A polynomial unweighted orbit counter with immediate stabilization

Checked reconstruction, 20 November 2026. The schedule was proposed by Astra 6.2 during the AHT audit. Sol 6.2 and Fable 5.2 independently reviewed its operations, complexity argument and conditional NP interface. This exposition incorporates their explicit definitions and local derivations, together with coordinator source checks. It is a scoped mathematical repair, not formal verification, an independent audit of all normal-surface theory, or a claim of priority over the literature.

The underlying interval operations are those of Agol–Hass–Thurston, *The Computational Complexity of Knot Genus and Spanning Area*, Section 4. The schedule changes, so the result below does not validate the printed algorithm's progress assertions or bound its exact sequence of cycles.

Statement and conventions

An input consists of an interval of integers $[1, N]$ and k signed isometries between subintervals. Endpoints are binary integers. Each pairing is an increasing translation or a decreasing reflection, and its domain and range have the same inclusive width. Orbits are the equivalence classes generated by all pairing edges.

Theorem. The modified schedule below returns the number of orbits in time polynomial in k and $B = 1 + \lceil \log_2(N+1) \rceil$.

Normalize a pairing by inversion so that its domain starts no later than its range. Delete identity pairings, including a singleton paired to itself. On a nonidentity singleton pairing retain its supplied or composed sign. This ambiguity is harmless: two distinct singleton sides cannot both satisfy the strict-closeness condition used below.

For an overlapping reflection $x \mapsto H - x$, with normalized minimum start a and maximum end d, restrict to

$$[a, \lceil H/2 \rceil - 1] \longrightarrow [\lfloor H/2 \rfloor + 1, d].$$

Here $H = a + d$. Discard an empty restriction. This **trim** deletes duplicate undirected edges and, when present, a fixed-point loop. It leaves disjoint sides.

An increasing pairing with domain start a, range start $c = a + t$, and width w is **periodic** when $0 < t \leq w$. Its span is the interval from the domain's start to the range's end. Two periodic pairings of periods t_1, t_2 may be merged when their spans overlap in at least $t_1 + t_2$ integers. Replace them by the periodic pairing of period $\gcd(t_1, t_2)$ on their union.

The modified schedule

Maintain the current endpoint n, the pairings, and an isolated-orbit counter, initially zero. Keep coordinates fixed: **never close an interior static gap**.

1. Normalize, remove identities and trim. Count and delete uncovered terminal suffixes. If no pairings remain, add all remaining n points to the counter and finish.
2. If a periodic merger is available, perform one and restart.
3. Choose a widest range ending at n. Denote its pairing by Q and its range by $[c, n]$. Break ties deterministically and hold Q fixed.
4. For each other pairing P whose range is contained in $[c, n]$, transmit that range by the largest legal inverse power of Q. Independently do the same to its domain if contained in $[c, n]$. Normalize and trim immediately. Repeat on this P while its range is contained in $[c, n]$. Delete an identity and restart the entire schedule if one arises.
5. Once all other ranges start below c, truncate the largest terminal suffix within $[c, n]$ disjoint from every side of every other pairing. Restrict Q accordingly, or delete it if its entire range is removed. Restart.

Call a completed fixed-Q stabilization followed by truncation a **macrostep**. Identity-deletion restarts and merger-only iterations decrease the pairing count and are charged separately. The algorithm finds suffixes from endpoints; it never enumerates their points.

Orbit correctness

An increasing periodic pairing has precisely the residue classes modulo its period on its span. In an overlap of length at least $t_1 + t_2$, a block of length t_1 and its translate by t_2 both fit. The second pairing identifies the first span's residues under addition of t_2. Interchanging the roles gives exactly the residue classes modulo $\gcd(t_1, t_2)$ on the union. Thus the merger preserves the full equivalence relation, including its interaction with other pairings.

A transmission connects each moved endpoint to its former endpoint using the unchanged Q. Replacing each edge by the edge between these transformed endpoints therefore preserves connectivity in both directions. Normalization, identity removal and trimming also preserve the equivalence relation.

For truncation, put $D = \max_{P \neq Q} d_P$, or zero when there is no competitor. Normalization makes every competitor's domain end at most its range end. The retained endpoint is therefore

$$n' = \max(c - 1, D).$$

For an increasing Q, the removed points are terminal portions of translation chains; every such chain attaches to a retained representative. Removing these tails cannot disconnect retained points. For a trimmed reversing Q, each removed point has its partner below c, in the retained interval. No whole orbit disappears in either case.

An uncovered terminal point is an isolated orbit of the current system, so counting and deleting it changes the required answer by exactly one. Keeping interior static points until they become terminal does not change the answer. Consequently the running isolated counter plus the number of current orbits is invariant.

Fixed-pivot stabilization

For an increasing Q of period $t > 0$, a contained side beginning at s is shifted left by

$$t(1 + \lfloor (s - c)/t \rfloor).$$

Its new start lies in $[c - t, c - 1]$. For a trimmed reversing Q, its inverse image lies wholly below c. Any side left unshifted already starts below c. Thus immediately after transmission both starts are below c; only a new trim can make the range eligible again.

At a repetition the pairing is a trimmed reflection whose domain starts below c. Only its range moves. If Q reverses orientation this turns P into a translation, so there is no further repetition. Such a pivot needs at most two transmissions per competitor.

If Q is increasing, write the repeating reflection as $x \mapsto H - x$. Its trimmed range begins at $u = \lfloor H/2 \rfloor + 1 \geq c$. The next transmission shifts it left by some $m \geq u - c + 1$, changing its constant to $H - m$. Neither normalization nor trimming changes this constant. At each repetition the integer

$$e = H - 2c + 2 \geq 0$$

therefore satisfies, if another repetition follows,

$$e_{\text{new}} \leq \lceil e/2 \rceil - 1.$$

Initially $e \leq 2N$, giving $O(B)$ transmissions and trims per competitor. Processing another competitor cannot undo this stabilization because Q and all previously processed pairings stay fixed. A macrostep uses $O(kB)$ transmissions and trims.

Every competitor originally ending at n was eligible by maximality, and had width no larger than Q's. After stabilization it starts below c and has not grown, so it cannot still end at n. Ineligible competitors are unchanged. Hence $n' < n$ and truncation makes progress.

If Q survives, $D = n' \geq c$, and some other range ends at n' with start strictly below c. There is no terminal gap or pending trim. At constant pairing count, the next pivot is different, its range starts strictly lower, and it transmits the shortened Q. A merger or identity-deletion restart falls outside this constant-count assertion and is charged to its count drop.

Potential and the fixed-coordinate telescope

Define, explicitly,

$$X = 4^{k'} \prod_{P \text{ live}} w_P,$$

where k' is the current number of pairings and the empty product is one. Transmissions preserve widths; restrictions do not increase them. A deletion loses a factor at least four. For a merger the union span has width at most $w_1 + w_2$, so the new pairing has width at most $w_1 + w_2 \le 2w_1w_2$. Its potential ratio is at most one half.

Start a block with reference Z equal to the current endpoint after cleanup. End the block after a pairing-count drop or after a completed macrostep with X at most half its starting value. A halving occurring within a macrostep may be charged at its completion since X cannot rebound.

In a constant-count block let $n_0 = Z$, n_i be the endpoint after step i, and c_i the starting point of its pivot's range. The previous section gives $c_{i+1} < c_i$, with no endpoint change between successive selections. The truncation ratio at step i is

$$\frac{n_i - c_i + 1}{n_{i-1} - c_i + 1}.$$

For $i \le r$, replacing c_i by the smaller c_r increases this ratio. Multiplication therefore telescopes:

$$\frac{X_{\text{after r}}}{X_{\text{start}}} \le \prod_{i=1}^{r} \frac{n_i - c_i + 1}{n_{i-1} - c_i + 1} \le \prod_{i=1}^{r} \frac{n_i - c_r + 1}{n_{i-1} - c_r + 1} = \frac{n_r - c_r + 1}{Z - c_r + 1}.$$

Extra trims only decrease the left-hand side. The argument is valid because the schedule retains interior gaps: the endpoint after one truncation really is the endpoint at the next selection.

Call a side $[u, v]$, of width $w = v - u + 1$, **strictly close** if

$$2w > Z - u + 1, \qquad \text{equivalently } w > Z - v.$$

If a surviving truncated pivot range is not strictly close, the telescope already gives a halving. In a block that continues, that range is close.

Close-side losses or a forced merger

At constant pairing count the number of strictly close sides cannot increase. A restriction reduces width and cannot move the right endpoint rightward. An inverse transmission moves a shifted side leftward, including for a reversing pivot. Normalization only exchanges the two side labels. Thus every lost close side is a permanent loss within this fixed-Z block.

Suppose no block end or close-side loss occurs. The next pivot Q transmits the preceding close range J, and its image stays close. Q cannot be a trimmed reverser: if its domain has width v, disjointness leaves at least v points to the right of that domain up to Z. No subinterval of it can be close. Therefore Q is increasing. Its range contains the close J and its domain contains the close image, so both sides are close. Without a loss they remain close after truncation.

Consider two successive such increasing transmitters P, Q. After its own truncation, P has close range J of start j and width w. Write Q's domain start as a, its range start as $c = a + t$, and its period as $t > 0$. Transmission moves J to start $j' = j - rt \ge a$ for an integer $r \ge 1$. After any domain transmission and normalization, P has starts $L \le R < c$, one equal to j'. Its sign remains increasing. Equality $L = R$ causes an identity deletion, ending the block; otherwise both sides remain close and their equal widths imply $R - L < w$. Thus P is periodic, of period $R - L$. Both sides of Q remain close after its truncation, so Q also remains periodic.

The old J fits before Z, and the new side at L is close with width w. Consequently $w > j - L$. Since $L \le j'$, $j - j' \ge t$ and $j \ge a + t$, this gives

$$w \ge t + \max(0, a - L).$$

After truncation, Q's span still covers the endpoint $R + w - 1$ of the live pairing P. Also $R \ge j' \ge a$. The overlap of their spans therefore has width

$$R + w - \max(L, a) \ge (R - L) + t.$$

This is exactly the periodic-merger condition, using the single period t, even if $r > 1$ and even if both sides of P moved. No later competitor operation changes P or Q. A merger must be available before the next macrostep, unless an intervening deletion has already ended the block.

Thus four consecutive macrosteps with no block end and no close-side loss are impossible: after step one, steps two and three give the two transmitters just analyzed. There are at most $2k'$ close-side losses in a block. Splitting at those losses gives $O(k' + 1)$ macrosteps per block. Loss of a transmitted image, loss of a pivot domain on truncation, identity deletion, and a vanishing pairing are all explicitly covered by this accounting.

Since $1 \le X \le (4N)^k$, there are $O(1 + kB)$ blocks. The total number of macrosteps is $O((k + 1)^2 B)$ and their transmissions/trims total $O((k+1)^3 B^2)$. All merger candidates can be scanned in pairs. Maximal powers use division, not successive unit shifts. Endpoint arithmetic and gcds use $O(B)$-bit numbers; the orbit counter does too. Even explicitly tracking X requires only $O(kB)$ bits. The bookkeeping therefore has polynomial bit cost. Combined with orbit correctness, this proves the theorem.

What this supplies to the NP verifier

For an admissible normal surface in an exterior triangulation with t tetrahedra, the original compressed construction gives at most $12t$ face-arc pairings. A normal curve on s triangles gives at most $3s$. Pairing endpoints are obtained from normal coordinates and edge-prefix sums, without enumerating individual intersection points. Disk boundaries connect their corners; attaching disk interiors neither joins distinct components nor creates a component without such corners. Therefore these orbits count components.

The verifier needs only three integer counts: one component for F, one for ∂F, and two for the normal surface with coordinates $2v(F)$. The last test detects two-sidedness; in an orientable ambient manifold it detects orientability of the already connected F. Ambient orientability must be a stated input promise or explicitly checked. Empty F or empty boundary has count zero and is rejected. No weighted orbit output is needed.

If W is the sum of disk coordinates, the number N of distinct edge intersections satisfies $N \le 4W$; they are not identical quantities. The supplied fundamental-coordinate bound gives

$$N \le 28t^2 2^{7t+2}.$$

Doubling coordinates only doubles this bound. Thus k and $\log N$ are polynomially bounded on the asserted short witnesses, and the replacement counter supplies the required polynomial-time component tests.

The genus calculation must also use **boundary-aware** Euler characteristic. With V distinct surface vertices, T triangular disks, Q quadrilateral disks and b boundary arcs,

$$E = (3T + 4Q + b)/2, \qquad \chi(F) = V - E + T + Q = V - T/2 - Q - b/2.$$

These values come directly from compressed face and edge counts. For connected orientable F with one boundary component, test $(1 - \chi(F))/2 \le g$. The printed closed-surface expression omits $b/2$ and cannot be retained.

Parity certifies boundary essentiality, but does not by itself identify the longitude. That implication additionally uses the integral null-homology of K and the fact that the accepted F is oriented. A spanning surface supplies a longitude λ zero in exterior homology; intersection with that surface shows a meridian μ has infinite order. The primitive boundary class $p\mu + q\lambda$, being zero in exterior homology, has $p = 0$ and $q = \pm 1$. It therefore extends to K through a collar without changing genus.

Conditional conclusion. The repair supplies the component-counting bottleneck in the NP argument, with the boundary correction above. The existence of a minimum-genus fundamental normal witness, its coordinate bound, normal-surface realization and the topological certificate premises remain imports. This is not a fresh proof of all of them. Weighted downstream algorithms need a separate audit, and the independent NP-hardness branch has not been reverified here. The ambient three-manifold is part of the input; no fixed-S^3 hardness conclusion follows.

Evidence and provenance

The exact checker verifies the full equivalence relation after every move on 53,667 exhaustive small systems and 10,000 deterministic random systems. Large exact reflection-family cases include endpoints with about 2,000 bits; their point sets

are not expanded. All checks passed. These finite checks support implementation fidelity; the universal bound is the proof above.

All six fresh scoped model responses completed. Their automatic ingestion failed because source labels included explanatory prose, and two responses listed omitted background material as unreadable. An explicit derivative reconciliation retains the raw failures, partial source declarations and omitted material; no mathematical field or verdict is changed. This is not a retry. The separate operations and complexity audits support the replacement; both NP-interface audits require the displayed boundary correction.

The public review includes the six readable audit records and the source-label reconciliation. Download the exact-check bundle or return to the review.