 386 computed pairs. We note also that the cost of the direct approach grows as $2^{\tau(n)-1}$ times the polynomial work, so that the natural next shape $n = p^3q^3$, with $\tau(n) = 16$, is 128 times more expensive at the same pq and is out of reach beyond very small q by this method.

Sections 8 and 9 change that last item. Two of the classes listed in Remark 7.2 are certified below in the sense of Theorem 3.1, and together they say more than either says alone.

8. Two residue classes at $p = 11$

Conjecture 7.1 predicts an agreement, and an agreement at the value $\min\{p^2, q^2\} = p^2$ is no evidence for it: (5) forces $B(p^2q^2) \geq p^2$, so such an agreement says only that the bound is attained on both sides. The two classes treated here are chosen against that. Take $p = 11$, so $p^2 = 121$; in each of the residue classes 10 and 16 modulo 121 there are exactly two primes below 500, and in both classes both of them exceed p^2:

$$q \equiv 10 \pmod{121}: \quad q = 131,\ 373; \qquad q \equiv 16 \pmod{121}: \quad q = 137,\ 379. \tag{10}$$

In each class the two members differ by exactly $2 \cdot 121$, so an agreement between them is evidence for the modulus p^2 itself and not merely for some multiple of it: 131 and 373 are incongruent modulo $p^3 = 1331$, and so are 137 and 379. All four pairs lie far outside the range $2 \leq p < q < 60$ of [9] and outside the set $\mathcal{K}$ of (9).

Theorem 8.1. *Let $p = 11$ and let $q \in \{137, 379\}$, the two primes below 500 congruent to 16 modulo 121. Then Conjecture 3.3 of [9] holds at $(11, q)$, and*

$$B(11^2q^2) = H_1(11, q) = 161, \qquad H_2(11, q) = 121 = 11^2.$$

In particular $B(11^2 \cdot 137^2) = B(11^2 \cdot 379^2)$, which is the assertion of Conjecture 7.1 for this class, and the common value 161 is strictly larger than p^2, so the agreement is not the one (5) forces.

Theorem 8.2. *Let $p = 11$ and let $q \in \{131, 373\}$, the two primes below 500 congruent to 10 modulo 121. Then Conjecture 3.3 of [9] holds at $(11, q)$, and*

$$B(11^2q^2) = H_1(11, q) = H_2(11, q) = 121 = 11^2.$$

Proof sketch for both. As in Theorem 3.1, and by the same code: for each of the four pairs the nine coefficient arrays are built, the tower certificate (7) is evaluated over the nine divisors of 11^2q^2, all $2^9 = 512$ subset products are formed by a depth-first walk over the power set and their heights compared, and the verdict (8) is recorded; Lemma 2.2 turns the passing certificate into the statement that the nine arrays are the nine cyclotomic polynomials and Lemma 2.1 turns the maximum over subset products into $B(11^2q^2)$. The arithmetic is exact throughout. The four instances are

$$n = 2\,076\,481, \quad 2\,271\,049, \quad 16\,834\,609, \quad 17\,380\,561$$

for $q = 131, 137, 373, 379$; the last is the largest instance certified in this paper, 10.9 times the largest instance of Theorem 3.1, and there the subset products have degree up to n and the largest single factor Φ_n has degree $\varphi(n) = 15\,758\,820$. For $q = 379$ and $q = 373$ the walk over the power set was too long to be run as one unit and was split into eight independent parts by Lemmas 9.1–9.2, the parts being recombined by Lemma 9.3; for $q = 137$ and $q = 131$ it was run as one unit, and at $q = 137$ it was also run in the split form, the two agreeing (Proposition 8.5(5)). □

Theorem 8.3. *Among the four primes $131, 373, 137, 379$, all larger than $p^2 = 121$,*

$$B(11^2 \cdot 131^2) = B(11^2 \cdot 373^2) = 121, \qquad B(11^2 \cdot 137^2) = B(11^2 \cdot 379^2) = 161,$$

while $131 \equiv 373 \equiv 10$ and $137 \equiv 379 \equiv 16 \pmod{121}$. So $B(11^2q^2)$ is constant on each of the two classes and takes different values on them.

Proof. The four values are those of Theorems 8.1 and 8.2; the four congruences and the inequality $121 \neq 161$ are arithmetic. □

Proposition 6.1(5) delimits Conjecture 7.1 from above, by refuting a law modulo p^2 once members $q < p^2$ are admitted; Theorem 8.3 delimits it from below, by showing that what is left is not empty. On the range covered, $B(11^2q^2)$ is neither eventually constant in q — which would make the residue irrelevant — nor everywhere equal to $\min\{p^2, q^2\} = p^2$ — which would make it a restatement of (5).

A finer invariant. The 512 subsets split into eight blocks of 64 according to which of the three lowest-degree factors $\Phi_1, \Phi_p, \Phi_{p^2}$ they contain. For $0 \le j \le 7$ write $j = 4\varepsilon_1 + 2\varepsilon_p + \varepsilon_{p^2}$ in binary, let T_j be the subset of $\{1, p, p^2\}$ that contains 1, p, p^2 according as ε_1, ε_p, ε_{p^2} is 1, and set

$$\beta_j(p,q) = \max\Big\{H\Big(\prod_{d\in S}\Phi_d\Big) : S \subseteq \{d : d \mid p^2q^2\},\ S \cap \{1,p,p^2\} = T_j\Big\}, \tag{11}$$

so that $B(p^2q^2) = \max_{0\le j\le 7}\beta_j(p,q)$ by Lemma 2.1. Call $\beta(p,q) = (\beta_0, \dots, \beta_7)$ the *block profile* of the pair. The first product of Conjecture 3.3 lies in the block $j = 2$ (it contains Φ_p and neither Φ_1 nor Φ_{p^2}) and the second lies in the block $j = 3$.

Theorem 8.4. *With $p = 11$,*

$$\beta(11, 137) = \beta(11, 379) = (68, 99, 161, 121, 121, 121, 121, 121)$$

and

$$\beta(11, 373) = (109, 110, 121, 121, 121, 121, 121, 121).$$

Proof sketch. Each β_j is an exhaustive maximum over the 64 subset products of its block, computed by the walk of the proof of Theorem 8.1 with its first three inclusion decisions fixed; Lemma 9.2 makes the eight computations independent of each other, and Lemma 9.1 says that their maximum is the value of B certified in Theorems 8.1 and 8.2, which is how those two theorems are proved at $q = 379$ and $q = 373$. □

So on the class 16 mod 121 not only the maximum but every one of the eight sub-maxima is constant, a strictly finer invariant than B; and the profile separates the two classes, which the value alone does at $j = 2$ only. The maximum sits in the block of Conjecture 3.3's first product in both classes — alone at $q = 137, 379$, jointly with five other blocks at $q = 373$, where all of $\beta_2, \dots, \beta_7$ equal p^2.

Proposition 8.5. (1) (Not smaller.) $B(11^2q^2) \le 160$ *fails at* $q = 137$ *and at* $q = 379$*, and* $B(11^2 \cdot 131^2) \le 120$ *fails.*

(2) (Not the other class's value.) $B(11^2q^2) = 121$ *fails at* $q = 137$ *and at* $q = 379$*, and* $B(11^2q^2) = 161$ *fails at* $q = 131$ *and at* $q = 373$*: each of the two certified classes refutes the other's value.*

(3) (Not larger.) $162 \le B(11^2q^2)$ *fails at* $q = 137$ *and at* $q = 379$*, and* $122 \le B(11^2 \cdot 373^2)$ *fails; the certified values are exact and not lower bounds.*

(4) (Not vacuous.) $11, 131, 137, 373, 379$ *are prime,* $137 \neq 379$ *with* $137 \equiv 379 \pmod{121}$*, and* $131 \neq 373$ *with* $131 \equiv 373 \pmod{121}$*: each class has two witnesses.*

(5) (The splitting is faithful.) *At* $(11, 137)$ *the eight blocks* (11)*, recombined, give exactly the value that the undivided walk over the power set certifies.*

Proof sketch. Items (1)–(3) are comparisons of numerals against the certified verdicts (8) of the four pairs. In (4) primality is decided by trial division, and the tower certificate is an independent check on it: the array offered as Φ_q is $1 + x + \cdots + x^{q-1}$, which is Φ_q only for prime q, so a composite q would fail (7) rather than return a plausible value of B. Item (5) is the comparison of two separately computed numbers, one produced by the undivided walk and one assembled from the eight blocks by Lemmas 9.1 and 9.2. □

Remark 8.6 (computations outside the formal development). The C implementation of Remark 2.5, extended to print the eight block maxima (11) rather than only their maximum, gives $\beta(11,131) = \beta(11,373)$, so the class 10 mod 121 has a constant block profile as well, of which Theorem 8.4 certifies the member at $q = 373$. It also gives a third class, 46 mod 121, whose two primes below 500 are 167 and 409:

$$B(11^2 \cdot 167^2) = B(11^2 \cdot 409^2) = 253,$$

$$\beta(11,167) = \beta(11,409) = (71, 66, 253, 121, 121, 121, 121, 121).$$

A third class-invariant profile, a third value, and again the maximum in the block of Conjecture 3.3's first product. Certifying that class in the sense of Theorem 8.1 was priced at about 3.4 processor-hours and not attempted here. In all six pairs of these three classes the four blocks containing $\Phi_1 = x - 1$ cap at exactly p^2; that is an observation about those six pairs, all of which have $q > p^2$, and we do not claim it beyond them.

Remark 8.7 (computations outside the formal development). Remark 7.2 priced the test $q = 233$ against $q' = 571$ at $p = 13$, with $n = 13^2 \cdot 571^2 = 55\,100\,929$, at about 2 300 processor-seconds and 2.2 gigabytes. Both figures were too high: the two runs cost 731 processor-seconds together, of which 701 belong to the single run at $q = 571$, and 275 megabytes at their peak, and they give

$$B(13^2 \cdot 233^2) = B(13^2 \cdot 571^2) = 169 = 13^2,$$

so Conjecture 7.1 survives at $p = 13$ as well. Counting that class, twelve residue classes have now been tested in the same implementation, over $p = 2, 3, 5, 7, 11, 13$, each of them a class all of whose computed members exceed p^2, with no exception — 50 pairs, the largest being $n = 55\,100\,929$, at a cost of 3 595 processor-seconds. Two further observations from the same runs. First, on such a class not only B but the whole family of maximising subsets agrees between the members. Second, the point of Remark 4.3 persists at other primes: the fold $q \equiv \pm r$ of [11, Theorem 1.8], which holds modulo p for $B(pq^b)$, fails modulo p^2 even when both members exceed p^2, as

$$\begin{aligned} B(3^2 \cdot 23^2) = 12 \neq 9 = B(3^2 \cdot 13^2), &\qquad 23 \equiv 5,\ 13 \equiv 4 \equiv -5 \pmod 9, \\ B(5^2 \cdot 107^2) = 49 \neq 45 = B(5^2 \cdot 43^2), &\qquad 107 \equiv 7,\ 43 \equiv 18 \equiv -7 \pmod{25}, \\ B(7^2 \cdot 151^2) = 63 \neq 69 = B(7^2 \cdot 241^2), &\qquad 151 \equiv 4,\ 241 \equiv 45 \equiv -4 \pmod{49} \end{aligned}$$

show; one accidental agreement, $B(7^2 \cdot 59^2) = B(7^2 \cdot 137^2) = 98$ with $59 \equiv 10$ and $137 \equiv 39 \equiv -10 \pmod{49}$, is recorded so that the control is not overread.

Remark 8.8 (a reformulation, outside the formal development). Conjecture 7.1 can be stated without a congruence. For primes $p < q$ let ρ, σ be the unique non-negative integers with $\rho p + \sigma q = (p-1)(q-1)$; necessarily $\rho \leq q-2$ and $\sigma \leq p-2$, a larger value of either forcing the other below zero. These are the parameters of the description of the coefficients of Φ_{pq} due to Lam and Leung, which we take in the form [11, Proposition 2.3] states and attributes to [4]: writing $\Phi_{pq} = \sum_k a_k x^k$, one has $a_k = 1$ exactly for $k = ip + jq$ with $0 \leq i \leq \rho$ and $0 \leq j \leq \sigma$, $a_k = -1$ exactly for $k + pq = ip + jq$ with $\rho < i \leq q-1$ and $\sigma < j \leq p-1$, and $a_k = 0$ otherwise. (That statement calls ρ and σ positive; $\sigma = 0$ does occur, at $q \equiv 1 \pmod p$.) For $q > p^2$ the map $q \bmod p^2 \mapsto (\sigma, \rho \bmod p)$ is well defined and, checked exhaustively for $p = 3, 5, 7, 11, 13$, injective; so Conjecture 7.1 says exactly that $B(p^2q^2)$ is a function of the shape of Φ_{pq}. The value σ alone does not suffice: at $p = 11$ the four primes $131, 373, 197, 439$ all have $\sigma = 9$, and $B(11^2q^2)$ is $121, 121, 166, 166$ along them, the two pairs being separated by $\rho \bmod p = 0$ against 6.

The classes reachable next were priced and left: the class 46 mod 121 of Remark 8.6 at about 3.4 processor-hours, the class 64 mod 169 at $p = 13$ (that is, $q = 233$ against 571) at about 4.0, and the first test at $p = 17$, $q = 359$ against $q' = 937$ with $n = 253\,733\,041$, at 3.4 to 6.7 processor-hours in the C implementation alone. Certifying a whole class at a fixed p is of course out of reach of this method: a class is infinite, and the cost of a single pair grows quickly with pq.

9. Splitting one exhaustive maximum

Every value of B above is computed the same way: a depth-first walk over the power set of the factor list, carrying a running maximum. Written as a recursion on a list $L=(g_0,\dots,g_{k-1})$ of polynomials, a current partial product c and a current best $b\in\mathbb{Z}$,

$$S([\,],c,b)=\max\{b,H(c)\},\qquad S(g{:}L,\ c,\ b)=S\big(L,\ cg,\ S(L,c,b)\big), \tag{12}$$

so that $B(n)=S(L_n,1,0)$ where L_n lists the $\tau(n)$ cyclotomic factors of x^n-1 in order of increasing degree, by Lemma 2.1. The order affects the cost and not the value: with the largest factors last, each of them is multiplied only into products of the smaller ones, which are the shortest partial products available.

Read literally, (12) is one indivisible computation: the second recursive call consumes the first call's output as its seed. At $(p,q)=(11,379)$ it is about 65 minutes of processor time, which is more than we were willing to spend inside a single check. Two elementary identities make it eight computations that can be run separately and combined afterwards.

Lemma 9.1. *Let $L=(g_0,\dots,g_{k-1})$ with $k\ge 3$. For $0\le j\le 7$ write $j=4\varepsilon_0+2\varepsilon_1+\varepsilon_2$ in binary, put $c_j=g_0^{\varepsilon_0}g_1^{\varepsilon_1}g_2^{\varepsilon_2}$ and $\lambda_j(b)=S\big((g_3,\dots,g_{k-1}),c_j,b\big)$. Then for every $b\in\mathbb{Z}$*

$$S(L,1,b)=\lambda_7\big(\lambda_6\big(\lambda_5\big(\lambda_4\big(\lambda_3\big(\lambda_2\big(\lambda_1\big(\lambda_0(b)\big)\big)\big)\big)\big)\big)\big).$$

Proof. Three unfoldings of the second equation of (12). The recursion computes its exclusion branch first at every level, so the eight subtrees below the first three decisions appear in the plain binary order of j, with the most significant bit belonging to g_0. No induction and no hypothesis on the g_i is needed; the identity holds by definition once L is matched as $g_0{:}g_1{:}g_2{:}L'$. □

Lemma 9.2. *For every list L, every polynomial c and all $a,b\in\mathbb{Z}$,*

$$S\big(L,c,\max\{a,b\}\big)=\max\big\{a,\ S(L,c,b)\big\}.$$

Consequently, in the notation of Lemma 9.1, $\lambda_j(s)=\max\{s,\lambda_j(0)\}$ for every $s\ge 0$.

Proof. Induction on L. For $L=[\,]$ both sides are $\max\{a,b,H(c)\}$, by associativity of max. For $g{:}L$ the left side is $S(L,cg,S(L,c,\max\{a,b\}))$, which by the induction hypothesis applied twice — once to the inner call and once to the outer one — is $\max\{a,S(L,cg,S(L,c,b))\}$, the right side. The consequence follows from $\max\{s,0\}=s$ for $s\ge 0$. □

Lemma 9.2 is what makes the eight parts independent. Without it the parts would have to be evaluated in order, each one quoting the previous one's output as its seed; with it, each $\lambda_j(0)$ is a self-contained computation, and the chain of Lemma 9.1 is rebuilt from the eight resulting integers by arithmetic alone. The eight numbers are the block profile (11): $\lambda_j(0)=\beta_j(p,q)$ when L is the factor list of p^2q^2 and $q>p^2$, so that g_0,g_1,g_2 are Φ_1,Φ_p,Φ_{p^2}. What was a device for splitting a computation is therefore also the invariant of Theorem 8.4, which is why that invariant is available at no extra cost.

Lemma 9.3. *The verdict* (8) *of a pair (p,q) is determined by separately established values of its four components: if the tower certificate* (7) *holds at (p,q) and $B(p^2q^2)=b$, $H_1(p,q)=h_1$, $H_2(p,q)=h_2$, then the verdict is* $(\text{true},\ b=\max\{h_1,h_2\},\ b,\ h_1,\ h_2)$.

Proof. Unfolding the definition of the verdict, whose four fields are computed independently of one another. □

So a pair whose maximum was computed in eight parts re-enters the tables of Section 3 exactly as a pair swept in one piece does, and the lemmas that read formulas off those tables apply to it unchanged. This is how Theorems 8.1 and 8.2 are obtained at $q=379$ and $q=373$.

Remark 9.4 (measurements outside the formal development). Three practical points about the pairs of (10).

The split is even. At $(11,379)$ the three lowest-degree factors $\Phi_1, \Phi_{11}, \Phi_{121}$ have 2, 11 and 111 coefficients, against six others of up to 15 758 821; the cost of (12) is dominated by the multiplications at the deepest levels, where a partial product is more than 10^6 coefficients long, so the eight parts perform the same number of coefficient operations to within a part in 10^4. This is a property of the ordering by degree, not luck. It is a count and not a timing: the four files that carry the parts took between 16 and 27 minutes of wall clock each, for the reason given at the end of this remark.

The split is free. The eight parts together perform the same 2^9 subset products as the undivided walk. At $(11,137)$, where both were run, the split form — a file carrying the eight parts, their recombination and one cheap check on the length of the factor list — was checked in 230.1 seconds of wall clock, of which we count about six seconds as fixed overhead. The undivided walk was checked in a file that also carries three certificates for $(11,379)$, and that file took 322.1 seconds; so the undivided walk is bounded by that figure rather than measured by it, and the two forms were not timed against each other. What the split buys is not speed but the ability to stop: at $(11,379)$ the undivided walk is about 65 minutes of processor time and the split replaces it by four checks of about 16 minutes of processor time each.

Prices. The certified computations of Sections 8 and 9 cost 2.2 processor-hours of proof checking; the C computations belonging to the same two sections, including those of Remark 8.6 which were not repeated in Lean, cost 0.38 processor-hours. Compared pair by pair, the proof checking costs 19.4 times the C computation at $n = 1.7 \cdot 10^7$ and 24.9 times at $n = 2.3 \cdot 10^6$; the ratio falls as n grows, both implementations becoming memory-bound at the top of the recursion. Peak memory never exceeded 1.6 gigabytes. One caution about all such figures: processor time on our machine depends on how loaded it is. The same two C runs, at $(11,137)$ and $(11,379)$, were timed at 302.2 processor-seconds on a heavily loaded machine and at 207.3 on a lightly loaded one, a difference of 46%, so a price extrapolated from a measurement made under load runs high. The 302 processor-seconds quoted in Remark 7.2 is an unrelated measurement, of the single pair $(11,373)$; that run was re-timed at 139.1 seconds on the lightly loaded machine, and the near-agreement of the two numerals is a coincidence.

10. Verification

Machine-checked in Lean 4 [5] against Mathlib [6], in the form stated, are Lemmas 2.1, 2.2, 9.1, 9.2 and 9.3, Theorems 3.1, 4.1, 4.2, 5.1, 8.1, 8.2, 8.3 and 8.4, and Propositions 6.1, 6.2 and 8.5. Proposition 3.2 is an inspection rather than a separate formal theorem: the table entries it quotes are machine-checked as part of Theorem 3.1, and the comparison of two of their columns across all 138 rows was made by reading them. The statements labelled as computations outside the formal development — Remarks 2.5, 3.3, 4.4, 5.2, 7.2, 8.6, 8.7, 8.8 and 9.4 — are not machine-checked, and neither are the two conjectures of Section 4 nor Conjecture 7.1. The two lemmas of Section 2 are ordinary Mathlib developments and are checked by the proof kernel alone, depending on no axioms beyond propositional extensionality, quotient soundness and choice; the identification of H_2 with p^2 in the tables of Section 3 for the 42 pairs with $p \geq 5$ and $q \geq 61$ is checked with no axioms at all. Every other statement above depends in addition on finite evaluations carried out by Lean's compiled evaluator — one per block of pairs, one per part of a split walk, or one per directly evaluated control: twenty-two for the material of Sections 3–6 and thirty-five more for that of Sections 8 and 9, fifty-seven in all — each contributing one evaluation axiom; what those evaluations assert is exactly that the compiled arithmetic engine, run on a pair (p,q), returns the verdict (8), the block value (11) or the height that the table claims. Everything above that layer is kernel work: that a passing tower certificate identifies the arrays with the cyclotomic polynomials (Lemma 2.2), that the maximum over subset products is $B(n)$ (Lemma 2.1), that a split walk equals the undivided one (Lemmas 9.1–9.3), and that the certified tables imply the displayed formulas. Since the tables carry the numbers produced by the independent C implementation of Remark 2.5,

a disagreement between that implementation, which uses fixed-width integers, and the Lean engine, which uses arbitrary-precision ones, would make the evaluation fail and the development not compile; each checked block is therefore also an assertion that two independently written implementations agree on every one of its pairs. There is no `sorry` in the development and it declares no axiom of its own. Of the statements added for Sections 8 and 9, none depends on choice: Lemma 9.1 depends on propositional extensionality alone, Lemmas 9.2 and 9.3 on that together with quotient soundness and on no evaluation, and the primality, distinctness and congruence assertions of Proposition 8.5(4) about 11, 137 and 379 depend on no axiom at all. One gap is left open and we state it plainly: nothing here proves that the array arithmetic used by the engine computes multiplication of polynomials. A correctness proof relating the array representation to Mathlib's $\mathbb{Z}[x]$ would close it, and would leave Lemmas 2.1 and 2.2, which already sit above that layer, unchanged.

One point about the record of what was checked, since that record is what a machine-checked paper rests on. Twelve files were added for Sections 8 and 9, and all twelve were checked with exit status zero, in import order. For three of them — three of the four carrying the parts of the walk at $(11, 379)$ — the sources that were checked differ from the ones we keep only inside a module comment, which was rewritten after the run to remove a cost estimate the run itself had superseded. Comments are discarded before elaboration, so this cannot change what was proved; rather than merely assert that, two files were re-checked byte for byte after their comments were edited, with the same outcome and, in the one of the two that carries values, the same values. Re-running the other three would have cost a further 0.8 processor-hours.

A second point about the record. In the development, the assertion that Conjecture 3.3 holds at every pair certified here is carried by a set of declarations, and this version adds two of them: `conj33_class16` and `conj33_class10`, which are Theorems 8.1 and 8.2 at $(11, 137)$, $(11, 379)$, $(11, 131)$ and $(11, 373)$. So that assertion is larger here than in version 1, by exactly those four pairs. The twenty declarations that carried it in version 1 — the fourteen certified tables behind Theorem 3.1 and the six statements read off them — are unchanged, and nothing version 1 asserted is withdrawn or weakened.

The material of Sections 3–6 cost about 1.1 processor-hours of proof checking against 0.93 processor-hours of C computation, and that of Sections 8 and 9 a further 2.2 against 0.38. The repository is private; repository reference to be added at submission.

References

[1] N. C. Ryan, B. C. Ward and R. E. Ward, *Cyclotomic database search*, 2010, cited as the data source of [9]; the address given there, `http://www.eg.bucknell.edu/~theburg/projects/data/wards/cyclo.py/index`, redirects to its `https` form and returned HTTP 500 when checked on 3 September 2026, and the Wayback Machine's index holds no capture of it or of any address under `eg.bucknell.edu/~theburg/projects/data/wards`.

[2] A. Decker and P. Moree, *Coefficient convexity of divisors of* $x^n - 1$, Sarajevo J. Math. **9(21)** (2013), no. 1, 3–28; `doi:10.5644/SJM.09.1.01`; arXiv:1010.3938. The description above follows [10]; the e-print's introduction likewise records that $x^{p^2q}-1$ has 64 monic divisors and that its theorems cover them all. The printed version was not consulted.

[3] N. Kaplan, *Bounds for the maximal height of divisors of* $x^n - 1$, J. Number Theory **129** (2009), no. 11, 2673–2688; `doi:10.1016/j.jnt.2009.04.015`. Not on arXiv and behind the publisher's paywall; not consulted here. Its results are quoted above as [9] and [10] state them, and the reference to its Theorem 6 is [9]'s, not ours.

[4] T. Y. Lam and K. H. Leung, *On the cyclotomic polynomial* $\Phi_{pq}(X)$, Amer. Math. Monthly **103** (1996), no. 7, 562–564; `doi:10.1080/00029890.1996.12004786`. Cited from [9], which points at it for an explicit formula for Φ_{pq}, and not consulted here: it is not on arXiv and the journal is paywalled. The coefficient description used in Remark 8.8 is quoted from [11, Proposition 2.3], which states it and credits it to this paper, and it agrees with zbMATH's review of this paper (Zbl 0868.11016).

[5] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, Cham, 2021, pp. 625–635; `doi:10.1007/978-3-030-79876-5_37`.

[6] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, New York, 2020, pp. 367–381; `doi:10.1145/3372885.3373824`.

[7] *The On-Line Encyclopedia of Integer Sequences*, sequence A114536, $B(n)$, the maximal height of a divisor of $x^n - 1$; cited as `[oeis]` by [9]. The comparisons of Remark 2.5 use its b-file, computed in PARI/GP by A. Karttunen and tabulating $B(n)$ for $1 \le n \le 719$.

[8] C. Pomerance and N. C. Ryan, *Maximal height of divisors of* $x^n - 1$, Illinois J. Math. **51** (2007), no. 2, 597–604; `doi:10.1215/ijm/1258138432`. The numbered references above — Lemma 2.1, Proposition 2.2 and display (1.5) — are to the copy the first author posts at `https://math.dartmouth.edu/~carlp/PDF/cyclo12.pdf`; the journal's typeset version was not consulted and its numbering may differ.

[9] N. C. Ryan, B. C. Ward and R. Ward, *Some conjectures on the maximal height of divisors of* $x^n - 1$, Involve **3** (2010), no. 4, 451–457; `doi:10.2140/involve.2010.3.451`; arXiv:1009.5970. All quotations above are from the e-print, whose arXiv listing carries the single version v1 of 29 September 2010, and every numbered cross-reference above is to that version's numbering, read off the labels its own LaTeX source resolves: Theorem 3.1, Lemma 3.2, Conjecture 3.3, Conjectures 3.4 and 3.6, Theorem 3.5, and Table 1. The printed *Involve* version was not consulted; its numbering may differ.

[10] C. Sanna, *A survey on coefficients of cyclotomic polynomials*, Expo. Math. **40** (2022), no. 3, 469–494; `doi:10.1016/j.exmath.2022.03.002`; arXiv:2111.04034. Quotations above are from the e-print, version v2 of 15 December 2021.

[11] S. Wang, *Maximal height of divisors of* $x^{pq^b} - 1$, Int. J. Number Theory **11** (2015), no. 1, 67–79; `doi:10.1142/S1793042115500049`; arXiv:1312.3108. The quotation above and the numbering of Theorems 1.5 and 1.8 are from the e-print, version v2 of 10 January 2014; version v1 carries the same theorem numbering and the same quoted sentence. The printed version was not consulted and its numbering may differ.