

Optimal searches for a moving target on paths and grids

Angel Ivanov Raychev

September 2026

Research manuscript in preparation

Abstract

An invisible target starts at an unknown vertex of a finite graph. Each day a searcher inspects at most m vertices simultaneously; after a miss, the target must traverse exactly one edge. We study feasibility, optimal strategies, and the minimum number of days guaranteeing capture. We give exact formulas for every budget on paths and rectangular grids with at most five rows. On every odd-sided box, in arbitrary dimension, a proved isoperimetric nesting theorem reduces the unrestricted game to two counts, giving an exact algorithm polynomial in the number of rooms. For rectangles we determine the exact neighborhood profiles and give explicit time formulas above quadratic budget thresholds for all side parities. At every feasible budget on an odd rectangle, a bounded joint frontier and a stabilization argument evaluate the exact time in a number of arithmetic stages depending only on width and budget. At the minimum budget, every odd rectangle has exact time $2wn - C_w$, with a width-only constant evaluated in $O(w)$ arithmetic stages. At any fixed feasible budget, sufficiently long odd rectangles need only the two pure solo residuals to decide the last day. For every even-area cylinder with a bipartite Hamiltonian cross-section, an optimal search decomposes into bounded local transitions between physical fronts. This gives a finite graph whose interior edge weights have explicit height formulas. Corner profiles and propagation bounds retain necessary geometric information at the two ends. No compatibility of arbitrary consecutive static minimizers is assumed. For any fixed transverse Cartesian box with A rooms and fixed $m > A/2$, put $D = 2m - A$ and $g = \gcd(A, D)$. We prove $T_m(n + 2D/g) = T_m(n) + 4A/g$ for all sufficiently large n , of both parities, with an explicit polynomial sufficient threshold. The argument classifies efficient boundaries of arbitrary optimal searches and replaces their interior crossings by direct height descent. For a fixed deadline t , a separate spatial transfer proves that the minimum budget is eventually An/t plus a periodic correction. We also determine every budget on the $3 \times 3 \times 3$, $4 \times 4 \times 4$, and $3 \times 4 \times 4$ boxes. Five inspections take exactly $18n - 36$ days on every $3 \times 3 \times n$ box with $n \geq 3$. Paths, two-row grids, and the first two cubes have complete Lean proofs quantifying over actual target walks. The account reconstructs the author's October 2019–March 2020 path work and documents subsequent AI-assisted development. A complete explicit classification for arbitrary finite boxes remains open here.

Contents

1	Introduction and results	3
2	The game and its possible positions	7
3	Compressing entire strategies	8
4	One quadrant profile for every forbidden initial triangle	12
5	Exact neighborhood profiles on odd rectangles	14
6	Every budget on odd rectangles: a one-day determination	19
7	A bounded joint-state calculation for every odd rectangle	23
8	Explicit formulas at large budgets on every odd rectangle	30
9	Exact profiles on every even-area rectangle	39
10	Corner restrictions and memory on an even-width half-strip	40
11	Uniform time bounds and exact large-budget searches on even widths	45
12	Exact symbolic transitions between physical fronts	49
13	Odd widths and even lengths above a quadratic budget	52
14	An exact isoperimetric order in every all-odd box	57
15	Paths: matching geometry and a probe-count potential	63
16	Two rows: an exact pair of counters	66
17	Three rows: a historical rank and the exact time	67
18	Four rows as a uniform-theorem corollary	70
19	Five rows	70
20	Seven rows with five inspections per day	80
21	Minimum-budget searches on every odd width	82
22	Eventual affine periods in odd cylinders	87
23	Higher-dimensional boxes	94
24	Eventual affine periods for every fixed transverse box	107
25	Exact finite interfaces for even-area cylinders	118
26	A spatial transfer matrix for a fixed deadline	122
27	Formal verification, reproducibility, and open problems	125

1 Introduction and results

A princess moves through a building every night. Each morning a searcher may open several doors. The princess knows the search plan, may start anywhere, and must move to an adjacent room after each unsuccessful day. The problem is to guarantee capture as quickly as possible. Even on a single corridor the answer has a parity correction: dividing the amount of uncertainty by the apparent daily progress is not always enough.

This is a minimum-time version of the Hunters and Rabbit game. A path P_n has n vertices, and $P_{n_1} \square \cdots \square P_{n_d}$ denotes the Cartesian box with those side lengths. Write $T_m(G)$ for the optimal guaranteed number of inspection days, with value ∞ when no finite guarantee exists. Section 2 specifies the order of inspection and compulsory movement precisely.

1.1 How the results fit together

The general two-dimensional results precede their narrow-grid specializations. Odd rectangles have exact compatible neighborhood profiles and an exact two-count search recurrence. At every feasible budget, Theorem 6.1 reduces their optimal time to $2\tau - 1$ or 2τ , using one explicit solo recurrence; Theorem 7.1 decides between those values with state and arithmetic-stage bounds depending only on width and budget. The simpler scalar rule is now proved for all sufficiently long odd rectangles at every feasible budget (Theorem 7.5). At the minimum budget, Theorem 21.1 gives $T_{b+1}(P_{2b+1} \square P_n) = 2(2b + 1)n - C_{2b+1}$ for every odd $n \geq 2b + 1$, with a width-only constant evaluated in $O(b)$ stages. The remaining scalar-rule question is confined to shorter rectangles at intermediate budgets. For every even-area rectangle, Theorem 9.1 gives its exact neighborhood profile without assuming dynamic nesting. Theorems 8.1, 11.3, and 13.1 give closed optimal-time formulas for all side parities above their stated quadratic budget thresholds. For every even-area rectangle, the physical interface theorem 25.1 supplies an exact finite local description at every feasible budget. Its more general statement covers even-area cylinders with bipartite Hamiltonian cross-sections. Theorem 12.3 and its balanced height construction evaluate the interior connections explicitly; the finite boundary graph retains the remaining shared-budget choices.

The corner analysis has a common capacity formula: $C_h(s) = \max\{s(s - 1)/2, s(s - h)\}$ is the largest quadrant support of surplus at most s after its first permitted diagonal is fixed at height h . Theorem 10.1 applies it to prescribed corner omissions in even-width half-strips. Efficient small survivors are forced into specific triangles, yielding lower bounds on the time before the opposite corner can become possible again.

The narrower complete classifications below retain the low-budget cases not covered by those uniform theorems. Their overlapping proofs are replaced by specializations wherever possible. The later box and cylinder sections preserve the higher-dimensional results, while the formalization section states precisely which ordinary proofs have complete Lean counterparts. None of these results is presented as a complete explicit solution for all rectangles at all budgets. In particular, the finite interface theorem can require large preprocessing as width and budget grow. Its contribution is an exact geometric normal form and strategy reconstruction; it is not a claim of a practical implementation for every input or an asymptotic improvement over the previous effective eventual-period theorem.

1.2 The exact narrow-grid formulas

On a path, $T_m(P_n) = 1$ if $n \leq m$. With one inspection per day, $T_1(P_2) = 2$ and $T_1(P_n) = 2n - 4$ for $n \geq 3$. For $n > m \geq 2$,

$$T_m(P_n) = \left\lceil \frac{2n - 4}{2m - 1} \right\rceil + \varepsilon, \quad (1)$$

where $\varepsilon = 1$ exactly when $n \equiv m + 1 \pmod{2m - 1}$ and n, m are not both even. This theorem is the domain-corrected result of the author's original 2019–2020 work. We give a shorter construction and a matching-based lower proof, and formalize the complete statement in Lean.

For two rows and $n \geq 2$, one daily inspection is insufficient, and $m \geq 2n$ gives one day. In the remaining range put $N = n - 1$, $M = m - 1$. Then

$$T_m(P_2 \square P_n) = \left\lceil \frac{2N}{M} \right\rceil + \mathbf{1}_{\{M \text{ even and } N \equiv M/2 \pmod{M}\}}. \quad (2)$$

For three rows and $n \geq 2$, the one-day threshold is $3n$, the two-day threshold is $\lfloor 3n/2 \rfloor$, and one inspection is insufficient. For $2 \leq m < \lfloor 3n/2 \rfloor$, write $3n - 4 = q(2m - 3) + r$, $0 \leq r < 2m - 3$. The answer is

$$\begin{cases} 2q, & r = 0, \\ 2q + 1, & r > 0 \text{ and } m \geq r + 4 - \mathbf{1}_{\{n \text{ even, } q \text{ odd}\}}, \\ 2q + 2, & \text{otherwise.} \end{cases} \quad (3)$$

For four rows and $n \geq 4$, the one-day threshold is $4n$, the two-day threshold is $2n$, and budgets at most two are insufficient. For $3 \leq m < 2n$, write $2n - 2 = q(m - 2) + r$, $0 \leq r < m - 2$. The answer is

$$\begin{cases} 2q, & r = 0, \\ 2q + 1, & r = 1 \text{ or } (r \geq 2 \text{ and } m \geq 2r + 4), \\ 2q + 2, & \text{otherwise.} \end{cases} \quad (4)$$

Rotating a shorter four-row board reduces it to a preceding family. In particular the minimum feasible budgets give $T_2(P_3 \square P_n) = 6n - 8$ and $T_3(P_4 \square P_n) = 4n - 4$ in their stated nondegenerate ranges. For five rows and $n \geq 5$, the minimum feasible budget is three, and

$$T_3(P_5 \square P_n) = 10n - 20. \quad (5)$$

The even-length proof uses a finite symbolic boundary certificate valid for every length, together with a history-dependent potential. The odd-length proof follows from the explicit neighborhood profiles. For every even $n \geq 6$, putting $h = 5n/2$ also gives

$$T_4(P_5 \square P_n) = 2 \left\lceil \frac{2h - 8}{3} \right\rceil.$$

Theorem 19.5 gives a closed quotient-and-remainder formula for every larger budget on these even-length boards. Theorem 19.7 supplies the closed answer for every odd length as well. In particular, four inspections per day take $2 \lceil (5n - 9)/3 \rceil + 2\mathbf{1}_{\{3 \nmid n\}}$ days when $n \geq 5$ is odd. Together with rotation of shorter boards, these give explicit formulas and attaining strategies for every five-row board and every budget. The wider odd-rectangle theorem extends the closed analysis to every width $2b + 1$ when $m \geq b^2 + 1$.




Figure 1: Optimal time at the smallest feasible daily budget for four narrow families and the five-row three-probe theorem. Each line is a proved formula, not a fit to search data.

1.3 The structural results

The common difficulty is temporal. A small possible-position set may have a particularly favorable next neighborhood, but that neighborhood need not contain an equally favorable set for the following day. We retain a history bit where necessary, and also prove a more general comparison theorem. A monotone, cardinality-preserving set map C with

$$N(C(A)) \subseteq C(N(A))$$

compresses a whole strategy without increasing any day's inspection budget. Idempotent maps supply an exact optimal-strategy normal form. This assertion is stronger than a one-step isoperimetric inequality.

We construct such operators from odd paths and square diagonal shifts. They give exact recurrences on smaller families of states in arbitrary dimension. More strongly, on every box whose side lengths are all odd, we prove that the restricted simplicial orders minimize open neighborhoods and that neighborhoods of prefixes are prefixes. An induction using fiber compressions and a weighted exchange lemma reduces the unrestricted game to two cardinalities. The resulting exact algorithm is polynomial in the number of rooms and returns an optimal schedule. A recurrence remains distinct from a closed expression: the general box problem is not declared solved by the existence of an exponential search over all possible-position sets.

The three-dimensional example has a particularly concise answer:

m	0:4	5	6	7:8	9:11	12	13:26	27 or more
$T_m(P_3^3)$	∞	18	10	6	4	3	2	1.

Its complete lower and upper proofs are checked in Lean. A separate height construction proves

$h(P_3 \square P_3 \square P_n) = 5$ for $n \geq 3$, where $h(G)$ is the minimum feasible daily budget. Moreover, Theorem 23.5 gives the exact time $18n - 36$ at that budget for every $n \geq 3$, of either parity.

The even cube P_4^3 is also completely classified, with physical Lean proofs at every budget. Eight inspections per day require exactly forty days; nine require twenty. Its eight-probe lower bound retains the shape of the possible-position set: using only its cardinality would falsely suggest thirty-two days. The complete tables for P_4^3 and $P_3 \square P_4 \square P_4$ appear in Theorem 23.4.

For a fixed bipartite Hamiltonian cross-section H of order A , every sufficiently long cylinder has hunting number $\lfloor A/2 \rfloor + 1$. For each fixed $m > A/2$, its optimal time is $2An/(2m - A) + O_{A,m}(1)$, with no restriction on side parity. For every transverse Cartesian box we prove more: the exact time has an effective eventual affine period for both longitudinal parities. Efficient searches must have one of finitely many translated frontiers; controlling the remaining days permits insertion and deletion of a physical interval in any optimal strategy. A separate fixed-deadline theorem applies to every finite cross-section: the smallest budget for capture within t days is An/t plus an eventually periodic correction. These are distinct parameter regimes, and neither substitutes for a uniform closed formula for all boxes, budgets, and deadlines.

1.4 Prior work and provenance

Britnell and Wildon [6] studied the princess puzzle with one daily inspection, including the path value $2n - 4$. Haslegrave studied the associated evasion game [9]; related moving-target search appears in Beluhov and Kolev [2]. Abramovskaya, Fomin, Golovach, and Pilipczuk [1] developed Hunters and Rabbit results, including the rectangular-grid feasibility threshold $\lfloor \min(a, b)/2 \rfloor + 1$. Bolkema and Groothuis proved isoperimetric nesting and the hunting number for hypercubes [5]. We use precise versions of these results; minimum-time claims do not follow from feasibility alone. Parity-restricted boundary minimization on the binary cube goes back to Körner and Wei [12]. Compression proofs and local-to-global principles have a substantial earlier literature, including Bezrukov and Serra [4]; their closed-neighborhood theorem is distinct from the open-neighborhood argument given here.

Dmitry Kamenetsky recorded conjectural two- and three-inspection path formulas in March 2018 [10, 11]. Equation (1) proves those special cases and extends them to arbitrary budgets. The author’s independent path work took place between October 2019 and March 2020 and was written in two Bulgarian conference manuscripts [14, 15]. The one-probe restriction omitted from a printed general formula is made explicit here. The original conference documents are preserved.

The time-budget inverse parameter is also considered in recent work [3]; we do not claim the general optimization question as new. Recontamination can matter substantially on general graphs [7], so a presumed monotone search is never used as an unproved lower-bound assumption. Classical closed-neighborhood isoperimetry for Cartesian grids [13] supplies one of our higher-dimensional applications. Our use of that theorem is distinguished from the open-neighborhood argument for odd rectangles.

Remark 1.1 (Status of this manuscript). This is a research manuscript in preparation, not an arXiv submission. The ordinary theorems have internal independent proof reviews; exact Lean coverage is stated in Section 27. Attribution of the 2026 extensions is still undergoing a focused literature review. The manuscript makes no blanket claim that every derived ingredient is new to the literature.

2 The game and its possible positions

Let $G = (V, E)$ be a finite simple graph without isolated vertices, and let $m \geq 1$ be the daily inspection budget. Every connected board with at least two rooms satisfies this graph assumption; a one-room board takes one inspection and is treated separately. The target chooses an unknown initial vertex. On day $t = 1, 2, \dots$, the searcher inspects a set $S_t \subseteq V$ with $|S_t| \leq m$. If the target occupies a vertex of S_t , it is caught. Otherwise it must move along exactly one edge before the next day. Inspections within a day are simultaneous. The searcher receives no information about an unsuccessful inspection other than the miss.

Write $T_m(G)$ for the minimum number of days in which capture can be guaranteed, and put $T_m(G) = \infty$ if there is no finite guarantee. The associated feasibility threshold is

$$h(G) = \min\{m \geq 1 : T_m(G) < \infty\}.$$

We consider deterministic guarantees. Before capture, every observation is a miss. Thus an adaptive strategy has only one continuing observation history and is represented by a sequence $(S_t)_{t \geq 1}$. This sequence formulation does not restrict the searcher's guaranteed performance.

For $A \subseteq V$, its open neighborhood is

$$N(A) = \{v \in V : \text{some } u \in A \text{ satisfies } uv \in E\}.$$

Define the possible-position sets immediately before inspection by

$$B_1 = V, \quad B_{t+1} = N(B_t \setminus S_t).$$

The following elementary interpretation is used in every subsequent argument.

Lemma 2.1. *For each $t \geq 1$, a vertex v belongs to B_t if and only if there is a walk $v_1, \dots, v_t = v$ in G such that $v_i \notin S_i$ for every $1 \leq i < t$. Consequently capture is guaranteed by day t exactly when $B_t \subseteq S_t$.*

Proof. The assertion about walks follows by induction. At $t = 1$ every vertex is an admissible initial position. For the inductive step, $v \in B_{t+1}$ means that some neighbor $u \in B_t$ avoids S_t . Append v to a walk ending at u supplied by the induction hypothesis. Conversely, the penultimate vertex of any such walk belongs to $B_t \setminus S_t$, so its final vertex belongs to B_{t+1} .

A walk ending outside S_t avoids all inspections through day t , whereas every surviving walk is caught that day if its endpoint lies in S_t . This is exactly the stated inclusion. $\square$

Since G has no isolated vertices, a nonempty set has a nonempty open neighborhood. Hence the capture condition is also equivalent to $B_{t+1} = \emptyset$. We use this equivalent form only under that hypothesis. A board with a single room is handled directly: one inspection on the first day suffices. Thus an inability to make a move is never silently treated as capture.

Two monotonicity observations will be useful. Starting from fewer possible positions cannot make a fixed inspection sequence less successful, and adding inspected vertices cannot make it less successful. Both follow by induction from monotonicity of N and set difference in their respective arguments. In particular, probes outside the current possible-position set may be wasted, but they cause no problem for the model or the lower bounds.

2.1 Two initial parity classes

Suppose G is bipartite, with vertex classes V_0, V_1 . Keep track of the two possible initial colors separately:

$$R_1^s = V_s, \quad R_{t+1}^s = N(R_t^s \setminus S_t), \quad s \in \{0, 1\}.$$

Then

$$B_t = R_t^0 \dot{\cup} R_t^1, \quad R_t^s \subseteq V_{s+t-1 \pmod{2}}.$$

Indeed, neighborhood and deletion distribute over unions, and every edge reverses color. The two sets therefore stay disjoint while exchanging their current colors each night. Allocations to the two initial classes sum to at most the day's inspection budget. We will state explicitly whether a pair of counts is indexed by initial color or by current color; the latter convention swaps the two counts after every move.

This separates the bookkeeping common to the path, ladder, and grid arguments. The work specific to a graph is to determine how small a neighborhood can be after a given number of inspections, and when a sequence of such bounds can actually be attained.

3 Compressing entire strategies

An isoperimetric inequality bounds one neighborhood. To preserve an optimal search, the replacement sets must also respect containment and remain compatible after successive moves. The following criterion provides that stronger conclusion.

Definition 3.1. A *strategy compression* on G is a map $C : 2^V \rightarrow 2^V$ satisfying

$$A \subseteq B \implies C(A) \subseteq C(B), \tag{6}$$

$$|C(A)| = |A|, \tag{7}$$

$$N(C(A)) \subseteq C(N(A)). \tag{8}$$

Idempotence is an additional property, not part of this definition.

Theorem 3.2 (Strategy comparison and a fixed-state normal form). *Let C be a strategy compression. Every successful strategy from A with prescribed daily budgets $m_1, \dots, m_T$ has a successful strategy from $C(A)$ with the same budgets. If C is idempotent, every starting set fixed by C has an optimal strategy whose possible-position sets and post-inspection survivor sets are all fixed by C .*

Proof. Let A_t be the original possible-position set before day t , and R_t its survivors, so $A_{t+1} = N(R_t)$ and $|A_t| - |R_t| \leq m_t$. Maintain a new actual possible-position set $B_t \subseteq C(A_t)$, initially $B_1 = C(A_1)$. Inspect

$$Q_t = B_t \setminus C(R_t).$$

Monotonicity and equal cardinalities give

$$|Q_t| \leq |C(A_t) \setminus C(R_t)| = |A_t| - |R_t| \leq m_t.$$

The new survivors are $B_t \cap C(R_t)$, and

$$B_{t+1} = N(B_t \cap C(R_t)) \subseteq N(C(R_t)) \subseteq C(N(R_t)) = C(A_{t+1}).$$

When the original survivors are empty, $C(R_t) = \emptyset$, so the new strategy also captures every remaining possibility.

Fixed sets are closed under intersections: if X, Y are fixed, monotonicity gives $C(X \cap Y) \subseteq X \cap Y$, and cardinality forces equality. They are closed under neighborhoods for the same reason: $N(X) \subseteq C(N(X))$ when X is fixed. If C is idempotent, $C(R_t)$ is fixed. Starting the preceding construction from a fixed B_1 , every actual survivor $B_t \cap C(R_t)$ and every next possible-position set is therefore fixed. Restricted strategies are actual strategies, so their optimum equals the unrestricted optimum. $\square$

3.1 Products and simultaneous compression

Lemma 3.3 (Product lift). *A strategy compression on a graph A lifts to $A \square H$ by applying it separately in every H -indexed A -fiber. Idempotence is preserved.*

Proof. Write S_h for the fiber at h . Its neighborhood is

$$N(S)_h = N_A(S_h) \cup \bigcup_{u \sim h} S_u.$$

For the lifted operator $\widehat{C}$, the corresponding neighborhood is

$$\begin{aligned} N(\widehat{C}(S))_h &= N_A(C(S_h)) \cup \bigcup_{u \sim h} C(S_u) \\ &\subseteq C(N_A(S_h)) \cup \bigcup_{u \sim h} C(S_u) \\ &\subseteq C\left(N_A(S_h) \cup \bigcup_{u \sim h} S_u\right). \end{aligned}$$

The last inclusion uses monotonicity separately on each input subset of the union. Cardinality, monotonicity and idempotence hold fiber by fiber. $\square$

Compositions of strategy compressions are strategy compressions, since

$$N(C(D(S))) \subseteq C(N(D(S))) \subseteq C(D(N(S))).$$

Suppose finitely many such operators strictly decrease a common nonnegative integer potential whenever they change a set. If that potential has a uniform upper bound B , cycling through all operators $B+1$ times reaches a common fixed point for every input. A cycle with a change decreases the potential, and a cycle without a change fixes every operator. The resulting map is a composition with a *uniform* number of factors, hence is monotone and cardinality preserving; it is also idempotent. This avoids assuming that an input-dependent stopping rule preserves monotonicity.

3.2 Odd path fibers

Lemma 3.4 (Odd-path compression). *On P_{2q+1} , $q \geq 1$, number vertices $0, \dots, 2q$. Replace the selected vertices of each parity by the same number of vertices at the beginning of*

$$0, 2, \dots, 2q \quad \text{or} \quad 1, 3, \dots, 2q - 1.$$

This is an idempotent, parity-preserving strategy compression.

Proof. Monotonicity, cardinality and idempotence follow from the definition. An even-parity k -set has at least $\min(k, q)$ odd neighbors. If $k \leq q$, omit an unselected even vertex $2j$ and match $2i$ to $2i+1$ for $i < j$, and to $2i-1$ for $i > j$. These matches are distinct and include every selected even vertex. If $k = q+1$, all q odd vertices are neighbors. A nonempty odd-parity k -set has at least $k+1$ even neighbors: its consecutive blocks in the spacing-two order each have one more neighbor than their size, and different blocks have disjoint neighbor intervals. The two prefixes attain these bounds, and their neighborhoods are prefixes of the opposite parity. Consequently each part of $N(C(S))$ is contained in the corresponding same-cardinality prefix of $N(S)$, proving (8) even when S contains both colors. $\square$

For several odd coordinates, combine their lifted operators. Every changing operation decreases the sum of coordinates, so simultaneous compression has the normal form of Theorem 3.2. Its fixed sets satisfy $v \in S \implies v - 2e_i \in S$ whenever the latter vertex belongs to the box. These are lower ideals separately in each coordinate-parity pattern. They need not be lower ideals in the whole checkerboard color: for example, $\{(1, 1)\}$ in a 3×3 square is fixed by the two path operators but does not contain $(0, 0)$.

3.3 Exact recurrences on odd cylinders

Let $G = H \square P_{2q+1}$, where H is any finite graph and $q \geq 1$. The product lift and Theorem 3.2 give an exact normal form from the full board. It uses two vectors indexed by $V(H)$: a_v^0 counts even path coordinates in fiber v , and a_v^1 counts odd path coordinates. Their ranges are

$$0 \leq a_v^0 \leq q+1, \quad 0 \leq a_v^1 \leq q.$$

Thus they refer to the parity of the *path coordinate*, and do not require H to be bipartite. Define

$$\phi_0(k) = \min(k, q), \quad \phi_1(k) = \begin{cases} 0, & k = 0, \\ k+1, & k > 0. \end{cases}$$

For survivor vectors $b^0 \leq a^0$, $b^1 \leq a^1$, the exact next state $F(b)$ is

$$F(b)_v^0 = \max \left(\phi_1(b_v^1), \max_{u \sim v} b_u^0 \right), \quad (9)$$

$$F(b)_v^1 = \max \left(\phi_0(b_v^0), \max_{u \sim v} b_u^1 \right), \quad (10)$$

with an empty maximum interpreted as zero. Path edges supply the first term, and H -edges the second. All contributions in a fiber are prefixes of the same parity, so their union has exactly the maximum length. The required number of inspections is

$$d(a, b) = \sum_{v \in V(H)} (a_v^0 - b_v^0 + a_v^1 - b_v^1).$$

They are realized by inspecting the removed suffix of each prefix.

For a fixed budget m , form an edge $a \rightarrow F(b)$ whenever $b \leq a$ and $d(a, b) \leq m$. Then $T_m(a)$ is exactly the shortest-path distance to zero in this state graph, or infinity if zero is unreachable. Equivalently,

$$T_m(a) = 1 + \min_{\substack{b \leq a \\ d(a, b) \leq m}} T_m(F(b)), \quad T_m(0) = 0, \quad (11)$$

where the shortest-path interpretation specifies the solution of the possibly cyclic equations. For all budgets simultaneously, let $W_t(a)$ be the least constant budget that wins within t days. Then

$$W_{t+1}(a) = \min_{b \leq a} \max\{d(a, b), W_t(F(b))\}, \quad (12)$$

with $W_0(0) = 0$, $W_0(a) = \infty$ for $a \neq 0$. Backpointers in either recurrence give actual inspection sets.

If $A = |V(H)|$, the number of states is exactly

$$K = ((q+1)(q+2))^A \leq (2q+2)^{2A}.$$

There are at most K^2 state-survivor pairs. Explicit construction of their transitions takes $O((A + |E(H)|)K^2)$ elementary operations, apart from integer bit costs. For fixed H , this is polynomial in the odd longitudinal length. Any successful path can have cycles removed, so a finite optimum is at most $K - 1$. This proves an exact algorithm, not a uniform polynomial bound in dimension or a closed time formula.

When H is bipartite, one may instead index vectors by global color. For a coloring χ , put $\varepsilon_p(v) = p - \chi(v) \pmod{2}$. A single current color has transition

$$[\Phi_p(b)]_v = \max\left(\phi_{\varepsilon_p(v)}(b_v), \max_{u \sim v} b_u\right).$$

The two current-color vectors then become $(\Phi_1(b^1), \Phi_0(b^0))$. This convention is useful when comparing the cylinder recurrence with the two-color grid arguments below.

3.4 Square fibers and equal-sided boxes

Abramovskaya, Fomin, Golovach and Pilipczuk [1, Section 3.1, Lemma 2] prove that diagonal shifts do not increase the neighborhood size of a one-color set in a square. The following fiber calculation strengthens the scalar comparison to (8) and applies it to both colors. That stronger statement is what allows whole strategies to be compressed.

Lemma 3.5 (Square diagonal compression). *On $\{0, \dots, n-1\}^2$, compress each line $x+y=k$ toward increasing priority for smaller y , keeping its selected cardinality. This is an idempotent, parity-preserving strategy compression. The same is true on lines $x-y=k$, again toward smaller y .*

Proof. Only the neighborhood inclusion needs proof. Consecutive square diagonals have lengths differing by one. Index both from the common increasing- y end. From a diagonal of length L to one of length $L+1$, a selected index j has neighbors $j, j+1$. A nonempty a -set therefore has at least $a+1$ neighbors; a prefix attains this bound with a prefix. From length L to length $L-1$, the indices are $j-1, j$, clipped to the valid range. A proper nonempty a -set has at least a neighbors, and the full set has $L-1$. Indeed, if it has c consecutive blocks and occupies e endpoints, its neighbor count is $a+c-e$; for a proper set $c \geq e$. Prefixes again attain the bounds with prefixes.

Fix an output diagonal. Its two input diagonals, after compression, contribute prefixes from the same end. Their union is the longer prefix. Each contributing length is no larger than its original contribution, hence no larger than the original union's size. The compressed union is therefore contained in the prefix selected by compressing that original neighborhood. This argument works on either parity and on their union. Reflecting x proves the other diagonal direction. $\square$

Theorem 3.6 (Normal form on equal-sided boxes). *On P_n^d , $n \geq 2$, an optimal strategy from the full board can be chosen so that its possible-position sets and survivors are lower ideals for the legal moves*

$$v \mapsto v + e_i - e_j, \quad v \mapsto v - e_i - e_j, \quad i < j.$$

The conclusion respects arbitrary prescribed daily budgets. The same operators are available within every group of equal-length coordinates in an arbitrary Cartesian box. Odd-coordinate path compressions can be included simultaneously.

Proof. Lift both square operators to every pair $i < j$ of equal coordinates. They compress toward smaller x_j on lines $x_i + x_j = k$ and $x_i - x_j = k$. Use the potential

$$\mathcal{E}(S) = \sum_{v \in S} \sum_{i=1}^d i v_i.$$

A nontrivial plus-diagonal transfer by $\delta > 0$ decreases it by $(j - i)\delta$, and a minus-diagonal transfer by $(i + j)\delta$. An odd-coordinate prefix transfer also strictly decreases it. The potential is uniformly bounded; on P_n^d the bound $n^d(n - 1)d(d + 1)/2$ suffices. Uniform cycling therefore gives an idempotent strategy compression. Its fixed sets are precisely the ideals for the displayed local moves, with the additional $-2e_i$ conditions when odd-coordinate operators are included. Apply Theorem 3.2. $\square$

In dimension two the two-diagonal fixed sets are the downward pyramidal sets of the square isoperimetric argument. They differ from ordinary checkerboard corner ideals. In any dimension they give an exact finite recurrence over their fixed family, since both neighborhoods and intersections remain fixed. The theorem does not bound the number of these ideals by a polynomial in d or in n .

Remark 3.7 (Why even path factors require another idea). Every parity-preserving strategy compression on an even path is the identity. Each endpoint is the unique degree-one vertex in its color. Applying (8) to its singleton forces that singleton to be fixed. Fixed sets are closed under neighborhoods, so every N^t of either endpoint is fixed. These are the spacing-two prefixes from the two ends. Intersections of suitable prefixes isolate every singleton; fixed intersections follow from monotonicity and cardinality, without idempotence. Thus all singletons, and consequently all sets, are fixed. This rules out a nontrivial even-path factor reduction within this specific comparison framework. It does not rule out richer state labels or different comparison theorems.

4 One quadrant profile for every forbidden initial triangle

The square-root and pronic-root bounds have a common extension that retains prescribed omissions near a corner. Neighborhoods below are always taken in the whole nonnegative quadrant. Forbidden rooms constrain the support; they are not deleted from the graph.

Theorem 4.1 (Punctured-quadrant profile). *Fix a checkerboard color $p \in \{0, 1\}$ and an integer $r \geq 0$. Require a finite color- p support to omit the diagonals $x + y = p, p + 2, \dots, p + 2(r - 1)$, and put $h = 2r + p$. For every integer $s \geq 0$, the largest possible support with neighborhood surplus at most s has cardinality*

$$C_h(s) = \max \left\{ \frac{s(s - 1)}{2}, s(s - h) \right\}. \quad (13)$$

Every cardinality up to this maximum is attainable with surplus at most s . Consequently the exact minimum neighborhood size of a k -room support, $k > 0$, is

$$k + \min\{s \geq 0 : k \leq C_h(s)\}.$$

Thus its surplus is the smaller of two integer quadratic roots.

For $h = 0$ and $h = 1$ the formula gives respectively $\lceil \sqrt{k} \rceil$ and the least s with $k \leq s(s-1)$. The case $h = 2$ forbids the even origin; $h = 3$ forbids both odd neighbors of the origin. These exact static profiles do not assert that their minimizers can be chained inside arbitrary earlier beliefs.

Proof. Use Lemma 3.5 inside a square large enough that neither the original nor compressed support or neighborhood reaches its far edges. It compresses each diagonal toward smaller y , preserving cardinality and all the forbidden diagonals, without increasing the neighborhood. Let a_j be the resulting count on $x + y = 2j + p$, whose length is $\ell_j = 2j + p + 1$, and put $\epsilon_j = \mathbf{1}_{\{a_j = \ell_j\}}$. Counts with $j < r$ vanish. The exact output count on $x + y = 2j + p + 1$ is

$$n_j = \max\{a_j + \mathbf{1}_{\{a_j > 0\}}, a_{j+1} - \epsilon_{j+1}\}.$$

The two contributions are prefixes of the same diagonal. The upper input contracts by one only when full. There is also an initial output count $n_{-1} = a_0 - \epsilon_0$; for $p = 0$ this is identically zero, correctly representing the nonexistent lower diagonal.

Put $\delta_j = n_j - a_j$ for $j \geq 0$, and $\delta_{-1} = n_{-1}$. These numbers are nonnegative and sum to the compressed surplus δ . Moreover

$$\delta_j \geq \mathbf{1}_{\{a_j > 0\}}, \quad \delta_j \geq a_{j+1} - a_j - \epsilon_{j+1}.$$

List the occupied diagonals as $j_1 < \dots < j_L$ and let f_t count the full ones through j_t . Telescoping below j_t gives surplus at least $a_{j_t} - f_t$; the remaining output has surplus at least $L - t + 1$. Hence

$$a_{j_t} \leq \delta - L + t - 1 + f_t. \tag{14}$$

If no occupied diagonal is full, the initial contribution is positive, so $L \leq \delta - 1$. Summing (14) with $f_t = 0$ gives

$$|S| \leq L\delta - \frac{L(L+1)}{2} \leq \frac{\delta(\delta-1)}{2}.$$

If the first full diagonal has occupied rank u , then $f_u = 1$ and $j_u \geq r + u - 1$, so its size is at least $h + 2u - 1$. Equation (14) gives $L \leq \delta - h - u + 1 \leq \delta - h$. Summing the same equation with $f_t \leq t$ now gives $|S| \leq L\delta \leq \delta(\delta - h)$. The empty set is separate. Since C_h is nondecreasing on nonnegative integers, compression proves the upper bound for every original set of surplus at most s .

Two nested constructions prove sharpness. For $h \geq 1$, use the proper diagonal counts

$$a_{r+i} = i + 1 \quad (0 \leq i \leq s - 2).$$

Their initial surplus is one, and each occupied diagonal contributes one more. Their size is $s(s-1)/2$ and their surplus is s . For $s \geq h + 1$, instead use the $s - h$ full consecutive diagonals

$$a_{r+i} = h + 1 + 2i \quad (0 \leq i < s - h).$$

Their initial surplus is h , and each full diagonal contributes one. Their size is $s(s - h)$ and surplus is again s .

A partial last diagonal interpolates between consecutive capacities in either family. After a nonempty completed ramp or full block, its downward contribution fits in the existing neighborhood, and a positive partial fill increases surplus by exactly one. For the first full block, a partial first diagonal of $v < h + 1$ rooms has surplus $v + 1 \leq h + 1$; the full first diagonal has surplus $h + 1$. The full-block branch dominates when $s \geq 2h - 1$. For $h \geq 2$ its first dominant positive level is within its valid range $s \geq h + 1$; for $h = 1$ it starts at $s = 2$. Thus whichever family supplies the new maximum also fills every gap above the preceding maximum. For $h = 0$, the full blocks alone give s^2 , starting with the origin, and partial final diagonals fill every intermediate size. This proves all asserted attainments. $\square$

Lemma 4.2 (Two candidates for convex capacities). *Let $A, B : \mathbb{N} \rightarrow \mathbb{N}$ be nondecreasing, unbounded and discretely convex, with $A(0) = B(0) = 0$. Define $\alpha(k) = \min\{s : k \leq A(s)\}$ and $\beta(k) = \min\{s : k \leq B(s)\}$. For integers $0 \leq l \leq u \leq S$, the minimum of $\alpha(y) + \beta(S - y)$ over $l \leq y \leq u$ is attained at one of*

$$y_1 = \min\{u, A(\alpha(l))\}, \quad y_2 = \max\{l, S - B(\beta(S - u))\}.$$

Proof. Take any feasible y of cost t , and put $a = \alpha(y)$, $b = \beta(S - y)$. Then $a + b = t$, $a \geq \alpha(l)$, $b \geq \beta(S - u)$ and $A(a) + B(b) \geq S$. Discrete convexity puts the maximum of $A(z) + B(t - z)$ on the integer interval $\alpha(l) \leq z \leq t - \beta(S - u)$ at an endpoint. At the first endpoint, $A(\alpha(l)) + B(t - \alpha(l)) \geq S$. The point y_1 has first cost at most $\alpha(l)$ and second count at most $B(t - \alpha(l))$: if clipped to u , use $\beta(S - u) \leq t - \alpha(l)$ and monotonicity of B . It therefore has total cost at most t . The other endpoint gives y_2 symmetrically; if clipped to l , use $\alpha(l) \leq t - \beta(S - u)$. Both candidates lie in $[l, u]$. Applying this to a minimizing y proves the assertion, including flat capacities, zero roots and singleton intervals. $\square$

Each C_h in (13) is the maximum of two convex quadratics and is nondecreasing on $\mathbb{N}$. The lemma therefore optimizes sums of any two punctured-quadrant root costs without searching their allocation interval. It also applies to the capacities $s(s + 1)$ and s^2 used in the inverse-profile calculation below.

5 Exact neighborhood profiles on odd rectangles

Let $G = P_{2a+1} \square P_{2b+1}$, where $a \geq b \geq 1$, with coordinates (x, y) ; thus x is the long coordinate. Write V_0, V_1 for its even and odd checkerboard classes, and

$$E = |V_0| = \frac{|G| + 1}{2}, \quad M = |V_1| = E - 1.$$

Define the two open-neighborhood profiles by

$$g_p(k) = \min\{|N(S)| : S \subseteq V_p, |S| = k\}.$$

It is convenient to use the integer functions

$$R(k) = \lceil \sqrt{k} \rceil, \quad Q(k) = \min\{r \geq 1 : k \leq r(r - 1)\}.$$

In particular $R(0) = 0$ and $Q(0) = 1$.

Theorem 5.1 (Odd-rectangle profiles and compatible orders). *Both profiles have $g_p(0) = 0$. For nonempty sets,*


$$g_0(k) = k + \min\{R(k), b, R(E - k) - 1\}, \quad 1 \leq k \leq E, \quad (15)$$

$$g_1(k) = k + \min\{Q(k), b + 1, Q(M - k)\}, \quad 1 \leq k \leq M. \quad (16)$$

Order each parity by increasing $(x + y, x)$, and let $I_p(k)$ be its prefix of size k . Every prefix attains the corresponding minimum, and

$$N(I_p(k)) = I_{1-p}(g_p(k)).$$

The order fills the short direction first within a diagonal. Its direction is significant on an unequal rectangle. For example, the opposite tie order on 7×5 gives the five-element even prefix $\{(0, 0), (0, 2), (1, 1), (2, 0), (4, 0)\}$, with eight neighbors; the set $\{(0, 0), (0, 2), (0, 4), (1, 1), (1, 3)\}$ has seven.



Filled circles: source set. Open circles: its possible neighbors.

Figure 2: Minimizing even prefixes on a 9×7 rectangle. The three regimes are controlled by a corner, the short width, and the complement.

5.1 Quadrant bounds

Apply the odd-coordinate compressions of Lemma 3.4 in both directions. Their common fixed sets are lower closed under decreasing either coordinate by two, and compression does not increase the neighborhood size. We may therefore assume that each coordinate-parity pattern of S is a lower Ferrers diagram. Let e_i and o_i be the selected counts in rows $y = 2i$ and $y = 2i + 1$, respectively. These are separately nonincreasing sequences. For source color 0 they count even- x and odd- x vertices, whereas for source color 1 they count odd- x and even- x vertices.

Lemma 5.2 (Quadrant area and row bounds). *Regard a finite set with this Ferrers property as a subset of the infinite nonnegative quadrant. Put $\delta = |N(S)| - |S|$. If S has color 0, then*

$$e_i > 0 \implies e_i + i \leq \delta, \quad o_i > 0 \implies o_i + i + 1 \leq \delta, \quad |S| \leq \delta^2.$$

If S has color 1, then

$$e_i > 0 \implies e_i + i + 1 \leq \delta, \quad o_i > 0 \implies o_i + i + 1 \leq \delta, \quad |S| \leq \delta(\delta - 1).$$

A nonempty set has surplus at least one in color 0 and at least two in color 1.

Proof. The sequences eventually vanish. A nonempty odd- x prefix of length t has $t + 1$ even- x neighbors in the quadrant; an even- x prefix has t odd- x neighbors. Contributions from neighboring rows are also prefixes, so their unions are given by maxima. Subtracting the source row sizes therefore yields, in color 0,

$$\delta = (o_0 - e_0)_+ + \sum_{i \geq 0} \left[\max\{e_i - o_i, \mathbf{1}_{\{o_i > 0\}}\} + (o_i - e_{i+1})_+ \right]. \quad (17)$$

For color 1 the identity is

$$\delta = \max\{\mathbf{1}_{\{e_0 > 0\}}, o_0 - e_0\} + \sum_{i \geq 0} \left[(e_i - o_i)_+ + \max\{o_i - e_{i+1}, \mathbf{1}_{\{e_{i+1} > 0\}}\} \right]. \quad (18)$$

All terms displayed are nonnegative.

Suppose first that S has color 0 and $e_i > 0$. In (17), each of the first i paired summands is at least one. If its odd count is positive, the first term is at least one; otherwise its even count is positive. The tail from i onward is at least the telescoping sum

$$\sum_{j \geq i} ((e_j - o_j) + (o_j - e_{j+1})) = e_i.$$

Thus $\delta \geq i + e_i$. If $o_i > 0$, every preceding pair contributes at least one, as does the first term of pair i . The second term of pair i and the subsequent pairs have sum at least o_i , by the same telescoping argument. This gives $\delta \geq i + 1 + o_i$.

For color 1 and $e_i > 0$, the initial term of (18) is at least one, and the first i pairs each contribute at least one because $e_{j+1} > 0$ for $j < i$. The remaining tail is at least e_i . If $o_i > 0$, the initial term is again at least one. Every preceding pair contributes at least one: either $e_{j+1} > 0$, or the positive o_j appears in its second term. The tail after the first term of pair i is at least o_i . These are the two claimed bounds for color 1.

Finally sum the row bounds. In color 0 the maximum possible total is

$$\sum_{i \geq 0} \max(\delta - i, 0) + \sum_{i \geq 0} \max(\delta - i - 1, 0) = \delta^2.$$

In color 1 both sums have the second form, giving $\delta(\delta - 1)$. The row inequalities also give the stated positive surplus bounds for nonempty sets. $\square$

5.2 The two far edges

Lemma 5.3 (Corner, strip, or complement). *Let $S \subseteq V_p$ have the Ferrers property. Put $\delta = |N_G(S)| - |S|$, and let F and Z count its vertices on the far edges $x = 2a$ and $y = 2b$, respectively. If exactly one of F, Z is positive, then $\delta \geq b + p$. If neither is positive, the quadrant bounds apply to S without change. If both are positive, $N_G(S)$ covers both near edges in color $1 - p$, and the set*

$$V_{1-p} \setminus N_G(S),$$

reflected in both coordinate midlines, has the Ferrers property and no far-edge vertices.

Proof. Viewing the same S in the quadrant adds exactly one neighbor beyond the rectangle for every selected far-edge vertex. These $F + Z$ vertices are distinct, even when the far corner is selected. Hence

$$\delta_{\text{quadrant}} = \delta + F + Z. \quad (19)$$

If $F = 0$ and $Z > 0$, the last even row has $e_b = Z$. The row bound in Lemma 5.2 gives $\delta + Z \geq Z + b + p$. If $F > 0$ and $Z = 0$, swap the axes to get the stronger bound $\delta \geq a + p$. If $F = Z = 0$, there is no clipping.

Suppose now that both are positive. In color 0, the Ferrers property forces every even- x vertex of row $y = 0$ and every even- y vertex of column $x = 0$ into S . They cover both opposite-color near edges. In color 1, far- x occupancy forces the full even- x row $y = 1$, covering the even- x near edge $y = 0$. Far- y occupancy forces $(1, 2i) \in S$ for every $0 \leq i \leq b$, covering the even- y near edge $x = 0$.

The neighborhood of a fixed set is fixed under the two path operators, so its complement is upper closed within each coordinate-parity pattern. The two reflections convert it to a lower Ferrers set. They preserve checkerboard color because both side lengths are odd. The near-edge coverage just proved means the reflected complement has no far-edge vertices, as required. $\square$

5.3 Lower bounds in Theorem 5.1

Let $S \subseteq V_0$, $|S| = k > 0$, and put $\delta = |N(S)| - k$. If $\delta \geq b$ the desired lower bound is immediate. Otherwise Lemma 5.3 leaves two possibilities. With no clipping, $k \leq \delta^2$. With both far edges occupied, let T be the reflected opposite-color complement of $N(S)$. It has size $t = M - k - \delta$. Before reflection, its neighborhood is contained in $V_0 \setminus S$, so its quadrant surplus is at most $\delta + 1$. For $\delta \geq 1$, Lemma 5.2 gives

$$M - k - \delta = t \leq \delta(\delta + 1), \quad k \geq E - (\delta + 1)^2.$$

If $\delta = -1$ or 0 , T must be empty, since a nonempty quadrant color-1 set has surplus at least two. This yields respectively $k = E$ or $k = M$, the same conclusion. No smaller δ is possible: pairing along each row, then vertically in the unpaired last column, gives a matching covering all vertices except the far even corner, and hence $|N(S)| \geq k - 1$. In all cases,

$$\delta \geq \min\{R(k), b, R(E - k) - 1\}.$$

For $S \subseteq V_1$, the same matching gives $\delta \geq 0$. Below $b + 1$, the unclipped case gives $k \leq \delta(\delta - 1)$. In the complement case, T has size $E - k - \delta$ and quadrant color-0 surplus at most $\delta - 1$. The case $\delta = 0$ is impossible: T is nonempty because $k \leq M < E$, but would have negative surplus. For $\delta \geq 1$,

$$E - k - \delta \leq (\delta - 1)^2, \quad k \geq M - \delta(\delta - 1).$$

The unclipped nonempty case also has $\delta \geq 1$. Thus

$$\delta \geq \min\{Q(k), b + 1, Q(M - k)\}.$$

These prove the lower bounds in both formulas.

For comparison, the two exact profiles of any bipartite graph obey the complement-inverse identity

$$g_1(y) = E - \max\{x : g_0(x) \leq M - y\}. \quad (20)$$

Indeed, the existence of nonadjacent sets of sizes x and y in the two colors is equivalent to $g_0(x) \leq M - y$, and also to $g_1(y) \leq E - x$. Thus the second formula can alternatively be checked from the first by exact inversion.

5.4 Attainment and neighborhood nesting

The vertices on diagonal $x + y = d$ have x in the interval

$$L_d = \max(0, d - 2b) \leq x \leq \min(2a, d),$$

whose length is

$$\ell(d) = \min(2a, d) - \max(0, d - 2b) + 1, \quad 0 \leq d \leq 2a + 2b.$$

A nonempty prefix ends after j vertices of diagonal d , for some $1 \leq j \leq \ell(d)$, and has size

$$k = \sum_{\substack{s < d \\ s \equiv d \pmod{2}}} \ell(s) + j. \quad (21)$$

All earlier opposite-color diagonals belong to its neighborhood. On diagonal $d + 1$, each selected x contributes x or $x + 1$ when the corresponding step stays inside the board. Their union is the initial interval there, of length

$$f(d, j) = j + 1 - \mathbf{1}_{\{d \geq 2b\}} - \mathbf{1}_{\{L_{d+j-1} = 2a\}}. \quad (22)$$

The first indicator removes the unavailable step past the short-direction boundary, and the second removes the step past the long-direction boundary. At the far corner both are present and $f = 0$. The initial minority diagonal $d = 1$ additionally covers the corner on diagonal zero, already included among the earlier opposite diagonals. Consequently

$$|N(I_p(k))| = \sum_{\substack{s < d \\ s \not\equiv d \pmod{2}}} \ell(s) + f(d, j),$$

and this neighborhood is exactly a prefix of the other color.

To evaluate its surplus, let $\Delta(d)$ be the earlier opposite-color sum minus the earlier same-color sum. The recurrence $\Delta(0) = 0$, $\Delta(d + 1) = \ell(d) - \Delta(d)$ gives

$$\Delta(d) = \begin{cases} \lceil d/2 \rceil, & d \leq 2b, \\ b + (d \bmod 2), & 2b \leq d \leq 2a, \\ a + b - \lfloor d/2 \rfloor + (d \bmod 2), & d \geq 2a. \end{cases}$$

The formulas agree at their shared endpoints. The exact surplus is $\Delta(d) + 1$ minus the two indicators in (22). Complete even diagonals through sum $2r - 2$ contain r^2 vertices, and complete odd diagonals through sum $2r - 1$ contain $r(r + 1)$. Substitution into (21) yields

Color	Size range	Prefix surplus
0	$1 \leq k \leq b^2$	$R(k)$
0	$b^2 < k < E - b^2$	b
0	$E - b^2 \leq k \leq E$	$R(E - k) - 1$
1	$1 \leq k \leq b(b + 1)$	$Q(k)$
1	$b(b + 1) < k < M - b(b + 1)$	$b + 1$
1	$M - b(b + 1) \leq k \leq M$	$Q(M - k)$

For a square, the shared minority endpoint has surplus $b + 1$ in both rows; empty ranges are omitted. The table equals the two lower bounds. It proves attainment and the compatible-neighborhood assertion, completing the proof of Theorem 5.1.

5.5 An exact two-count game

Corollary 5.4. *For every daily budget on every odd-by-odd rectangle, the unrestricted minimum capture time is exactly the shortest-path distance from (E, M) to $(0, 0)$ in the following two-count state graph:*

$$(u, v) \longrightarrow (g_1(j), g_0(i)), \quad 0 \leq i \leq u, \quad 0 \leq j \leq v, \quad u + v - i - j \leq m.$$

A shortest path gives an optimal physical strategy by deleting suffixes of the two current-color prefix orders. The reduction also respects arbitrary prescribed daily budgets.

Proof. Define

$$C(S) = I_0(|S \cap V_0|) \cup I_1(|S \cap V_1|).$$

It is monotone, cardinality preserving and idempotent. By the exact profiles and nesting,

$$N(C(S)) \cap V_{1-p} = I_{1-p}(g_p(|S \cap V_p|)) \subseteq I_{1-p}(|N(S \cap V_p)|).$$

The last set is the corresponding part of $C(N(S))$. Apply Theorem 3.2. Every displayed transition is also the exact neighborhood of the indicated survivors, so it has a physical realization. Shortest paths, with infinity for an unreachable zero state, settle both feasibility and minimum time. $\square$

The corollary is an exact recurrence for every parameter, rather than a closed expression for its shortest-path distance. It avoids imposing an unproved sequential allocation to the two initial colors: both may receive inspections on any day.

6 Every budget on odd rectangles: a one-day determination

Retain $E = M + 1$ and the exact profiles of Theorem 5.1. The following reduction has no restriction on width, length, or feasible budget. It reduces the answer to one binary choice. Section 7 resolves that choice by an exact bounded-state calculation. The same section proves a scalar criterion after an explicit onset at every budget; only its remaining short-board range is open.

Write $\iota_p(z) = \max\{k : g_p(k) \leq z\}$ for the inverse profiles. Define $D_0(0) = D_1(0) = 0$ and

$$D_0(t+1) = \iota_0(\min\{M, D_1(t) + m\}), \quad D_1(t+1) = \iota_1(\min\{E, D_0(t) + m\}). \quad (23)$$

Here $D_p(t)$ is the largest possible deficit in current color p after t inspection-and-movement steps devoted entirely to that cohort. Compatible prefixes attain these values. Let τ be the first t for which $D_0(t) = E$ or $D_1(t) = M$.

Theorem 6.1 (Uniform one-day determination). *For every feasible budget on every odd rectangle,*

$$2\tau - 1 \leq T_m \leq 2\tau.$$

Moreover $T_m = 2\tau - 1$ whenever

$$E + M - D_0(\tau - 1) - D_1(\tau - 1) \leq m. \quad (24)$$

Necessity follows after the explicit onset in Theorem 7.5, as well as in the all-length minimum-budget and high-budget ranges proved later. It remains open outside those ranges. The exact two-count algorithm determines the answer in all cases. The solo calculation itself uses at most $4b + 2$ arithmetic stages, as shown below.

6.1 Concentrating the deficit before the first possible capture

Put $q(z) = Q(z) - 1$. Complement duality gives

$$\begin{aligned} \iota_0(z) &= z - A(z), & A(z) &= \min\{q(z), b, q(M - z)\}, & 0 \leq z < M, \\ \iota_0(M) &= E, \\ \iota_1(z) &= z - B(z), & B(z) &= \min\{R(z), b + 1, 1 + R(E - z)\}, & 0 \leq z \leq E. \end{aligned}$$

In particular $\iota_0(0) = \iota_1(0) = \iota_1(1) = 0$.

Lemma 6.2 (Inverse concentration). *For $x, y \geq 0$ with $x + y \leq M$,*

$$\iota_0(x) + \iota_1(y) \leq \iota_0(x + y).$$

If $x \geq 2$, also $\iota_1(x) + \iota_0(y) \leq \iota_1(x + y)$.

Proof. The function q is subadditive: $u \leq r(r + 1)$ and $v \leq s(s + 1)$ imply $u + v \leq (r + s)(r + s + 1)$. The minimum A remains subadditive below M . Indeed, if either minimizing branch is its cap or its decreasing reflected branch, that branch alone bounds $A(x + y)$; otherwise use subadditivity of q . Each branch of $B(y)$ dominates the corresponding branch of $A(y)$. Hence $A(x) + B(y) \geq A(x + y)$, proving the first inequality below M . At $x + y = M$, each proper inverse is at most its argument, whereas $\iota_0(M) = M + 1$; the endpoint $x = M, y = 0$ is equality.

For the second inequality, $x \geq 2$ and $y < M$. If the cap or reflected branch minimizes $B(x)$, then $B(x + y) \leq B(x)$. Otherwise $B(x) = R(x) \geq 2$. If $A(y) = q(y)$, writing $r = R(x)$ and $s = q(y)$ gives $x + y \leq r^2 + s(s + 1) \leq (r + s)^2$. If $A(y) = b$, use $R(x) + b \geq b + 1 \geq B(x + y)$. Finally, if $A(y) = q(M - y)$, put $z = M - y \geq x$. Since $R(z) \leq Q(z)$,

$$B(x + y) \leq 1 + R(z + 1 - x) \leq 1 + Q(z) \leq R(x) + Q(z) - 1.$$

These are exactly the required cost inequalities. □

Lemma 6.3 (Exact total deficit before τ). *Let $d_p(t)$ be the deficit in current color p for an arbitrary physical strategy after t inspections and moves. For $t < \tau$,*

$$d_0(t) + d_1(t) \leq \max\{D_0(t), D_1(t)\}.$$

Equality is attainable by devoting the budget to one cohort.

Proof. Always $d_p(t) \leq D_p(t)$ by the neighborhood lower bounds and monotonicity. Induct on t . When $t + 1 < \tau$, the endpoint values of the inverses imply

$$\max\{D_0(t), D_1(t)\} + m \leq M. \tag{25}$$

For useful quotas p, q with $p + q \leq m$, put $X = d_0(t) + p$, $Y = d_1(t) + q$. Complement duality bounds the new deficits by $\iota_0(Y)$, $\iota_1(X)$, and $X + Y \leq \max D_p(t) + m \leq M$. If $D_1(t) \geq D_0(t)$, the first inverse inequality bounds their sum by $\iota_0(D_1(t) + m) = D_0(t + 1)$. Otherwise, when $X \geq 2$, the second bounds it by $\iota_1(D_0(t) + m) = D_1(t + 1)$. When $X \leq 1$, the latter inverse vanishes and the individual bound on $d_1(t)$ suffices. The induction begins with zero deficits. □

6.2 The midpoint argument and attaining searches

Proof of Theorem 6.1. For $\tau = 1$ the lower bound is immediate. Otherwise put $L = \tau - 1$ and suppose a schedule wins in $2L$ days, padding a shorter schedule by empty inspections if necessary. Its first L inspections, each followed by movement, leave a belief F on day $L + 1$. The preceding lemma gives $|F| \geq E$.

Read the last L inspections backward along the undirected edges. Use $L - 1$ inspection-and-movement steps and then one inspection without movement. The resulting set R is also on day $L + 1$. Its deficit is at most $\max D_p(L - 1) + m \leq M$, by (25) with the noncapturing time L . Thus $|R| \geq E$. Since $|G| = 2E - 1$, the sets intersect. Concatenating their forward and backward avoiding walks contradicts success. Hence $T_m \geq 2\tau - 1$.

If $A_1, \dots, A_\tau$ captures one initial parity, use

$$A_1, \dots, A_\tau, A_\tau, \dots, A_1.$$

The first half captures that parity. Reversing an avoiding walk in the second half gives a walk of that initial parity avoiding the first half, so the opposite cohort is captured as well. This proves the upper bound with an explicit legal schedule.

Finally prepare one physical color optimally for L steps before a central inspection, and the other physical color for L reversed steps after it. Each half leaves the other cohort full. Their intersection therefore consists of the two solo residual sets, of total size $E + M - D_0(L) - D_1(L)$. Inspect that intersection centrally. Condition (24) makes this a legal $(2\tau - 1)$ -day search. $\square$

6.3 Evaluating the solo recurrence without iterating over days

Put $C_0 = E, C_1 = M$, $f_p(k) = g_p((k - m)_+)$, and $P_p = f_{1-p} \circ f_p$. The two-day map returns a cohort to its original physical color. Extend $\iota_p(z) = C_p$ for $z \geq C_{1-p}$.

Proposition 6.4 (A bounded number of arithmetic stages). *For any canonical one-cohort prefix, its exact solo capture time and its remaining count at any prescribed day can be evaluated using at most $4b + 2$ translation intervals and integer floor divisions. The number of stages is independent of the longer side and the number of search days. This evaluates the solo part of Theorem 6.1; it does not decide the unresolved scalar midpoint condition.*

Proof. Every change in the surplus $g_p(k) - k$ at a positive argument belongs to

$$\begin{aligned} \mathcal{B}_0 &= \{r^2 + 1 : 0 \leq r < b\} \cup \{E - r^2 : 0 \leq r \leq b\}, \\ \mathcal{B}_1 &= \{r(r - 1) + 1 : 1 \leq r \leq b\} \cup \{M - r(r - 1) : 1 \leq r \leq b\}. \end{aligned}$$

Include 1 and discard points outside the appropriate profile domain. These are the change points of the two capped root terms; their minimum cannot change elsewhere. Set $c_p = \min\{C_p, m + \iota_p(m)\}$. Exactly the counts $k \leq c_p$ disappear in at most two days. Above c_p , the change points of P_p are contained in

$$\{m + s : s \in \mathcal{B}_p\} \cup \{m + \iota_p(m + s - 1) + 1 : s \in \mathcal{B}_{1-p}\}.$$

The second set is the inverse image of the first argument at which the second surplus changes. Add $c_p + 1, C_p + 1$ and clip to this interval. Between consecutive endpoints $l, u + 1$, the map is $k \mapsto k - d$, where $d \geq 2m - 2b - 1 \geq 1$ and every output is positive. There are at most $4b + 2$ such intervals.

Process them from top to bottom. If the current count $k \geq l$, exactly $1 + \lfloor (k - l)/d \rfloor$ consecutive pairs have their source in $[l, u]$. Subtract this multiple of d and add twice that many days. The next count is below l , so no interval is revisited. On reaching $k \leq c_p$, add zero, one, or two final days according as $k = 0$, $0 < k \leq m$, or $k > m$. For a prescribed day, truncate the appropriate number of pairs and use one application of f_p if needed. All operations are exact integer arithmetic. Sorting the endpoints uses $O(b \log b)$ comparisons; the number of division stages is $O(b)$. $\square$

Lemma 6.5 (Two candidates for opposed roots). *Write $q(y) = Q(y) - 1$, with $q(0) = 0$. For integers $0 \leq l \leq u \leq S$, the minimum of $q(y) + R(S - y)$ over $l \leq y \leq u$ is attained at one of*

$$\min\{u, q(l)(q(l) + 1)\}, \quad \max\{l, S - R(S - u)^2\}.$$

Proof. Apply Lemma 4.2 to $A(s) = s(s + 1)$ and $B(s) = s^2$, whose inverse capacities are q and R . It includes zero roots and singleton intervals; the two candidates need not be the original endpoints l, u . The module `RootConvolution` verifies the exact attaining minimum with internally defined integer roots. $\square$

6.4 Exact transfer through the affine middle

This second reduction preserves joint states and does not assume that optimal searches finish one cohort before starting the other. Set $K = b^2 + 1$, $L = K + m$, $U = E - K$. Both profiles have constant surplus on every residual obtained by allocating at most m inspections to a count in $[L, U]$. Track the initial cohorts, the first currently in color $z \in \{0, 1\}$, and put $B_z(t) = bt + \lfloor (t + z)/2 \rfloor$.

Theorem 6.6 (Exact middle transfer). *Assume $m \geq b + 1$ and $(a, d), (a', d') \in [L, U]^2$. A t -day full-budget trajectory remaining in this band joins the two states if and only if*

$$a' + d' = a + d + (2b + 1 - m)t, \quad 0 \leq a + B_z(t) - a' \leq mt.$$

The statement includes a construction of all daily allocations.

Proof. Write $S_j = a + d + (2b + 1 - m)j$. The attainable first counts at time j form exactly the integer interval

$$[\max\{L, S_j - U, a + B_z(j) - mj\}, \min\{U, S_j - L, a + B_z(j)\}].$$

For one step the first count changes by $b + (z + j \bmod 2) - p$, $0 \leq p \leq m$. Taking the union over this integer interval and intersecting the two band constraints gives the displayed interval at $j + 1$: both alternating surpluses lie between zero and m . The intervals are nonempty whenever $2L \leq S_j \leq 2U$; pairwise comparison of their three lower and upper bounds proves this directly. Since S_j is affine, the initial and final band conditions ensure these inequalities at every intermediate time. For a prescribed next count x' , choose the preceding count to be the larger of the current lower endpoint and $x' - b - (z + j \bmod 2)$. It lies in the current interval and gives an allocation between zero and m . Backtracking constructs the trajectory. Necessity follows by summing its allocations. $\square$

The Lean module `MiddleIntervalTransfer` verifies this interval statement, including constructive sufficiency. The identification of the band with physical rectangle profiles remains an ordinary proof. Replacing maximal interior excursions by these transfers gives an exact boundary graph with $O((b^2 + m)E)$ retained states: retain states outside $[L, U]^2$ and a collar of width $m + 1$ inside its boundary. Every one-day entry or exit crosses that collar, and the theorem expands each added transfer back into legal moves. This reduction does not itself evaluate the remaining boundary optimization in closed form.

7 A bounded joint-state calculation for every odd rectangle

The one-day ambiguity in Theorem 6.1 can be resolved without exploring a state space that grows with the longer side. The reduction below retains the exceptional mixed histories. We then prove that, after an explicit width- and budget-dependent length threshold, the two solo endpoints alone decide the answer. Scalar necessity for all shorter rectangles at intermediate budgets remains open.

Fix odd $n \geq w = 2b + 1 \geq 3$ and a feasible budget $m \geq b + 1$. Retain $E = M + 1$, the inverse profiles ι_p , and the solo deficits $D_p(t)$. Define the following constants, depending only on w, m :

$$\begin{aligned} d &= 2m - w \geq 1, & H &= b^2 + 1, & \Gamma &= m - 1, \\ K &= H - 1 + 2m(\lfloor \Gamma/w \rfloor + 1), & W &= \Gamma + K + 1, \\ J &= 2K + \Gamma + 2m + H + 2. \end{aligned} \tag{26}$$

Theorem 7.1 (Exact evaluation with a length-independent state bound). *The minimum capture time on $P_w \square P_n$ can be evaluated using at most*

$$F = 2 + 2\Gamma(K + 1)$$

retained joint states and at most

$$4 \lceil J/d \rceil + W + 8$$

ordinary state updates, together with integer floor divisions and a final optimization over pairs of retained states. Each update considers at most $m + 1$ quota splits per state. All bounds are independent of n . The calculation determines the exact choice between $2\tau - 1$ and 2τ and gives a winning strategy through the compatible prefixes. The bit lengths of the arithmetic inputs still depend on $\log n$.

The endpoint cases are immediate: one day suffices exactly when $m \geq E + M$, and $M \leq m < E + M$ gives two days by inspecting one entire color and then the surviving other cohort. We therefore discuss $m < M$.

7.1 The exact central optimization

Write $A(t) = \max D_p(t)$ and $B(t) = \min D_p(t)$. A deficit pair is *exceptional* at time $t < \tau$ if its total exceeds $B(t)$. At each layer retain the two pure solo pairs $(D_0(t), 0), (0, D_1(t))$ and the attainable exceptional mixed pairs, optionally discarding pairs dominated componentwise. Their exact transition, for quota p in current color zero, is

$$(x, y) \mapsto (\iota_0(y + m - p), \iota_1(x + p)), \quad 0 \leq p \leq m. \tag{27}$$

The noncapture bound (25) ensures the inputs are proper whenever the successor time is below τ .

Lemma 7.2 (Discarding low-total mixed states). *Iterating (27), retaining only the types just described, gives the exact minimum central inspection requirement*

$$C_L = \min_{(x,y),(u,v) \in \mathcal{R}_L} ((E - x - u)_+ + (M - y - v)_+), \quad L = \tau - 1. \tag{28}$$

Here $\mathcal{R}_L$ is the retained frontier. The answer is $2\tau - 1$ exactly when $C_L \leq m$, and is 2τ otherwise.

Proof. First, a mixed exceptional successor cannot come from a pair whose total is at most $B(t)$. Apply both inequalities of Lemma 6.2 to the common total input, at most $B(t) + m$. The second applies because a positive new minority deficit requires its inverse argument to be at least two. The resulting total is at most both next solo capacities. This is the persistence property used below.

By induction every attainable exceptional pair is dominated by a retained attainable pair: mixed exceptional successors have exceptional predecessors, and pure successors are dominated by the corresponding solo endpoint. The transition is monotone. Conversely every retained pair is attained by the exact prefix recurrence of Corollary 5.4. A discarded pair has total at most B , while every other pair has total at most A , by Lemma 6.3. Its capped union with any other pair is therefore at most $A + B = D_0 + D_1$, already attained by the two opposite solo endpoints. It cannot improve the central cut.

For completeness, cut a putative $(2L + 1)$ -day winning schedule at its central inspection. The first L inspections and moves give a forward belief; the last L , read backward along the undirected edges, give a reverse belief on that same day. Their intersection must be inspected. In color p its size is at least the positive part of the color size minus the two deficits. Canonical domination and the preceding retention argument give the lower bound (28). Conversely realize one retained pair by prefixes and the other by reflecting both coordinates of its prefix construction. On an odd rectangle the reflection preserves colors and reverses their orders. The two beliefs therefore intersect in exactly the displayed positive parts. Inspect that intersection centrally and reverse the second half-schedule. This attains the bound whenever it is at most m . The uniform one-day theorem gives the alternative 2τ . $\square$

Lemma 7.3 (The ancestry of an exceptional state). *Every retained exceptional mixed pair has a realization whose most recent pure ancestor belongs to the currently faster initial cohort. All such pairs at a fixed time therefore have the same primary ancestry. The secondary deficit means the coordinate of the other initial cohort; it need not a priori be the numerically smaller coordinate.*

Proof. Replace every pure successor by its solo endpoint, and discard low-total mixed successors by Lemma 7.2. Label a mixed history by its most recent pure ancestor. A slower pure endpoint has total at most B and cannot produce an exceptional mixed successor.

Suppose an exceptional source has primary physical color zero. Its total s is at most $D_0(t)$. The second inverse concentration inequality bounds a mixed successor's total by $\iota_1(s + m) \leq \iota_1(D_0(t) + m) = D_1(t + 1)$. If the successor is exceptional, color one must consequently be strictly faster. With primary color one, the first concentration inequality instead bounds its total by $D_0(t + 1)$, forcing color zero to be faster. The primary color flips in both cases, exactly as its initial cohort does under movement. Proper inputs hold before τ ; positive minority output supplies the extra hypothesis of the second concentration inequality. Persistence excludes low-total mixed predecessors. This proves the assertion inductively, including rebirth from a pure endpoint. At a solo tie no exceptional pair exists. $\square$

7.2 Only bounded endpoint collars can be exceptional

Lemma 7.4 (Uniform collar bound). *For $t < \tau$, $|D_0(t) - D_1(t)| \leq \Gamma$. Every attainable exceptional mixed pair satisfies $\min(x, y) \leq K$.*

Proof. The solo recurrence is monotone from its zero initial pair. Write its proper step as $(a', c') = (\iota_0(c + m), \iota_1(a + m))$. Every input is at least $m \geq 2$. The three deficit branches of ι_0 are all at

least one, and those of ι_1 are all at least two. Thus $a' \leq c + m - 1$ and $c' \leq a + m - 2$. Since $a' \geq a$ and $c' \geq c$,

$$a' - c' \leq m - 1, \quad c' - a' \leq m - 2.$$

The initial tie satisfies the symmetric bound as well. This argument includes the final precapture layer; no full inverse endpoint is used.

Put $s = x + y$. If $x, y \geq H$ at a source layer with a precapture successor, the inverse inputs X, Y satisfy $X, Y \geq H$ and $X + Y = s + m \leq M$. These inequalities put every lower and reflected root argument beyond its corner: $q(Y), q(M - Y) \geq b$, $R(X) \geq b + 1$, and $1 + R(E - X) \geq b + 1$. Thus

$$\iota_0(Y) = Y - b, \quad \iota_1(X) = X - (b + 1), \quad s' = s + m - w.$$

Each solo capacity gains at least $d = 2m - w$ over two proper steps, since its inverse deficit is at most b or $b + 1$. Consequently two successive bulk source states give

$$(A - s)(t + 2) \geq (A - s)(t) + w.$$

An exceptional state has $0 \leq A - s < A - B \leq \Gamma$. A consecutive run of bulk exceptional states therefore has at most $2 \lfloor \Gamma/w \rfloor + 1$ transitions.

For a target exceptional mixed state, trace backward to the most recent state with smaller coordinate below H ; such a state exists at time zero. Every intervening bulk state is mixed and exceptional by the persistence property in Lemma 7.2. At entry its smaller coordinate is at most $H - 1 + m$. Each subsequent step increases the smaller coordinate by at most m , because each proper inverse is at most its argument. The run-length bound gives precisely $H - 1 + 2m(\lfloor \Gamma/w \rfloor + 1) = K$. $\square$

There are at most Γ integer totals strictly between B and A including the upper endpoint. For each, at most $2(K + 1)$ pairs have a coordinate at most K . Adding the two pure endpoints proves the state bound F in Theorem 7.1.

7.3 The long middle forgets its initial mixed histories

Call a layer *safe* when $J \leq D_0, D_1 \leq M - J$. These layers form an interval. At a safe layer encode a retained mixed pair by (j, k, ℓ) : j is its larger coordinate's color, k its smaller deficit, and $\ell = D_j - d_j$ its loss relative to the solo capacity in that color. Then

$$0 \leq k \leq K, \quad 0 \leq \ell \leq \Gamma + K.$$

The primary coordinate is greater than K , so its color is unambiguous. Define the bottom inverses

$$h_0(z) = z - \min\{q(z), b\}, \quad h_1(z) = z - \min\{R(z), b + 1\}.$$

If r of the inspections are devoted to the smaller-deficit cohort, the exact normalized transition is

$$(j, k, \ell) \mapsto (1 - j, h_j(k + r), \ell + r), \quad 0 \leq r \leq m. \quad (29)$$

The primary input and the corresponding solo input lie in the same affine profile band; their difference therefore increases by r . The smaller input is at most $K + m$ and uses only the bottom inverse. To check all margins explicitly, the primary input is between H and $M - H$, its output remains greater than K , and safety implies $M \geq 2J$, excluding the reflected branch for an input at

most $K + m$. The definition of J ensures these inequalities even on the last safe source step. If its successor remains mixed and exceptional, the collar lemma forces its smaller output to be at most K .

The solo transition in this region is

$$(D'_0, D'_1) = (D_1 + m - b, D_0 + m - b - 1), \quad D_p(t + 2) = D_p(t) + d.$$

Thus its color difference and the exceptional-state filter $\ell - k < D_j - \min D_p$ are periodic with period two. Moreover, whenever the smaller output is positive, $h_j(z) \leq z - 1$, and (29) gives

$$\ell' - k' \geq \ell - k + 1.$$

At entry $\ell - k \geq -K$, whereas exceptionality requires $\ell - k < \Gamma$. An initial mixed history cannot persist for $\Gamma + K$ safe transitions. Becoming pure resets its relevant history to the dominating solo endpoint; becoming low-total erases its future relevance until a pure state is reached. A new mixed history born from a solo endpoint starts from $k = \ell = 0$ and has bounded age by the same inequality. After $W = \Gamma + K + 1$ safe transitions, every relevant mixed history comes from a recent pure endpoint. Its rules, births and filters depend only on the two-day phase. Therefore the normalized retained frontier is exactly periodic with period two. Pareto pruning preserves this statement: dominance within a primary color depends only on (k, ℓ) , and opposite primary colors are incomparable.

Completion of Theorem 7.1. Iterate the retained recurrence until capture or until both solo capacities reach J . This takes at most $2 \lceil J/d \rceil + 1$ updates, because every proper pair of steps raises both capacities by at least d . If a safe interval is entered at t_0 , its last layer is explicit. For $p = 0, 1$, let $A_p = \max D_j(t_0 + p)$; omit a parity with $A_p > M - J$. Its final safe layer is

$$t_0 + p + 2 \lfloor (M - J - A_p)/d \rfloor.$$

The larger of these is the last safe layer. Perform at most $W + 2$ updates in the safe interval. If it is longer, stabilization permits an even jump to its last matching parity: for each skipped pair, add d to both solo capacities and to every retained primary coordinate, leaving smaller coordinates unchanged. Complete the at most one remaining safe transition normally.

After leaving the safe interval, some solo capacity exceeds $M - J$; that cohort has at most J rooms remaining, including the majority's extra room. At most $2 \lceil J/d \rceil + 2$ further steps give capture. If the safe interval was empty when the lower threshold was reached, the same upper-collar bound applies. Stop one layer before the first solo capture and use (28). The generous displayed update bound covers the two collars, warm-up, and parity endpoints.

Every retained history is attainable. At an accelerated layer recover its bounded recent history from a pure solo endpoint, realize that endpoint by the solo prefix construction, and then follow the recorded quota suffix. The central-cut construction in Lemma 7.2, or the solo palindrome when the cut is too expensive, supplies the winning strategy. This proves both exact evaluation and construction without expanding the long affine middle one day at a time. $\square$

This is an ordinary geometric and arithmetic proof. It uses the proved prefix reduction but has not been formalized as a complete Lean theorem.

7.4 An explicit onset for the scalar formula

Theorem 7.5 (Eventually only the solo endpoints matter). *With the constants in (26), put*

$$M_* = 2J + \Gamma + m(W + 3).$$

For every odd $n \geq w$ with $M = (wn - 1)/2 \geq M_$,*

$$T_m(P_w \square P_n) = \begin{cases} 2\tau - 1, & E + M - D_0(\tau - 1) - D_1(\tau - 1) \leq m, \\ 2\tau, & E + M - D_0(\tau - 1) - D_1(\tau - 1) > m. \end{cases} \quad (30)$$

Thus Proposition 6.4 evaluates the exact answer using $O(b)$ scalar arithmetic stages; no joint-state optimization is needed.

Proof. The maximum solo deficit increases by at most m per proper step, because each inverse is at most its argument. This concerns the maximum, not each fixed physical coordinate. At the first layer t_0 with both capacities at least J , their maximum is at most $J + \Gamma + m - 1$. This layer precedes capture: a predecessor with smaller capacity below J has maximum at most $J + \Gamma - 1$, and adding m still leaves both inverse inputs below the full endpoint. Monotonicity and eventual capture ensure that t_0 exists. For $0 \leq s \leq W + 1$ the maximum is at most $J + \Gamma + m - 1 + ms \leq M - J$, while the minimum remains at least J . Induction using the same proper-input bound excludes capture during these layers. There are therefore W consecutive safe transitions.

The age argument above erases every mixed history present at safe entry. After W transitions every retained mixed state has a pure ancestor within the safe interval. Its ancestry-primary is then the larger coordinate, greater than K , and its ancestry-secondary is at most K . This alignment is essential: Lemma 7.4 alone bounds only the numerical minimum.

The alignment persists through the final collar. Following an initial cohort across movement, a secondary deficit at most K becomes at most $K + m$, since this input is proper. If the new primary were at most K , the total would be at most $2K + m < J \leq B$, so the successor could not be exceptional. Thus an exceptional successor still has primary greater than K ; the numerical collar bound forces secondary at most K again. A fresh mixed birth from a pure endpoint obeys the same argument, with secondary output at most m . Pure states reset to solo endpoints, and persistence excludes revival of a discarded mixed history.

At the midpoint all exceptional states consequently share a primary ancestry by Lemma 7.3, and each has secondary at most K . A pair of exceptional states leaves at least $M - 2K > m$ rooms in the other color and cannot win centrally. If at least one state is nonexceptional, the two totals sum to at most $A + B = D_0 + D_1$; their central cost is at least $E + M - D_0 - D_1$. The two opposite solo endpoints attain precisely this latter cost. Lemma 7.2 now gives (30). $\square$

The threshold is deliberately generous. Theorem 7.1 still gives an exact bounded calculation below it. Neither this theorem nor the ancestry lemma asserts serial optimality from arbitrary partial states. The all-length results at minimum budget and at $m \geq b^2 + 1$ (Theorems 21.1 and 8.1) leave only intermediate budgets $b + 2 \leq m \leq b^2$ for the unrestricted scalar conjecture. The constants are nondecreasing in m , and $M_*(b, b^2) = O(b^5)$. Hence any counterexample at a fixed width must have $n = O(b^4)$; this is a finite obstruction bound, not a claim that those remaining cases satisfy the scalar criterion.

7.5 Minimum total probes in the unbounded corner model

The next lemma isolates a second scalar reduction valid at every feasible budget. Its conclusion concerns the number of probes, not a fixed horizon in the full finite rectangle. For $k > 0$ define

$$L_0(k) = k + \min\{R(k), b\}, \quad L_1(k) = k + \min\{Q(k), b + 1\}, \quad L_p(0) = 0.$$

Their inverses are the bottom maps h_0, h_1 above. Each is nondecreasing and 1-Lipschitz on the nonnegative integers, vanishes at zero, and has image all nonnegative integers.

Proposition 7.6 (Exact corner ammunition recurrence). *Let $\mu_p(k)$ be the least total quota in an alternating growth sequence $v' = h_p(v+r)$, $0 \leq r \leq m$, starting from zero in either phase and ending at least k in color p . The horizon is unrestricted. Then $\mu_p(0) = 0$, $\mu_p(k) - k$ is nondecreasing, exact targets are attainable, and*

$$\mu_p(k) = \min\{m, L_p(k)\} + \mu_{1-p}((L_p(k) - m)_+). \quad (31)$$

Every two positive recursive calls strictly decrease their argument. The minimum growth time is $\lceil \mu_p(k)/m \rceil$. At any prescribed horizon t , the target is feasible exactly when $\mu_p(k) \leq mt$; its minimum ammunition then remains $\mu_p(k)$. The initial zero phase is free here. With a fixed initial phase, the horizon and final phase must have the corresponding cyclic compatibility. For inspection-before-movement clearing with profiles L_p , the minimum total quota is

$$F_p(k) = \begin{cases} k, & k \leq m, \\ m + \mu_p(k - m), & k > m, \end{cases} \quad F_p(k) = m + F_{1-p}(L_p(k - m)) \quad (k > m).$$

Thus greedy full quotas minimize total probes in this unbounded model.

Proof. Feasibility follows since $h_0(z + m) \geq z + 1$ and $h_1(z + m) \geq z$. In a fixed growth schedule, deleting one probe changes the final output by at most one, by monotonicity and the Lipschitz bound. Deleting probes from an optimal schedule until its output first hits $\ell < k$ removes at least $k - \ell$ probes. The remaining cost is at least $\mu_p(\ell)$, proving monotone excess; the same deletion argument gives optimal schedules with exact output.

Put $z = L_p(k)$, the least input with $h_p(z) \geq k$. A last-step predecessor must be at least $v_0 = (z - m)_+$. For $v_0 \leq v \leq z$, monotone excess gives

$$\mu_{1-p}(v) + z - v \geq \mu_{1-p}(v_0) + \min\{m, z\}.$$

For $v \geq z$ the same bound follows from monotonicity of the excess. Attain v_0 optimally and use quota $\min(m, z)$; since $h_p(z) = k$, equality is attained. This proves (31). Its predecessor is strictly smaller than k in color zero and at most k in color one, so two calls strictly descend.

Every nonterminal positive call contributes m and the terminal one contributes an integer from one to m . Reading backward gives one possibly partial first quota followed by full quotas, in exactly $\lceil \mu_p(k)/m \rceil$ steps. A t -step schedule spends at most mt , proving the horizon lower bound. When t is larger, prepend zero steps: all maps fix zero, and choosing the initial phase to end in p aligns the optimal block correctly. For a fixed initial phase this requires the stated compatibility. Padding at the end of a positive block is not used. In particular,

$$D_p^{\text{bottom}}(t) = \max\{k : \mu_p(k) \leq mt\}.$$

Finally the inverse threshold identity, applied backward through the nonterminal moves of a clearing schedule, gives

$$F_p(k) = \min_{0 \leq r \leq \min(m, k)} (r + \mu_p(k - r)).$$

Monotone excess makes the largest allowed first quota optimal, giving the displayed formula. A full first quota followed by the optimal continuation gives the equivalent clearing recurrence. $\square$

Remark 7.7 (Cyclic growth controls). The proof requires no special root identities. Let the phases form any nonempty finite cycle, let $m \geq 1$, and let each $h_p : \mathbb{N} \rightarrow \mathbb{N}$ be nondecreasing and onto. These assumptions imply $h_p(0) = 0$ and increments in $\{0, 1\}$. With least-input inverse L_p , a target (p, k) is feasible exactly when the reverse threshold orbit

$$(p, k) \mapsto (p - 1, (L_p(k) - m)_+)$$

reaches zero. Necessity follows by propagating required predecessors backward through any actual schedule; sufficiency assigns the least predecessor and quota $\min(m, L_p(k))$ at each reverse step. For feasible targets the same deletion and last-step proof gives (31), with phase subtraction taken cyclically, and the same exact cost and horizon conclusions. Intermediate reverse targets may increase; termination, rather than per-step decrease, is the essential condition. The bottom root maps supply a simple terminating instance. With a fixed initial phase a , the prescribed-horizon assertion requires $a + t = p$ in the phase cycle.

This proposition omits the reflected end branches by definition. It does not by itself justify allocating full quotas serially in the two-cohort rectangle problem. The following conditional exchange states precisely where it does apply to the finite board.

Lemma 7.8 (Suffix exchange beyond the primary lower corner). *Consider a legal history up to $t < \tau$, with a pure solo endpoint for an initial cohort A at time $i \leq t$. Suppose that every subsequent inverse input of A is at least $H = b^2 + 1$. Its final deficit pair is componentwise dominated by a history of the same length consisting of full quotas on A , one possibly shared quota, then full quotas on the opposite initial cohort B . Empty and pure-block degenerations are allowed.*

Proof. On proper inputs $z \geq H$, the lower root in I_p has reached its cap and the reflected root is nonincreasing. Hence $I_p(z) - z$ is nondecreasing, so

$$I_p(v) - I_p(u) \geq v - u \quad (v \geq u \geq H).$$

Before τ , the two input deficits sum to at most M . Thus every old B input is at most $M - H$, where the finite inverse agrees with the bottom map. Let $r = t - i$, let k be the final B deficit, and let p be its final physical color. The cases $r = 0$ and $k = 0$ are immediate. Otherwise write $\mu = \mu_p(k)$ and let $Q \geq \mu$ be the old total quota on B . Pack the exact minimum-ammunition block as late as possible in the same r steps: zeros, one partial quota, then full quotas. Its cumulative quota at step s is $(\mu - m(r - s))_+$. The old cumulative quota is at least $(Q - m(r - s))_+$, and hence at least the new one.

Assign the complementary quota to A . If $\Delta_s \geq 0$ is its cumulative new-minus-old quota, induction gives new-minus-old primary deficit at least Δ_s : the next input difference is at least Δ_{s+1} , and the displayed expansion preserves this inequality. Both histories are legal from the full board, so coordinatewise solo domination and $t < \tau$ keep every new input proper independently of the induction.

The packed B block has nondecreasing inputs, because after its first partial quota it uses $m \geq b + 1$ at every step. Its final input is $L_p(k)$, no greater than the old final input and hence at most $M - H$. Every new B input therefore remains in the bottom range, and the finite-board output is exactly k . Beginning padding preserves the initial cohort: total length and final physical color are fixed. Thus A improves and B is unchanged. Before the packed block the new history simply extends the original pure solo trajectory, giving the asserted two-block form. $\square$

The lower-corner hypothesis is essential to this proof. No analogous exchange, or unrestricted scalar midpoint criterion, is asserted for histories entering that corner. The generic concentration, ancestry and midpoint implications have a Lean development with their inverse and secondary-bound hypotheses explicit; the physical corner-clock and safe-interval instantiations here and the conditional exchange remain ordinary proofs.

8 Explicit formulas at large budgets on every odd rectangle

The exact recurrence of Corollary 5.4 admits a closed solution in a uniform budget range, with no restriction on the length. Let

$$w = 2b + 1 \geq 3, \quad n \geq w \text{ odd}, \quad N = wn = 2h + 1, \quad B = b^2.$$

The majority and minority classes have sizes $K_0 = h + 1$ and $K_1 = h$. All neighborhood profiles and prefixes below are those of Theorem 5.1.

Define $L_0(0) = L_1(0) = 0$ and, for $k > 0$,

$$L_0(k) = k + \min\{R(k), b\}, \quad L_1(k) = k + \min\{Q(k), b + 1\}. \quad (32)$$

These are the corner profiles before the far boundary clips a neighborhood. In particular $g_z(k) \leq L_z(k)$. For $r > 0$ put

$$J(r) = \begin{cases} L_1(r/2), & r \text{ even}, \\ L_0((r+1)/2), & r \text{ odd}, \end{cases} \quad (33)$$

$$\delta(r) = \begin{cases} L_0(r/2 + 1) - L_1(r/2), & r \text{ even}, \\ L_1((r+1)/2) - L_0((r+1)/2), & r \text{ odd}. \end{cases} \quad (34)$$

Set $J(0) = \delta(0) = 0$. The two corner profiles interlace, so $\delta(r) \in \{0, 1\}$. These explicit square-root functions describe the small set left by the first sweep and the inspection cost for starting the second sweep on the same day.

Theorem 8.1 (All odd rectangles above a quadratic budget). *Suppose $m \geq b^2 + 1$. If $m \geq N$, then $T_m(P_w \square P_n) = 1$; if $h \leq m < N$, then $T_m(P_w \square P_n) = 2$. For $m < h$, write*

$$D = 2m - w, \quad N - w - 1 = qD + r, \quad 0 \leq r < D.$$

Then

$$T_m(P_w \square P_n) = \begin{cases} 2q, & r = 0, \\ 2q + 1, & r > 0 \text{ and } 2J(r) + \delta(r) \leq m, \\ 2q + 2, & r > 0 \text{ and } 2J(r) + \delta(r) > m. \end{cases} \quad (35)$$

The constructions use two prefix sweeps, with at most one shared day. The lower bound applies to arbitrary inspection schedules.

Corollary 8.2 (A formula using only integer division). *For $m \geq 2b^2 + 1$ and $m < h$, put $c_m = m - w - 1 + \mathbf{1}_{\{m=2b^2+1\}}$. With the same q, r ,*

$$T_m(P_w \square P_n) = \begin{cases} 2q, & r = 0, \\ 2q + 1, & 1 \leq r \leq c_m, \\ 2q + 2, & r > c_m. \end{cases}$$

8.1 An inspection potential and a solved partial-state recurrence

Write $\rho_z(k) = L_z(k) - k$, with $\rho_z(0) = 0$. For $k > 0$,

$$0 \leq L_0(k+1) - L_1(k) \leq 1, \quad 0 \leq L_1(k) - L_0(k) \leq 1.$$

Both L_z and their surpluses ρ_z are nondecreasing. Thus J is nondecreasing, $J(r+2) \geq J(r) + 1$, and $J(1) = 2$. Throughout the proof we assume $m \geq B + 1$. Then $D > 0$ and $J(D) = m$.

Assign a nonempty parity- z count k the rank $u = 2k + z$. Define

$$F(u) = \begin{cases} \lfloor u/2 \rfloor, & \lfloor u/2 \rfloor \leq m, \\ am + J(s), & u - w - 1 = aD + s, \quad 0 \leq s < D, \quad \lfloor u/2 \rfloor > m. \end{cases} \quad (36)$$

In particular $F(0) = F(1) = 0$. Empty supports always have potential zero. The function measures a lower bound on the total number of inspections still required against one cohort.

Lemma 8.3 (Interior inequality and its equality cases). *The function F is nondecreasing, $F(u) \geq \lfloor u/2 \rfloor$, and $F(u+2) \geq F(u) + 1$. For $0 \leq p \leq m$,*

$$F(2k + z) \leq p + F(2L_z((k - p)_+) + 1 - z). \quad (37)$$

Put $j(u) = \lfloor (u - w - 2)_+ / D \rfloor$. Every noncapture equality in (37) lowers j by exactly one. A capture has source quotient zero.

Proof. For $j \geq 1$, use the positive-remainder convention $u = w + 1 + jD + s$, $1 \leq s \leq D$. Then

$$F(u) = jm + J(s) = \left\lfloor \frac{u + jw}{2} \right\rfloor - \epsilon(s), \quad 0 \leq \epsilon(s) = \left\lfloor \frac{s + w + 1}{2} \right\rfloor - J(s) \leq b - 1.$$

At $j = 0$, $F(u) = \lfloor u/2 \rfloor$, without a correction. These formulas prove the stated monotonicity properties, including the period boundaries. A capturable source has potential k . Every nonempty survivor has $L_z(k - p) > k - p$, so a noncapture transition from $k \leq m$ is strict.

Suppose $k > m$ and put $s = k - p > 0$. On the plateau of L_z , the next rank is $v = u - 2p + w$. If $p \leq b$, then $v \geq u + 1$: monotonicity makes the inequality strict for $p > 0$, while $p = 0$ is strict because $v \geq u + 2$. If $p > b$, the quotient cannot increase and decreases by at most one. When it stays unchanged, the floor term, after including the p inspections, increases by at least b ; the correction can save at most $b - 1$. When the quotient falls once, its positive remainders obey $t = s + 2(m - p)$, with s, t now denoting those remainders. Hence $J(t) - J(s) \geq m - p$, as required. A target of quotient zero has its actual cardinality as potential, at least the corresponding $J(t)$.

Off the plateau, the survivor size is at most $b(b - 1)$ in phase one or $(b - 1)^2$ in phase zero. Its next count is at most $B < m$. Since $m > B$, the noncapturable source has quotient exactly one. Write $a = k - m > 0$ and $e = m - p$, so the survivor size is $a + e$. Substitution gives $F(2k + z) = m + L_z(a)$, and $L_z(a + e) - L_z(a) \geq e$ proves the inequality. This step goes from quotient one to zero. Finally a capture source has $k \leq p \leq m$, hence quotient zero. $\square$

Lemma 8.4 (Full-budget capacity). *Using the unbounded profiles L_z , the largest starting count in phase z that clears in $t \geq 1$ consecutive m -inspection days is*

$$A_z(t) = \left\lfloor \frac{tD + w + 1 - z}{2} \right\rfloor.$$

In particular an actual finite-board prefix of rank at most $tD + w + 1$ clears in t such days.

Proof. All queried capacities are at least $m \geq B+1$, beyond the last nonplateau output. The inverse of L_z there is subtraction by $b + z$. Thus $A_z(1) = m$ and $A_z(t+1) = m + A_{1-z}(t) - (b + z)$. The displayed formula solves these two affine recurrences. Finite-board neighborhoods are no larger, and compatible prefixes realize every step. $\square$

8.2 A finite correction for the far corner

For a count k in phase z , let $d = K_z - k$ be its deficit from a full cohort. For a noncapture step with p useful inspections, write $d' = K_{1-z} - g_z(k - p)$. The profiles imply

$$d' \leq d + p. \quad (38)$$

For majority survivors this uses $g_0(s) \geq s - 1$; for nonempty minority survivors it uses the strict Hall bound $g_1(s) \geq s + 1$. The inequality need not hold at capture, which will be treated separately. If $g_z(s) \neq L_z(s)$, the survivor deficit $x = K_z - s = d + p$ is at most B . Define the finite constant

$$C = \min_{\substack{z \in \{0,1\} \\ 0 \leq x \leq B}} [x + F(2g_z(K_z - x) + 1 - z)] \quad (39)$$

and the monotone potential

$$\Phi_z(k) = \min\{F(2k + z), (C - K_z + k)_+\}, \quad \Phi_z(0) = 0. \quad (40)$$

All sets in (39) are nonempty because $n \geq w$. The $x = 0$ cases give $C \leq F(N), F(N + 1)$.

Lemma 8.5 (Boundary envelope). *For every exact-profile transition,*

$$\Phi_z(k) \leq p + \Phi_{1-z}(g_z((k - p)_+)).$$

Consequently $T_m(P_w \square P_n) \geq \lceil 2C/m \rceil$.

Proof. For a noncapture step, (38) gives the inequality against the target's linear branch. If the far corner does not clip, Lemma 8.3 gives the inequality against its F branch. Otherwise $x = d + p \leq B$, and the defining minimum gives $C - d \leq p + F(\text{next rank})$. Thus the source is at most p plus each target branch. At capture, $k \leq p \leq m$ and $\Phi_z(k) \leq F(2k + z) = k \leq p$. Both initial potentials equal C and both terminal potentials are zero. Summing over the two cohorts gives the bound. Monotonicity of Φ and the isoperimetric lower bounds give the same comparison for arbitrary physical supports; no special shape of an uncompressed belief is assumed. $\square$

Lemma 8.6 (Corner correction). *For $m \geq B + 1$ and $m < h$,*

$$F(N) - 1 \leq C \leq F(N).$$

For $m \geq B + b$, the stronger equality $C = F(N)$ holds.

Proof. For $x > 0$, the two full-cohort departures in (39) have next ranks

$$N - d_E(x), \quad d_E(x) = 2(x - R(x)); \quad N - d_O(x), \quad d_O(x) = 2(x - Q(x)) + 1.$$

For $x \leq B$, the high-end term in the finite profile indeed gives these expressions. A negative gain, possible at $x = 1$ in the minority case, is harmless by monotonicity. Every nonnegative gain is at most $2B - 2b < D$, so at most one quotient boundary is crossed.

Without a quotient crossing, the potential drop is at most x . One can use the strict inequalities proved below in (43) when the target remainder is positive. If that remainder is zero, direct substitution gives $J(d_E(x)) \leq x$ and $J(d_O(x)) \leq x$ instead.

It remains to bound a crossing. Put $a = R(x)$ for an E departure, $a = Q(x)$ for an O departure, and $d = b - a$. Let Δ be $F(N)$ minus the departure cost. The following table writes its periodic endpoint expression; using the actual cardinality branch of F can only decrease Δ . Here the source remainder is $2k$ or $2k + 1$.

Case	Target argument ℓ	c	Δ
E , even	$m - b - x + a + k$	$\rho_1(k)$	$d + c - \rho_0(\ell)$
E , odd	$m - b - x + a + k$	$\rho_0(k + 1)$	$d + 1 + c - \rho_1(\ell)$
O , even	$m - b - x + a + k - 1$	$\rho_1(k)$	$d + 1 + c - \rho_1(\ell)$
O , odd	$m - b - x + a + k$	$\rho_0(k + 1)$	$d + 1 + c - \rho_0(\ell)$

All target arguments are positive. Suppose first that $c > 0$. If $\Delta \geq 2$, its bound on the target surplus is below that surplus's plateau, so the defining square or consecutive-product inequality applies without clipping. Substitute $x \leq a^2$ for E and $x \leq a(a - 1)$ for O , together with $k \geq (c - 1)(c - 2) + 1$ in an even-source row and $k \geq (c - 1)^2$ in an odd-source row. Write $v = 2d(c - a)$ and $I = \mathbf{1}_{\{a=b+1\}}$. Rearrangement gives the bounds in the middle column below. Assuming only $\Delta \geq 1$ increases the allowed target root by one and gives the last column.

Case	Upper bound on $m - B$ if $\Delta \geq 2$	Upper bound if $\Delta \geq 1$
E , even	$v - 3d - c + 1$	$v - d + c - 2$
E , odd	$v - 2d - c + 1$	$v + c - 1$
O , even	$v - d - b - bI$	$v + d + 2c - b - 2 - bI$
O , odd	$v - b - bI$	$v + 2d + 2c - b - 1 - bI$

The indicator comes from $x \leq B$, which improves $x \leq a(a - 1)$ by b when $a = b + 1$.

For E we have $d \geq 0$ and $c \leq a \leq b$, with $c \geq 2$ in the even row. The middle-column bounds are nonpositive; the last-column bounds are at most $b - 2$ and $b - 1$. For O with $a \leq b$, the same conclusions hold, with last-column bounds at most $b - 2$ and $b - 3$; the odd row has $c \leq a - 1$. For $a = b + 1$, the middle-column bounds become $3 - 2c$ and $2 - 2c$, and the last-column bounds both become -1 . Thus $\Delta \geq 2$ contradicts $m \geq B + 1$, while $\Delta \geq 1$ contradicts $m \geq B + b$.

If $c = 0$, only even-source rows occur, with $k = 0$. The assumption $\Delta \geq 2$ forces $d \geq 3$ and gives, respectively, $m - B \leq -2ad - 3d + 4$ and $m - B \leq -2ad - d - b + 3$, both nonpositive. The assumption $\Delta \geq 1$ forces $d \geq 2$ and gives $m - B \leq -2ad - d + 1$ and $m - B \leq -2ad + d - b + 1$, both below b . This completes both lower bounds on C . The $x = 0$ majority departure always supplies the upper bound $C \leq F(N)$. $\square$

8.3 The attaining sweeps

The upper construction works throughout $m \geq B + 1$. Suppose $m < h$. Start with the minority cohort. Every noncapture full-budget shot decreases its rank by at least D . If $r = 0$, its rank is $qD + w + 1$, so Lemma 8.4 clears it in q days. Here q is odd, since $N - w - 1$ and D are odd. The other untouched cohort is therefore also minority when its q -day sweep starts.

If $r > 0$, after $q - 1$ full shots the first rank is at most $2m + r + 1$. Its next shot leaves at most $J(r)$ rooms: the last survivor has phase one and size at most $r/2$ for even r , or phase zero and size at most $(r + 1)/2$ for odd r . Clear those rooms on day $q + 1$.

On the same day, the other cohort is minority for even r and majority for odd r . To put its next rank at most $qD + w + 1$, create a deficit of $r/2 + 1$ majority rooms in the even case, or $(r + 1)/2$ minority rooms in the odd case. Inspecting the neighbors of a terminal corner prefix achieves this at costs, respectively,

$$L_0(r/2 + 1) = J(r) + \delta(r), \quad L_1((r + 1)/2) = J(r) + \delta(r).$$

These neighborhoods are terminal prefixes too; the remaining initial prefix is a compatible next belief. The next q full shots suffice by Lemma 8.4. Thus a shared day works precisely at the stated sufficient budget $2J(r) + \delta(r) \leq m$. Otherwise either full phase has rank at most $(q + 1)D + w + 1$, because $r + 1 \leq D$, and two $(q + 1)$ -day sweeps give the other upper bound. Early capture during a sweep only reduces its actual inspection costs.

For $m \geq N$ one inspection of the entire board suffices. For $h \leq m < N$, inspect the minority class on day one; after movement the other cohort occupies at most the h minority rooms and clears on day two. One day is impossible when $m < N$.

8.4 The equality case and the missing shared-day inspection

For $m \geq B + b$, Lemmas 8.5 and 8.6 give the lower bound $\lceil 2F(N)/m \rceil$. It matches the construction except when

$$2J(r) = m, \quad \delta(r) = 1. \tag{41}$$

We now rule out $2q + 1$ days in this case. Assume first $b \geq 2$. Then $m > B$, $r \geq 2$, $r < D$, and $C = qm + m/2$. Moreover

$$F(N + 1) = C + 1, \quad C - h = \frac{(q - 1)w + m - r}{2} \geq 2. \tag{42}$$

For the latter inequality, $m - r \geq 4$ for even r and $m - r \geq 3$ for odd r . Every linear branch of Φ is consequently positive on nonempty supports; its zero cutoff cannot introduce a false equality.

We first record the strict inequalities used when a full-cohort corner departure leaves the same positive quotient. For $x > 0$, put $a = R(x)$ in the first two rows and $a = Q(x)$ in the last two. For positive displayed arguments and $x \leq B$,

$$\begin{aligned} x + L_1(k - x + a) &> L_1(k), & E, \ r = 2k, \\ x + L_0(k - x + a) &> L_0(k), & E, \ r = 2k - 1, \\ x + L_0(k - x + a) &> L_1(k), & O, \ r = 2k, \\ x + L_1(k - x + a - 1) &> L_0(k), & O, \ r = 2k - 1. \end{aligned} \tag{43}$$

Here E and O indicate the phase of the full cohort, while the source potential is $qm + J(r)$ in both cases.

To verify these inequalities, write c for the source surplus. In the first row, the cost difference is $a + \rho_1(\text{target}) - c$. It is positive if $a \geq c - 1$. Otherwise the bounds $k \geq (c - 1)(c - 2) + 1$ and $x \leq a^2$ place the target above $(c - a)(c - a - 1)$: their difference is at least $2c(a - 1) - 2a^2 + 3 > 0$. Its surplus is therefore at least $c - a + 1$. For the second row, only $a < c$ needs checking; the target exceeds $(c - a)^2$ by at least $2c(a - 1) - 2a^2 + a + 2 > 0$. For the third row, use $a \geq 2$, $x \leq a(a - 1)$, and $k \geq (c - 1)(c - 2) + 1$; when $a < c$, the target exceeds $(c - a)^2$ by at least $(2a - 3)c - 2a^2 + 2a + 3 > 0$. For the last row, its cost difference is $a - 1 + \rho_1(\text{target}) - c$. When $a < c$, the target exceeds $(c - a + 1)(c - a)$ by at least $(2a - 3)c - 2a^2 + 3a + 1 > 0$. The required surpluses are all at most their plateau caps. The remaining cases follow from the minimum positive surpluses one and two. This proves (43).

Suppose now that a strategy captured in $2q + 1$ days. Its available inspections total $2C$, so every step must be tight in the envelope inequality and every day must use all m useful inspections. A positive full departure has $d' < p$. A proper nonempty source has $d' < d + p$ on every noncapture step: its majority survivors have surplus at least zero, and its minority survivors have surplus at least two. Together with (42), these facts make a target linear branch strictly inefficient. Every tight target after departure therefore uses F , and its potential remains below C .

A positive full departure lowers j exactly once. For $p \leq B$ its rank gain is less than D , and (43) excludes no drop. For $p > B$ there is no corner clipping: a majority departure is strict because its initial F value is $C + 1$, while a tight minority departure has one drop by Lemma 8.3.

Consider a later clipped step, with proper-source deficit d and $x = d + p \leq B$. Its output equals that of a full-cohort x -inspection departure. Its source quotient is either q or $q - 1$, since $D \geq 2B - 1$ and $r \geq 2$. If both source and target quotients equal q , the strict corner inequalities give $p + F(\text{target}) > C - d \geq \Phi(\text{source})$. If the source quotient is $q - 1$, then

$$F(\text{source}) < C - d. \quad (44)$$

For completeness, when $r = 2s$ put $\alpha = \rho_1(s)$, so $m/2 = s + \alpha$; when $r = 2s - 1$ put $\alpha = \rho_0(s)$, so again $m/2 = s + \alpha$. Substitution in F bounds $C - d - F(\text{source})$ below by $\alpha - 1, \alpha, \alpha, \alpha$ for $E/\text{even}, O/\text{even}, E/\text{odd}, O/\text{odd}$, respectively. All are positive. At $q = 1$ the source quotient is zero and only even r occurs; its cardinality potential gives exactly $\alpha - 1$ and α . Thus (44) includes this endpoint. The full-corner lower bound gives $x + F(\text{target}) \geq C$, making this proper transition strict too.

Consequently every tight clipped step goes from q to $q - 1$. Every other tight noncapture step has a single quotient drop by Lemma 8.3, and capture starts at quotient zero. Each cohort therefore has exactly $q + 1$ consecutive active days, from its first positive allocation through capture. Two such intervals covering the $2q + 1$ -day horizon start on days 0 and q and overlap once. The first day spends all m on its only active cohort, which must be minority: a full-majority m -inspection departure is unclipped and strict. The first cohort spends qm before the shared day and therefore needs exactly $C - qm = m/2$ inspections on that final day.

For even $r = 2s$, the other cohort is minority on the shared day. Its necessary quotient drop requires at least $s + 1$ majority rooms to be absent from the next belief. All neighbors of those absent rooms must be inspected, requiring at least $g_0(s + 1) = L_0(s + 1) = J(r) + \delta(r) = m/2 + 1$ inspections. For odd $r = 2s - 1$, the other cohort is majority and the same argument uses s absent minority rooms and $g_1(s) = L_1(s) = m/2 + 1$. These profiles equal L_z : the complementary deficits of the absent sets are at least $m > B$, using $q \geq 1$ and the parity of q (for odd r , q is even and at least two). Thus the shared day needs more than m inspections, a contradiction.

For $b = 1$, the case $m = 2$ has $D = 1$ and $r = 0$, and follows directly from $C = F(N)$ and the

solo construction. At $m \geq 3$ the same equality argument applies; a critical remainder has $m \geq 4$, $r = m - 3 \geq 1$, and the clipped proper-source cases are immediate since $B = 1$: only a majority survivor of deficit one can clip, and its neighborhood is the full minority class, of potential C . Such a return to full is impossible after a tight departure. This proves the formula for every $m \geq B + b$.

8.5 The full range $m \geq B + 1$: a one-unit slack argument

It remains to add $B + 1 \leq m < B + b$. For $b = 1$ this interval is empty. For $b = 2$ its only budget is $m = 5$: then $q = n - 2$, $r = 4$, and $J(4) = 4$. The corner-correction lemma gives $C \geq 5q + 3$, so the initial potential $2C \geq 10q + 6$ exceeds the $10q + 5$ inspections available in $2q + 1$ days. The two sweeps above attain $2q + 2$ days. This proves that endpoint directly, without using a separate five-row classification. Henceforth assume $b \geq 3$. The minimum square size and $D \leq 2B - 3$ give $q \geq 2$. Put $\eta = F(N) - C \in \{0, 1\}$. At $r = 0$, the $c = 0$ positive-discount bounds in the proof of Lemma 8.6 are nonpositive, so $\eta = 0$ and the lower time is $2q$.

A positive corner discount necessarily has $\delta(r) = 0$. To see this, suppose $\delta(r) = 1$. In an even-remainder row of the corner table, $c = \rho_1(k) = \rho_0(k + 1)$, so $k \geq (c - 1)^2$, improving the earlier lower bound by $c - 2$. In an odd-remainder row, $\rho_1(k + 1) = c + 1$, so $k \geq c(c - 1)$, an improvement by $c - 1$. The four upper bounds on $m - B$ for a positive discount become

$$v - d, \quad v, \quad v + c - a, \quad v + 2d + c - b, \quad v = 2d(c - a).$$

The first three are nonpositive since $d \geq 0$ and $c \leq a$. The last is at most $-d - 1$, since $c \leq a - 1$. The exceptional last minority band $a = b + 1$ already has upper bound -1 before this improvement. Thus a discount contradicts $m > B$ when $\delta = 1$.

For $r > 0$, Lemma 8.6 gives $C \geq qm + 1$. The lower time is at least $2q + 1$, so only a proposed upper of $2q + 2$ needs attention. Its unresolved cases are precisely

	η	$2J(r) - m$	$\delta(r)$
A	0	0	1
B	1	1	0
C	1	2	0.

We exclude $2q + 1$ days in all three cases.

Exact counters and tight proper tails. By Corollary 5.4, it suffices to study exact profile counters. More explicitly, retain any alleged physical strategy's useful allocations and evolve the minimum-profile counters, capping allocations if the smaller counter has already been reached. Monotonicity keeps these counters below the actual cardinalities, and compatible prefixes attain them with no greater daily budget. They therefore capture by the alleged deadline; pad with zero allocations if capture occurs early. All following equalities use these exact g_z updates.

The nonnegative slack of one cohort step is $p + \Phi(\text{next}) - \Phi(\text{current})$. Add each day's unused budget. Over a $(2q + 1)$ -day successful schedule the total is $(2q + 1)m - 2C$, which is 0, 1, 0 in cases A, B, C , respectively. In A and B ,

$$C - h = \frac{(q - 1)w + m + (2J - m) - 2\eta - r}{2} > 1.$$

Indeed $m - r \geq 3$, $q \geq 2$, and $w \geq 7$ suffice. Every nonempty linear branch $C - d$ is thus positive and strictly exceeds the support size.

Every tight noncapture step from a proper nonempty support is positive, remains proper, and lowers j by one. Here are the adjustments needed to extend the preceding argument to the present budgets. A linear target is still strict because $d' < d + p$. An unclipped tight step has one quotient drop by Lemma 8.3. A return to full is strict because its potential C exceeds that of the source.

For a positive clipped step, $d + p \leq B$, so $d \leq B - 1$. Its source quotient is q or $q - 1$, because $D + r > 2B - 2$. For $b \geq 4$, in case A the even and odd bounds are, respectively, $D + r \geq 3B - 4b$ and $D + r \geq 3B - 4b + 1$; case B improves them by one. They all exceed $2B - 2$. For $b = 3$, direct substitution into J gives $J(4) = 4$, $J(5) = 5$, $J(6) = J(7) = 6$, $J(8) = 7$, and $\delta(5) = \delta(7) = 1$, $\delta(6) = 0$. Hence the possible critical residues are exactly

	m	r	$D + r$
A	10	5	18
B	11	6	21
C	10	6	19.

In particular the rows A, B needed here both exceed $2B - 2 = 16$. This strict inequality also handles the positive-remainder convention at zero. The target quotient is q or $q - 1$ since every full-corner gain is less than D .

If both source and target have quotient q , the full departure with $x = d + p$ inspections costs strictly more than $F(N)$ in A , by (43), or at least $F(N) > C$ in B . Subtracting d rules out equality. If the source quotient is $q - 1$, write $r = 2s$, $\alpha = \rho_1(s)$, or $r = 2s - 1$, $\alpha = \rho_0(s)$. Substitution gives lower bounds

$$C - d - F(\text{source}) \geq \alpha - \eta - 1, \quad \alpha - \eta, \quad \alpha - \eta, \quad \alpha - \eta$$

in the order E/even , O/even , E/odd , O/odd . They are positive: case A needs only $\alpha \geq 2$ in even parity and $\alpha \geq 1$ in odd parity. In B , $J = (m + 1)/2 \geq 6$ forces even $\alpha \geq 3$ and odd $\alpha \geq 2$. There is no cardinality exception because $q \geq 2$. The full-corner lower bound again rules out equality. Only a transition from q to $q - 1$ remains.

A zero allocation from a proper support is strict directly. A minority neighborhood adds at least two rooms. A majority neighborhood adds at least one, except at deficit one when it returns to the full minority class. Thus F increases strictly in the former cases, and the full potential C is strictly larger in the latter. The linear branch also increases, since $d' < d$. Finally, tight capture begins at quotient zero: its potential can equal the support size only in the cardinality branch. A proper tight tail consequently lasts exactly $j + 1$ consecutive days.

Case A: the phase obstruction persists. Every step and day is tight. The full-departure argument using (43) is unchanged, so each cohort has $q + 1$ consecutive active days. They overlap once, the first cohort is minority, and it requires $J = m/2$ inspections on the shared day. The fresh cohort requires $J + \delta = m/2 + 1$ by the omitted-set argument above. The small-set profiles are uncut: their complementary deficits are $(qD + w)/2$ for even r , when q is odd and at least three, and at least $(qD + w - 1)/2$ for odd r , when q is even and at least two. Both exceed B . Hence the shared day exceeds the budget.

Discounted cases: the first day uses the slack. A tight positive full departure in B or C uses the F target branch, has $p \leq B$, and crosses a quotient boundary. The linear branch is strict

by $d' < p$; an uncut departure costs at least $F(N) > C$. The crossing conditions imply

	E departure	O departure
r even	$p \geq J + 1$	$p \geq J$
r odd	$p \geq J$	$p \geq J + 1$.

For even $r = 2s$, an E crossing has $p - R(p) \geq s + 1$ with $R(p) \geq Q(s)$; an O crossing has $p - Q(p) \geq s$ with $Q(p) \geq Q(s)$. For odd $r = 2s - 1$, use $R(s) \leq R(p)$ in the majority case and $R(s) \leq Q(p) - 1$ in the minority case. These prove the displayed bounds, including the surplus caps since $p \leq B$.

The first day's total slack is at least one. Unused budget already contributes one; allocating all m to one cohort costs at least one slack because $m > B$ and $C < F(N)$. If all m are divided between two positive tight departures, their minima sum to at least $2J + 1 > m$, impossible. Case C , which has zero total slack, is excluded.

In case B , day one consumes exactly the single available unit. Every later day uses all m inspections and every later step is tight. If both outputs of day one are proper, their quotients are at most q : a proper output of a full departure has rank at most N . Both tails then finish by day $q + 2 < 2q + 1$, leaving an entirely unused later day, a contradiction.

If exactly one output is proper, the other remains full. Its probes and any unused budget contribute their entire amount to slack, and so total at most one. The proper allocation is therefore $p \geq m - 1 \geq B$. Define its gain here as N minus the target rank (the initial majority rank is $N + 1$). This gain lies between $D - 3$ and D : the majority gain is $2B - 2b$ at $p = B$ and $2p - 2b - 2$ at $p > B$; the minority gain is $2p - 2b - 1$. The opposite corner is on its plateau since $n \geq w$ and $p \leq m < B + b$. Here $r \leq m - 3 < D - 3$ and $r \geq 2$, so this output has quotient $q - 1$. Both outputs cannot remain full, since that would cost slack $m > 1$.

The proper cohort has exactly q consecutive tight days left; a fresh tight departure of the other begins $q + 1$ consecutive days. They cover the remaining $2q$ days and overlap once, since no later day's budget may be unused. Immediately after day one the proper potential is $C - m + 1$. Before the overlap it receives another $q - 1$ full allocations, so its capture requires

$$C - qm + 1 = J$$

inspections. The fresh tight departure requires at least J too. But $2J = m + 1$, again exceeding the budget. This excludes B and proves Theorem 8.1 for every $m \geq B + 1$.

8.6 Removing the root functions

For Corollary 8.2, the function $H(r) = 2J(r) + \delta(r)$ is strictly increasing for $r \geq 1$: its increments alternate between positive increments of L_0 and L_1 . If $m = 2B + 1$ and $b \geq 2$, put $k = b(b - 1)$. Then

$$H(2k) = L_1(k) + L_0(k + 1) = m, \quad H(2k + 1) = L_0(k + 1) + L_1(k + 1) = m + 2.$$

The last admissible remainder is therefore $2k = m - w$. If $m > 2B + 1$, substitution at $r = m - w - 1$ and $r + 1 = m - w$ puts the relevant arguments on the plateaus. The only endpoint is $m = 2B + 2$, where $L_1(B - b) = B$; there $H(m - w - 1) = m - 1$ and $H(m - w) = m + 1$. Otherwise $H(m - w - 1) = m$ and $H(m - w) > m$. For $b = 1$, the cutoffs at $m = 3, 4$ are zero and follow directly from $H(1) = 5$; the other cases use the same plateau calculation.

9 Exact profiles on every even-area rectangle

Let $G = P_w \square P_n$, where $2 \leq w \leq n$ and wn is even, and put $b = \lfloor w/2 \rfloor$ and $h = wn/2$. Reflection in an even-length coordinate interchanges the checkerboard classes, so their profiles agree. Write $\rho(d) = \min\{s \geq 0 : d \leq s(s+1)\}$.

Theorem 9.1 (Even-area neighborhood profile). *For every $0 \leq k \leq h$,*

$$g(k) = k + \min\{\lceil \sqrt{k} \rceil, b, \rho(h-k)\}.$$

Equivalently, every one-color set of surplus $s = |N(S)| - |S| < b$ satisfies

$$|S| \leq s^2 \quad \text{or} \quad h - |S| \leq s(s+1). \quad (45)$$

The minimum is attained for each cardinality and each color. These minimizers are not asserted to form a compatible dynamic order.

9.1 Matching contraction and a common component

Choose an even side length e and let o be the other side. Use physical coordinates (x, y) with $0 \leq x < e$, $0 \leq y < o$. Match rows $2i, 2i+1$ in every column. Identify each color- p vertex with (i, y) by $x = 2i + ((p-y) \bmod 2)$. A physical edge followed by the inverse matching gives a directed graph D_p on an $(e/2) \times o$ array. It has loops, bidirectional horizontal edges, and vertical rungs directed down in one column parity and up in the other. The two choices of p transpose the directed graph.

For the image R of S , put $B = N_{D_p}^+(R) \setminus R$. Then $|B| = s \geq 0$, and R is outgoing-closed in $D_p - B$. If $s < b$, at most s columns contain B . Since $o \geq 2b \geq 2s + 2$, two consecutive columns avoid B . They form a strongly connected ladder: its two rung directions allow travel both ways between rows. Every row avoiding B meets this ladder and is a full bidirectional path. Thus all such rows lie in a single strongly connected component C , and at least one exists because $e/2 \geq b > s$. An outgoing-closed R either contains C or avoids it.

9.2 A sharp area bound outside the component

Suppose $U \cap B = \emptyset$, $N^+(U) \setminus U \subseteq B$, and every row avoiding B is empty in U . We prove $|U| \leq s^2$. Choose an empty row r avoiding B . Above it, let P be the column parity with downward rungs and put

$$\alpha_i = |U_i \cap P|, \quad \beta_i = |U_i \setminus P|, \quad t_i = |B_i \cap P|, \quad w_i = |B_i|.$$

Rung closure gives $\alpha_i \leq \alpha_{i+1} + t_{i+1}$ and hence $\alpha_i \leq \sum_{i < j < r} t_j$. Horizontal matching gives $\beta_i \leq \alpha_i + t_i$. For even o use a perfect matching of the horizontal path. For odd $o = 2q + 1$, every proper subset of the majority column parity matches into its neighbors, as does every subset of the minority. The sole possible exception is the full majority. That would force $q \leq \alpha_i + t_i \leq \sum_{i \leq j < r} t_j \leq s$, contrary to $q \geq b > s$. Thus in all cases each occupied row satisfies

$$|U_i| \leq w_i + 2 \sum_{i < j < r} w_j. \quad (46)$$

Let $W = \sum_{i < r} w_i$ and let a be the number of occupied rows above r . Every occupied row contains a boundary vertex, so $a \leq W$. In the sum of (46), a boundary vertex in row j has

coefficient twice the number of earlier occupied rows, plus one if its own row is occupied. Reserve one vertex in each occupied row. Their coefficients are $1, 3, \dots, 2a - 1$, summing to a^2 ; every remaining vertex has coefficient at most $2a$. The total is therefore at most $a^2 + 2a(W - a) \leq W^2$.

Below r , reverse row order and use upward rungs. If W' counts the boundary there, this gives area at most $(W')^2$. Since $W + W' = s$, we obtain $|U| \leq W^2 + (W')^2 \leq s^2$. The argument includes $b = 1, s = 0$.

If R avoids C , apply this bound to R . Otherwise set $U = V(D_p) \setminus (R \cup B)$. It avoids every boundary-free row and is outgoing-closed outside B in the transposed graph: an edge from U to R there would be an edge from R to U originally. The same area bound gives $h - |R| = |U| + s \leq s(s + 1)$. This proves (45) and the profile lower bound for every physical support.

9.3 Attaining the three bounds

For the plateau bound, return to coordinates with w transverse rows and n longitudinal columns. Take an ideal of the predecessor relation $(x, y) \succ (x \pm 1, y - 1)$ on one checkerboard class. Its row cutoffs a_x have the required row parity and satisfy $|a_{x+1} - a_x| = 1$; virtual cutoffs $-2, -1$ represent empty rows. The neighborhood cutoff in row x is at most $a_x + 1$. Write $s_x = (p - x) \bmod 2$ for its source parity. The rowwise cardinality change is $(a'_x - a_x + 2s_x - 1)/2$. For even w , the parity correction sums to zero; for odd w , choose $p = 0$, when it sums to -1 . In both cases the total surplus is at most $b = \lfloor w/2 \rfloor$. Reflection in an even-length coordinate supplies the other color. Every size k occurs: take an initial segment of any linear extension of this finite predecessor order.

If $s = \lceil \sqrt{k} \rceil < b$, the even-color quadrant prefix of size k , ordered by increasing coordinate sum and then decreasing transverse coordinate, has exactly $k + s$ neighbors. It lies within coordinates $0, \dots, 2s - 2$, with neighborhood within $0, \dots, 2s - 1$, so both fit inside G . This attains the small-corner bound; $k = 0$ is immediate.

For the other end put $d = h - k$ and $s = \rho(d) < b$. Then $x = d - s$ satisfies $0 \leq x \leq s^2$. The small-corner construction provides an opposite-color set X of size x with at most $x + s = d$ neighbors. Choose k vertices outside $N(X)$ in the desired color. Their neighborhood avoids X , so has size at most $h - x = k + s$. Reflection in an even-length coordinate supplies either color in these constructions. Each branch that improves on b is thus attained, proving the theorem.

Together with Theorem 5.1, this settles the static neighborhood profiles for all nondegenerate rectangles. The scalar profile alone need not determine optimal time: on 6×6 with five inspections its full-budget solo recurrence permits 18, 15, 13, 11, 9, 6, 2, 0. Combining two six-step preparations with a central inspection would give thirteen days if those prescribed sizes were physically attainable. The exact physical optimum is fourteen, as independently certified in the research archive. This is an obstruction to dynamic attainment, not to the profile theorem.

10 Corner restrictions and memory on an even-width half-strip

Static minimizers need not be compatible with an earlier belief. Here we quantify one source of incompatibility without imposing a shape on that belief. Work in the *whole* half-strip

$$H_b = \{(x, y) : 0 \leq x < 2b, y \geq 0\}, \quad b \geq 1,$$

with square-grid adjacency. For color zero put $c = (0, 0)$ and $c' = (2b - 1, 0)$. These are the favorable bottom corners of the current and next colors. Reflection supplies the other color. A restriction

on a source set does not delete vertices from the graph: every neighborhood below is taken in H_b . Write

$$R(k) = \lceil \sqrt{k} \rceil, \quad Q(k) = \min\{s \geq 1 : k \leq s(s-1)\}, \quad C_h(s) = \max\{s(s-1)/2, s(s-h)\}.$$

The last function is the punctured-quadrant capacity of Theorem 4.1.

Theorem 10.1 (Exact conditional corner profiles). *For every $k > 0$, the least neighborhood size of a k -room color-zero set subject to each indicated condition is k plus the following surplus:*

condition on S	minimum surplus
none	$\min\{R(k), b\}$
$c \notin S$	$\min\{Q(k), b+1\}$
$c' \in N(S)$	$\min\{Q(k), b\}$
$c' \notin N(S)$	$\begin{cases} R(k), & k < b^2, \\ b+1, & k \geq b^2. \end{cases}$

For the joint condition $c \notin S$, $c' \notin N(S)$, the exact capacity at surplus at most s is, when $b \geq 2$,

$$\Gamma_b(s) = \begin{cases} C_2(s), & 0 \leq s \leq b, \\ \max\{b(b+1)/2, b^2 - 2\}, & s = b+1, \\ \infty, & s \geq b+2. \end{cases} \quad (47)$$

Thus its minimum surplus is $\min\{s : k \leq \Gamma_b(s)\}$. For $b = 1$ the joint minimum is instead two for every $k > 0$. Every asserted minimum is attained at every cardinality.

10.1 Matching rows and the marginal restrictions

Use the matching contraction of Section 9: the source room represented by (i, y) has physical coordinate $x = 2i + (y \bmod 2)$. The contracted graph has horizontal edges in both directions, upward rungs at even y , and downward rungs at odd y . If $B = N_D^+(S) \setminus S$, then $|B| = |N(S)| - |S| =: s$. A matched row avoiding B is horizontally closed in a ray and finite, hence empty in S . Its two empty physical rows separate the support into parts with disjoint neighborhoods. Each part has its full neighborhood unchanged when viewed as a quadrant; the empty row supplies its transverse neighbor layer. Surpluses therefore add. Such a row always exists if $s < b$.

The two unrestricted quadrant capacities are $C_0(u) = u^2$ and $C_1(u) = u(u-1)$, so their sum at total surplus s is at most s^2 . If c is omitted, the capacities become C_2 and C_1 , both at most $u(u-1)$; their sum is at most $s(s-1)$. If c' belongs to the neighborhood, the odd-quadrant part is nonempty and consumes surplus $r \geq 2$. Even allowing the other part its square capacity gives

$$|S| \leq (s-r)^2 + r(r-1) \leq s(s-1).$$

These prove the subcritical lower bounds in the first three rows.

For the omitted-current bound at $s = b$, only the absence of a boundary-free row remains. Every matched row then has exactly one boundary vertex. Its occupied set is empty or an initial interval of length c_i , because any other finite subset of a ray has two horizontal boundary vertices. Here $c_0 = 0$. Upward even rungs give $c_{i+1} \leq c_i + 2$, including when $c_i = 0$, and consequently

$$|S| \leq \sum_{i=0}^{b-1} 2i = b(b-1).$$

Thus omission of c requires surplus at least $b + 1$ beyond this size.

For later use, omission of the outgoing corner has the sharper critical bound

$$c' \notin N(S), s = b \implies |S| \leq b^2 - 1. \quad (48)$$

With a boundary-free row the two capacities are C_0 and C_3 , whose sum is at most b^2 . Equality forces all surplus into the unrestricted even quadrant and the full square extremizer there. Indeed, equality in the full-layer proof of Theorem 4.1 forces every occupied diagonal to be full and consecutive. That square reaches transverse coordinate $2b - 2$, whereas a component before an empty matched row reaches at most $2b - 3$. Equality is impossible. Without a boundary-free row, all occupied rows are prefixes, the last is empty, and downward odd rungs give $c_i \leq 2(b - i) - 1$ for $i < b - 1$. Their sum is $b^2 - 1$. This proves (48) and the fourth lower bound in the theorem.

We record the constructions together. Finite downward pyramids, meaning ideals under $(x, y) \succ (x \pm 1, y - 1)$, exist at every size and have surplus at most b : their neighborhood cutoffs are at most their source cutoffs plus one. Small even-quadrant prefixes give surplus $R(k)$; odd-quadrant prefixes at the opposite corner give surplus $Q(k)$ and touch c' . Fill a partial diagonal toward smaller transverse coordinate. Then every even prefix of size $k < b^2$ avoids the two neighbors of c' , attaining the fourth row below its threshold.

The following elementary bounds supply all larger constructions. If (v, t) is missing from a pyramid of source color zero, its cutoffs obey $a_x \leq t - 2 + |x - v|$, and its size is at most $\sum_x (t + |x - v| - (x \bmod 2))/2$. In particular, omission of $(1, 1)$ bounds the size by $b(b - 1) + 1$, and omission of $(2b - 2, 2)$ bounds it by $b^2 + 1$. The sharper bottom-root threshold in Lemma 11.2 is $b(b - 1)$. For $k > b(b - 1)$, take a pyramid of size $k + 1$ and remove c . The forced room $(1, 1)$ covers both neighbors of c , so this has surplus at most $b + 1$. A pyramid of size $k > b(b - 1)$ also contains $(2b - 2, 0)$ and hence touches c' , proving the remaining attainment with outgoing corner present. For $k \geq b^2$, take a pyramid of size $k + 2$ and remove the two neighbors $(2b - 2, 0), (2b - 1, 1)$ of c' . The forced room $(2b - 2, 2)$ and its predecessors cover every other neighbor of these rooms. Exactly c' disappears from the neighborhood, so the resulting surplus is at most $b + 1$. This completes all marginal profiles, including $b = 1$.

10.2 Both omissions, including the critical extra unit

Suppose first $b \geq 2$. The joint condition forbids exactly the three source rooms

$$(0, 0), \quad (2b - 2, 0), \quad (2b - 1, 1).$$

A boundary-free row splits their capacities into C_2 on the left and C_3 on the reflected right. Convexity, their zero values at zero, and $C_3 \leq C_2$ give total capacity at most $C_2(s)$. If $s = b$ and there is no such row, the preceding prefix argument has both endpoint rows empty and gives

$$c_i \leq \min\{2i, 2(b - i) - 1\} \quad (1 \leq i \leq b - 2), \quad |S| \leq b(b - 1)/2 - 1 \leq C_2(b).$$

This proves the first line of (47).

Now let $s = b + 1$. If both neighbors of c already belong to $N(S)$, adding c leaves the neighborhood unchanged and still omits c' . Equation (48) then gives $|S| + 1 \leq b^2 - 1$, or $|S| \leq b^2 - 2$. Assume henceforth that this origin-addition argument is unavailable.

If there is a boundary-free row and the right part is nonempty, its surplus is at least two. For $b \geq 3$, convexity bounds the total by

$$\max\{C_3(b + 1), C_2(b - 1) + C_3(2)\} \leq b^2 - 2.$$

If only the left part is nonempty, its sole larger possibility is $C_2(b+1) = b^2 - 1$. Equality in the punctured-quadrant proof forces full diagonals $2, 4, \dots, 2b-2$, which cannot fit before an empty matched row. Diagonal compression toward the smaller transverse coordinate preserves that width restriction, so equality is impossible even for an uncompressed source. For $b=2$ the direct quadrant bound is already $C_2(3) = 3$, as required.

It remains to consider $s = b+1$ with no boundary-free row. Exactly one row has two boundary vertices; all other occupied rows are prefixes. First suppose both endpoint rows are empty. Write α_i, β_i for the numbers of even and odd occupied columns and t_i, u_i for the respective boundary counts. Rung inclusion gives

$$\alpha_i \leq \sum_{j < i} t_j, \quad \beta_i \leq \sum_{j > i} u_j,$$

and therefore

$$|S| \leq \sum_j ((b-1-j)t_j + ju_j) \leq (b+1)(b-1) - (b-2) = b^2 - b + 1.$$

For the second inequality reserve one boundary vertex on each internal row, where its coefficient is at most $b-2$; every other coefficient is at most $b-1$. This is at most $b^2 - 2$ for $b \geq 3$, and is three when $b=2$. No interval assumption on the exceptional internal row was used.

If an endpoint row is occupied, it must be the exceptional row. Its source is a single interval avoiding column zero, since two components away from zero would require more than two horizontal boundary vertices. There are two endpoint cases.

- At the first row, an interval starting at one contains the physical room $(1, 1)$ and permits origin addition. In the remaining case its first column is at least two. The next row must be empty: otherwise its prefix sends an even rung to column zero in the first row, creating a third boundary vertex. At most one odd interval room can feed that empty next row, so the interval has length at most three. The other prefixes grow by at most two per row, giving total at most $(b-2)(b-3) + 3$.
- At the last row, the interval starts at least at two because both columns zero and one are forbidden. The earlier prefixes satisfy $c_i \leq 2i$. If the preceding prefix has length $c > 0$, its sole boundary is at c ; the interval's even coordinates are at most c , and its length is at most c . The total is at most $(b-2)(b-1) + 2(b-2) = b^2 - b - 2$. If that preceding row is empty, at most one even interval room can feed it. The interval has length at most three, and the earlier prefixes contribute at most $(b-3)(b-2)$.

Each bound is at most $b^2 - 2$ for $b \geq 3$; the endpoint cases for $b=2$ give at most three. Together these exhaust the extra-unit cases and prove the finite upper capacity in (47).

For sharpness at $s \leq b$, both the proper-ramp and full-layer constructions for $C_2(s)$ avoid all three forbidden rooms. The full layers end at diagonal $2s-4$; the ramp is filled toward small transverse coordinate. Their partial terminal layers give every intermediate size. For $b \geq 3$, fill the even diagonals $2, 4, \dots, 2b-2$ and delete $(2b-2, 0)$. This has $b^2 - 2$ rooms and surplus $b+1$. Every smaller size at this surplus bound comes from an even-quadrant prefix of size $k+1$, filled toward smaller transverse coordinate, with its origin deleted. For $b=2$ the ramp $\{(0, 2), (0, 4), (1, 3)\}$ and its initial subprefixes give sizes up to three.

For every larger k , take a downward pyramid of size $k+3$ and remove all three forbidden rooms. Its size forces $(1, 1)$ and $(2b-2, 2)$ by the two omission bounds above; their predecessors

ensure that all three rooms to be removed are present. Their other neighbors remain covered, so only c' disappears from the neighborhood. The surplus is at most $b + 2$. This proves every-size attainment, not just unboundedness along a subsequence. When $b = 1$, the two bottom forbidden rooms coincide; the joint restriction is exactly the outgoing-absent restriction already proved, with minimum surplus two. This finishes the theorem.

10.3 Localization gives more than availability bits

Corollary 10.2 (Triangular localization and propagation delay). *Let S be a nonempty color-zero set, with surplus $s < b$. If $|S| > s(s - 1)$, then S is connected under the relation of sharing a neighbor, contains c , and*

$$x + y \leq 2s - 2 \quad ((x, y) \in S), \quad c' \notin N^j(S) \quad (j < 2(b - s) + 1).$$

If $c \notin S$ and $|S| > C_2(s)$, then S is connected under the same relation, $c' \in N(S)$, and

$$(2b - 1 - x) + y \leq 2s - 3 \quad ((x, y) \in S), \quad c \notin N^j(S) \quad (j < 2(b - s) + 2).$$

Both exclusions remain valid after arbitrary intervening inspections. Here j counts movements from the survivor S .

Proof. Sharing-neighbor components have disjoint neighborhoods, so their sizes and surpluses add. A nonempty component omitting c has surplus at least two and capacity $u(u - 1)$ at surplus $u < b$. For the first claim, if no component contains c , the total capacity is at most $s(s - 1)$. If one contains c with surplus $a \geq 1$ and has companions of total surplus $d \geq 2$, its total capacity is at most $a^2 + d(d - 1) \leq (a + d)(a + d - 1)$. Both are contradictions. Thus S is one component containing c .


For the second claim, every component has surplus at least two. Two or more components have total capacity at most

$$(s - 2)(s - 3) + 2 \leq s(s - 2) \leq C_2(s), \quad s \geq 4;$$

this follows by merging all but one component and maximizing the convex two-part pronic sum at an endpoint. For $s < 4$ two components are impossible. Thus S is connected; the joint profile then forces $c' \in N(S)$.

A boundary-free matched row exists in both cases. Connectedness confines the whole support to the corresponding corner quadrant, with its actual neighborhood unchanged. In the first case its occupied even diagonals are $0, 2, \dots, 2L - 2$: a sharing-neighbor edge changes the coordinate sum by zero or two, so there is no gap. Diagonal compression preserves which diagonals are occupied. The layer inequalities in Theorem 4.1 charge at least one surplus unit per occupied diagonal; hence $L \leq s$. This bounds the original support, without claiming it is compressed. In the reflected odd quadrant of the second case, $c' \in N(S)$ forces the first diagonal to be one. The occupied diagonals are $1, 3, \dots, 2L - 1$, and the output origin supplies one more surplus unit, so $L + 1 \leq s$.

The bottom corners have distance $2b - 1$. Subtracting the respective triangle radii gives the two stated distances. Every later belief is a subset of the corresponding free neighborhood iterate, proving the inspection-independent exclusions. Full square and pronic triangles attain these distances, so the strict inequalities on j are intentional. $\square$



Width 8, surplus 3. Arrows show the first possible arrival; inspections can only delay it.

Figure 3: The sharp localization examples for $b = 4, s = 3$. Filled rooms are S ; open rooms are $N(S)$. Crosses mark the excluded opposite corner. Its shortest arrival times are three and four movements.

These statements concern arbitrary supports, but are not a sufficiency theorem for corner-availability bits. For example, on width fourteen a 30-room set omitting its current favorable corner and having surplus six must be the opposite odd triangle on diagonals 1, 3, 5, 7, 9. Indeed the separating-row capacities force all surplus into that quadrant; equality in its full-layer bound forces each raw diagonal to be full. Its 36-room neighborhood cannot reach the other bottom corner on the following movement. Remembering only current corner presence loses this distance information even though each individual conditional profile is sharp. No full-board capture time is being asserted by this example.

Finally, on a finite rectangle the half-strip profile applies to a source that avoids the far longitudinal row, so its neighborhood agrees with the half-strip neighborhood. That guard must be checked at every use of a profile inequality. Once localization has been established, the propagation exclusions also hold on the finite board, whose free neighborhoods can only be smaller.

11 Uniform time bounds and exact large-budget searches on even widths

Continue with $G = P_{2b} \square P_n$, $n \geq 2b$, and $h = bn$. The preceding static profile gives an all-budget lower bound without assuming that its minimizing sets can be chained.

Lemma 11.1 (A scalar lower bound with a central-day test). *More generally, let g be the symmetric profile of any even-area rectangle, with h vertices in each color. Starting with $c_0 = h$, define*

$$c_{t+1} = g((c_t - m)_+), \quad \tau = \min\{t : c_t = 0\}.$$

If τ is finite, then $T_m \geq 2\tau - 1$, and $T_m \geq 2\tau$ when $2c_{\tau-1} > m$.

Proof. The inverse profile is

$$I(z) = h - g(h - z) = z - \min\{\rho(z), b, R(h - z)\}, \quad 0 \leq z \leq h.$$

Its subtracted term is subadditive. If a cap or decreasing reflected branch minimizes either summand, that branch bounds the value at the sum. Otherwise use subadditivity of ρ , following from $u \leq r(r + 1)$, $v \leq s(s + 1)$ and $u + v \leq (r + s)(r + s + 1)$. Thus $I(x) + I(y) \leq I(x + y)$ when $x + y \leq h$.

Put $D_t = h - c_t$. Then $D_{t+1} = I(\min\{h, D_t + m\})$. For $t < \tau$, the total physical deficit of both cohorts after t inspections and moves is at most D_t . Indeed, useful allocations of total at most m give next deficit at most $I(d_0 + p) + I(d_1 + q) \leq I(D_t + m)$ whenever $t + 1 < \tau$. Here $D_t + m < h$ follows from $D_{t+1} < h$ and $I(h) = h$. This proves the invariant inductively from zero.

The case $\tau = 1$ is immediate. Otherwise, for a purported $2\tau - 2$ -day capture, set $L = \tau - 1$. After the first L inspections and moves the belief has more than h vertices. Reversing the last L inspections, using $L - 1$ moves and then a final inspection, also leaves more than h vertices: its deficit is at most $D_{L-1} + m < h$. The two sets occupy the same day and intersect, giving an avoiding walk. Thus $T_m \geq 2\tau - 1$.

At the central day of a $(2\tau - 1)$ -day search, the forward and backward beliefs each have deficit at most $D_{\tau-1}$. Their intersection has at least $2h - 2D_{\tau-1} = 2c_{\tau-1}$ vertices, all of which must be inspected centrally. This proves the second bound. $\square$

11.1 A common enlargement and erosion lemma

The following construction supplies the physical upper strategies in both parity cases. It also permits a nonpyramidal terminal survivor. Let Q be a finite connected bipartite graph with at least two vertices, with bipartition map $\chi : V(Q) \rightarrow \{0, 1\}$. In $Q \square P_n$, a *downward pyramid* is a one-color set closed under every valid predecessor $(u, y - 1)$ of (v, y) , where $uv \in E(Q)$ and $y > 0$. Write $\text{cl}(R)$ for the least such set containing R . These are ideals of a finite poset, since every predecessor lowers y . Put

$$H_Q = \max_{v \in V(Q)} \sum_{u \in V(Q)} \lceil \text{dist}_Q(v, u)/2 \rceil.$$

The sharper, phase-dependent threshold needed for erosion is

$$B_{Q,p} = \max_{v: \chi(v)=p} \sum_{u \in V(Q)} \lfloor \text{dist}_Q(v, u)/2 \rfloor.$$

Lemma 11.2 (Prescribed envelopes with arbitrary terminal supports). *Let R be any set in physical color p , and let k satisfy*

$$|\text{cl}(R)| \leq k \leq |V_p|, \quad k > B_{Q,p},$$

where V_p is the full physical color class in the cylinder. There are pyramids $A \supseteq R$ of color p and E of color $1 - p$ such that

$$|A| = k, \quad |E| = k - |\chi^{-1}(p)|, \quad N(E) \subseteq A.$$

Consequently, after a movement from any survivor contained in E , the designated inspections $A \setminus R$ leave an actual survivor contained in R , using exactly $k - |R|$ designated inspections.

Proof. Two predecessor moves show that every fiber of a pyramid is a spacing-two prefix. Put $s_v = (p - \chi(v)) \bmod 2$. If its fiber size is k_v , the virtual last height $a_v = s_v - 2 + 2k_v$ satisfies $|a_u - a_v| = 1$ on every Q -edge; empty fibers have heights -2 or -1 . If fiber v is empty, distance along Q gives

$$k_u \leq \frac{s_v + \text{dist}(v, u) - s_u}{2} = \begin{cases} \lfloor \text{dist}(v, u)/2 \rfloor, & s_v = 0, \\ \lceil \text{dist}(v, u)/2 \rceil, & s_v = 1. \end{cases}$$

Thus a pyramid with any empty fiber has at most H_Q rooms. If an empty fiber has $s_v = 0$, its sharper bound is $B_{Q,p}$. These are exactly the fibers whose bottom room contributes to erosion.

Extend the ideal $\text{cl}(R)$ to exactly k elements by repeatedly adding a minimal element of its complement. Call it A . Define the opposite-color cutoffs by

$$e_v = \max\{a_v - 1, (1 - s_v) - 2\}.$$

Both vectors inside the maximum are edgewise 1-Lipschitz with the required parity; their maximum has the same properties. Opposite parity on an edge makes its height difference exactly one. The cutoffs obey both physical boundaries and therefore define a pyramid E . Its exact cardinality loss is the number of occupied fibers with $s_v = 0$. Since $k > B_{Q,p}$, all these fibers are occupied, giving loss $|\chi^{-1}(p)|$. The other term in the maximum represents an empty fiber and contributes no actual room. Every neighbor of an actual room of E has height at most the appropriate cutoff of A , proving $N(E) \subseteq A$ even at the physical ends. Movement therefore enters A , and inspecting $A \setminus R$ leaves only rooms of R . $\square$

The all-fiber threshold H_Q is sharp across the two colors when $n \geq \text{diam}(Q)$: at a maximizing vertex v , choose $s_v = 1$ and heights $a_u = \text{dist}(v, u) - 1$. Bipartiteness makes adjacent distances differ by one, giving a valid pyramid of size H_Q with an empty fiber. For a path of order a ,

$$H_{P_a} = \lfloor a^2/4 \rfloor.$$

Indeed the distance sum is maximized at an endpoint, by the nondecreasing increments of $\sum_{j=1}^t \lfloor j/2 \rfloor$. The root threshold $B_{Q,p}$ is separately sharp when $n \geq \max\{1, \text{diam}(Q)\}$: at a maximizing vertex v of color p , the cutoffs $a_u = \text{dist}(v, u) - 2$ give a pyramid of size $B_{Q,p}$ with an empty bottom root. For paths the sharper thresholds are

$$B_{P_{2b},0} = B_{P_{2b},1} = b(b-1), \quad B_{P_{2b+1},0} = b^2, \quad B_{P_{2b+1},1} = b(b-1),$$

where color zero contains the endpoints of the odd path. Indeed, at vertex v of a path indexed from zero, the floor-distance sum is $\lfloor v^2/4 \rfloor + \lfloor (a-1-v)^2/4 \rfloor$; convexity maximizes it at an extreme vertex of the required parity. Taking $k = |R| + m$ gives the usual backward step whenever $m > B_{Q,p}$ and the requested envelope fits; the additional admission condition for a nonpyramidal R is exactly $|\text{cl}(R)| - |R| \leq m$.

11.2 The exact large-budget formula

Theorem 11.3 (Every even width above a quadratic budget). *Let $\max\{b+1, b(b-1)+1\} \leq m < h$, and write*

$$D = m - b, \quad h - b = qD + r, \quad 0 \leq r < D.$$

Set $J(0) = 0$ and $J(r) = r + \min\{\lceil \sqrt{r} \rceil, b\}$ for $r > 0$. Then

$$T_m(G) = \begin{cases} 2q, & r = 0, \\ 2q + 1, & r > 0 \text{ and } 2J(r) \leq m, \\ 2q + 2, & \text{otherwise.} \end{cases}$$

For $h \leq m < 2h$ the answer is two days, and for $m \geq 2h$ it is one. The constructions are physical searches; the lower bounds permit arbitrary interleaving and arbitrary support shapes.




Figure 4: An optimal seven-day search on 6×9 with ten daily inspections, constructed backwards from the final survivor envelope. Every displayed possible-position set and inspection has been replayed on the physical room graph.

Lower bound. Because $m > b(b-1)$, the reflected term in the full-budget scalar recurrence is saturated. Thus $c_{t+1} = L((c_t - m)_+)$, where $L(0) = 0$ and $L(k) = k + \min\{R(k), b\}$ for $k > 0$. Now $D \geq (b-1)^2$. When $D > (b-1)^2$ or $r > 0$, successive plateau steps give

$$c_t = h - tD \quad (0 \leq t \leq q-1), \quad c_q = J(r).$$

If $r = 0$, then $\tau = q$ and $c_{q-1} = m$. If $r > 0$, then $0 < J(r) < m$, so $\tau = q + 1$. The only new endpoint is $D = (b-1)^2$, $r = 0$, $b \geq 2$. Here $q \geq 2$, and the penultimate plateau step has survivor $D = (b-1)^2$, whose surplus is $b-1$. Consequently $c_{q-1} = m-1$ and $\tau = q$. Since $m \geq 3$, one still has $2c_{q-1} > m$, giving the same answer $2q$. For $b = 1$, $D \geq 1 > (b-1)^2$ already covers every feasible budget. Lemma 11.1 gives precisely the claimed lower bounds. Its argument also covers the two extreme budget ranges. $\square$

Physical upper construction. Apply Lemma 11.2 with $Q = P_{2b}$. Its root threshold is $B_{Q,p} = b(b-1)$, and erosion loses exactly b rooms in either color. The prescribed-size enlargement and physical neighborhood inclusion are supplied by that common lemma.

Choose a terminal pyramid R of size r with $|N(R)| \leq J(r)$. The small-corner construction, ordered toward the longitudinal near edge as above, supplies it when $\lceil \sqrt{r} \rceil < b$, and the plateau bound supplies it otherwise. Reflection in the even short axis permits either physical color. Enlarge R by m vertices and apply the clipped erosion. This constructs a preceding survivor of size $r + D$, whose neighborhood lies inside that enlarged envelope. Repeat $q - 1$ times, and enlarge by m once more. This first envelope has size $r + (q - 1)D + m = h$, so is the full color class. Every enlargement fits, and its size is at least $m > b(b - 1)$, giving the required erosion loss even if some other fibers are empty.

Read the construction forward. On each day inspect the difference between its envelope and prescribed survivor, using exactly m designated rooms. The actual belief stays inside the next envelope. After q days its survivor lies in R and its next belief has at most $J(r)$ rooms. When $r = 0$, this is a q -day solo capture; otherwise one more day suffices.

For $r > 0$ and $2J(r) \leq m$, join one such q -day preparation and the time reversal of another on opposite cohorts. Inspect both final neighborhood envelopes on the central day. Their colors are opposite and their total size is at most $2J(r)$, giving $2q + 1$ days. Otherwise concatenate two $(q + 1)$ -day solo searches. For $r = 0$ concatenate two q -day searches. Reflection lets each half choose its required physical color independently. Finally, when $m \geq h$, two inspections of one fixed physical color capture both initial cohorts; inspecting the entire board gives one day when $m \geq 2h$. $\square$

12 Exact symbolic transitions between physical fronts

The next results evaluate transitions between prescribed fronts, rather than assuming that every optimal intermediate support has this form. They include exact one-day costs at the physical ends and arbitrary prescribed daily quotas in the interior. A later application eliminates the translation coordinate from the finite-interface graph.

Let Q be connected and bipartite, with $A \geq 2$ vertices and bipartition map χ . For a downward pyramid of physical color p in $Q \square P_n$, write $s_v = (p - \chi(v)) \bmod 2$. Its virtual last occupied heights satisfy

$$a_v = s_v - 2 + 2k_v, \quad |a_u - a_v| = 1 \quad (uv \in E(Q)),$$

where k_v is the fiber cardinality. The values $-2, -1$ denote empty fibers. This is the height representation established in Lemma 11.2. All the results have reflected versions for upward pyramids.

Proposition 12.1 (Maximal erosion and exact one-day costs). *Let B be a downward pyramid of color p , with heights b_v . Its maximal graph erosion in the opposite color,*

$$E(B) = \{x \in V_{1-p} : N(\{x\}) \subseteq B\},$$

is a pyramid with heights

$$e_v = \max\{b_v - 1, (1 - s_v) - 2\} + 2\mathbf{1}_{\{T_v\}}, \quad (49)$$

where T_v means that $n - 1 \equiv 1 - s_v \pmod{2}$, $b_v = n - 2$, and $b_u = n - 1$ for every neighbor u of v in Q .

For any downward pyramid P of the opposite color, with heights a_v , the minimum useful inspections forcing the next belief to be contained in B are exactly

$$|P \setminus E(B)| = \frac{1}{2} \sum_v (a_v - e_v)_+. \quad (50)$$

The exact next belief B is attainable if and only if $N(P \cap E(B)) = B$; when it is, the same cost is minimal. These minima allow arbitrary competing survivors.

Proof. A candidate room below the top has an upward longitudinal neighbor, so its height is at most $b_v - 1$. The transverse inequalities add nothing: adjacent cutoffs are $b_v - 1$ or $b_v + 1$. The lower maximum in (49) represents an empty opposite-color fiber when no room is admissible. At the top, the upward neighbor is absent. That one additional room is admissible exactly under T_v , by its downward and transverse neighbors. This also handles $n = 1$, when the condition reduces to the transverse neighbor test.

The bottom correction raises a local minimum -3 to -1 , with neighboring new cutoffs -2 . A top correction raises a local minimum $n - 3$ to $n - 1$, with neighboring new cutoffs $n - 2$. Thus edge differences remain one, proving that the erosion is a pyramid. Its intersection with P has cutoffs $\min(a_v, e_v)$.

A survivor has neighborhood contained in B exactly when it is contained in $E(B)$. The largest admissible survivor is therefore $P \cap E(B)$, proving the minimum cost. If any survivor has neighborhood exactly B , its inclusion in this largest survivor forces $B \subseteq N(P \cap E(B)) \subseteq B$. The converse is immediate. $\square$

Lemma 12.2 (Exact height descent with a quota word). *Let a, z be integer height configurations on Q , each with edge differences one and checkerboard vertex parities. Fix $t \geq 0$ and nonnegative integer quotas $p_1, \dots, p_t$. In the virtual height model, a day consists of at most p_j legal local-maximum lowerings by two, followed by adding one to every height. The endpoint z is reachable from a in exactly t days if and only if*

$$z_v \equiv a_v + t \pmod{2}, \quad z_v \leq a_v + t \quad \text{for every } v, \quad (51)$$

$$F := \frac{\sum_v a_v - \sum_v z_v + At}{2} \leq \sum_{j=1}^t p_j.$$

Exactly F lowerings suffice. The same criterion holds with movement before each day's lowerings.

Proof. Each lowering reduces the height sum by two, while movement raises each coordinate by one. This proves necessity, including the componentwise and parity conditions.

Put $c = z - t\mathbf{1}$. It has the same vertexwise parities as a and satisfies $c \leq a$. From any current $x \neq c$, choose a vertex v of maximum current height among those with $x_v > c_v$. A higher neighbor u would also be above its target: otherwise $c_u = x_u = x_v + 1$ and $c_v \leq x_v - 2$, contradicting the target edge condition. Such a neighbor contradicts maximality, so all neighbors are one lower. Lowering x_v by two is legal and preserves $x \geq c$. Repeated descent reaches c in exactly F steps.

Choose integers $0 \leq q_j \leq p_j$ summing to F and divide this word into successive pieces of lengths q_j . Uniform additions commute with the local-maximum test. Inserting one addition per piece, before or after it, gives endpoint $c + t\mathbf{1} = z$. The case $t = 0$ forces $a = z$ and uses the empty word. $\square$

The following physical lower bound is what permits comparison with arbitrary, possibly non-pyramidal competitors. Suppose Q has a Hamiltonian path, indexed $0, \dots, A-1$, and take its parity as χ . All color indices in the following formulas are read modulo two. Put

$$C_A = \left\lfloor (A-2)^2/4 \right\rfloor, \quad \gamma_p = |\chi^{-1}(1-p)|.$$

For any finite color- p support S in $Q \square P_\infty$,

$$|N(S)| \geq |S|, \quad |S| > C_A \implies |N(S)| \geq |S| + \gamma_p. \quad (52)$$

Indeed, deleting the extra transverse edges leaves the spanning rectangle. Put the finite support in a sufficiently long auxiliary rectangle so that its neighborhood is unchanged and the far profile branch cannot lower the cap. For $A = 2b$, the proved profile becomes $k + \min(\lceil \sqrt{k} \rceil, b)$, whose plateau begins at $k > (b-1)^2$. For $A = 2b+1$, choose the auxiliary length odd. Its two profiles are $k + \min(R(k), b)$ and $k + \min(Q(k), b+1)$, with their common plateau beginning at $k > b(b-1)$. These are exactly the stated C_A and γ_p . Their small branches also prove nonnegative expansion.

Theorem 12.3 (Unrestricted physical transitions in a guarded interior). *Let Q be bipartite with a Hamiltonian path and $A \geq 2$. Let P, Z be downward pyramids in $Q \square P_n$, with heights a, z . Fix $t \geq 1$, quotas $p_1, \dots, p_t$, and $P_* := \sum_j p_j$. Assume*

$$t \leq \min_v a_v, \quad t \leq \min_v z_v, \quad \max_v a_v, \max_v z_v \leq n-1-t, \quad (53)$$

$$|P| - P_* > C_A.$$

An actual t -day block with these daily quotas and exact endpoint Z exists, without restrictions on intermediate supports, if and only if (51) holds. Whenever it exists, there is such a block with every survivor and belief a downward pyramid. Both longitudinal parities and arbitrary partial initial pyramids are included.

Proof. Free evolution from P has heights $a+j$ for $0 \leq j \leq t$ under the height margins. Every competitor is contained in this free belief, so the endpoint inclusion and parity conditions are necessary. Before each movement no survivor meets the far longitudinal row; hence its neighborhood agrees with that in the half-strip.

Nonnegative expansion in (52) bounds every survivor below by $|P| - P_* > C_A$. If the initial color is p , its j th expansion is therefore at least γ_{p+j-1} . For $s_v^p = (p - \chi(v)) \bmod 2$,

$$\sum_{j=1}^t \gamma_{p+j-1} = \frac{At + \sum_v s_v^p - \sum_v s_v^{p+t}}{2}, \quad |P| - |Z| = \frac{\sum_v a_v - \sum_v z_v - \sum_v s_v^p + \sum_v s_v^{p+t}}{2}.$$

Summing the useful inspections and canceling these parity terms gives $F \leq P_*$. This lower bound did not constrain the intermediate shapes.

For sufficiency use the descent in Lemma 12.2. Its unshifted heights stay between $z-t$ and a . The margins keep every surviving fiber nonempty and, before a movement, its height at most $n-2$. Its actual neighborhood consequently has cutoffs exactly one higher. Each lowering removes one actual top room; successive lowerings of the same fiber remove distinct rooms. Thus the virtual construction is an exact physical block with the prescribed daily quotas. $\square$

The cardinality guard is only a lower-bound hypothesis. The construction itself extends to arbitrarily long durations with fixed physical margins. Write $r = \text{diam}(Q)$ and $\mu(x) = A^{-1} \sum_v x_v$.

Proposition 12.4 (Balanced inspections in fixed physical margins). *Let Q be connected and bipartite with $A \geq 2$, and suppose the criterion (51) holds for a constant quota m and $t \geq 1$. If*

$$r + 2 \leq \mu(a), \mu(z) \leq n - r - 3, \quad (54)$$

there is an exact physical pyramid block from a to z using at most m inspections each day, regardless of its duration.

Proof. Use the legal descent word, with piece lengths

$$q_j = \lfloor jF/t \rfloor - \lfloor (j-1)F/t \rfloor \leq m.$$

The completed-day mean is

$$\mu(a) + j - \frac{2}{A} \lfloor jF/t \rfloor = (1 - j/t)\mu(a) + (j/t)\mu(z) + \frac{2}{A} \{jF/t\}.$$

Here braces denote the fractional part. It lies between the smaller endpoint mean and the larger plus $2/A$. The survivor immediately before a movement is one lower. Each partial inspection step lies coordinatewise between that day's source and survivor, and every height configuration has range at most r . All partial heights therefore lie in

$$[\min(\mu(a), \mu(z)) - r - 1, \max(\mu(a), \mu(z)) + r + 2/A].$$

The margin keeps them in the nonempty, unclipped physical range. Neighborhoods are exactly the global additions, completing the proof. $\square$

For $m > A/2$, define $D = 2m - A$. Between prescribed height endpoints, the least virtual duration is the smallest nonnegative integer t_0 of the required color parity satisfying

$$t_0 \geq \max \left\{ 0, \max_v (z_v - a_v), \left\lceil \frac{\sum_v a_v - \sum_v z_v}{D} \right\rceil \right\}. \quad (55)$$

Thus two maxima, an integer ceiling, and a parity adjustment give its cost. Under (54) the corresponding block is physical. An unrestricted lower bound for a long connection requires control of possible boundary excursions; it does not follow by applying (53) with a large mt_0 . The finite ports below handle those excursions explicitly.

13 Odd widths and even lengths above a quadratic budget

The remaining parity of rectangular boards requires remembering the orientation of an efficient boundary. The resulting formula has the same corner cost as the odd-board formula, but no middle-day parity penalty. Throughout this section let

$$w = 2b + 1 \geq 3, \quad n = 2\ell \geq w, \quad h = wn/2, \quad B = b(b+1).$$

Use L_0, L_1, J, F from (32), (33), and (36), with the present b, m .

Theorem 13.1 (Odd width, even length, every sufficiently large budget). *Suppose $m \geq b(b+1)+1$. For $m < h$, write*

$$D = 2m - w, \quad 2h - w - 1 = qD + r, \quad 0 \leq r < D.$$

Then

$$T_m(P_w \square P_n) = \begin{cases} 2q, & r = 0, \\ 2q + 1, & r > 0 \text{ and } 2J(r) \leq m, \\ 2q + 2, & r > 0 \text{ and } 2J(r) > m. \end{cases} \quad (56)$$

For $h \leq m < 2h$ the answer is two days, and for $m \geq 2h$ it is one. The upper bounds are explicit physical searches. The lower bounds allow arbitrary inspection locations and arbitrary interleaving of the cohorts.

13.1 A sharp orientation constraint

Call a survivor *critical* when its surplus is b , and *strictly middle* when its cardinality s satisfies

$$b^2 < s < h - B.$$

Match longitudinal rooms $(x, 2j)$ and $(x, 2j + 1)$, and identify a color- p room with (x, j) by $y = 2j + ((p - x) \bmod 2)$. The matching-contracted graph has bidirectional edges in x and longitudinal edges pointing toward $j = 0$ on even x , toward $j = \ell - 1$ on odd x , when $p = 0$; phase one reverses the directions. Put $U = \{0, 2, \dots, 2b\}$ and $V = \{1, 3, \dots, 2b - 1\}$.

Lemma 13.2 (Critical orientation and a corner consequence). *A phase-zero strictly middle critical survivor consists of prefixes of lengths t_x with*

$$1 \leq t_u \leq \ell - 1 \quad (u \in U), \quad t_v \leq t_u \leq t_v + 1 \quad (uv \text{ an edge}).$$

Phase one gives reflected suffixes. Two successive strictly middle survivors cannot both be critical. After a strictly middle critical survivor, any next survivor of size $s \leq b^2$ has at least $L_1(s)$ neighbors.

More generally, a survivor of size $h - x$, $x \leq B$, and surplus $a \leq b$ can precede a strictly middle critical survivor only if $x \leq a^2$.

Proof. Let H be the outside directed boundary, of size b . If two adjacent x -columns avoid H , they form a strongly connected ladder, and all horizontal rows avoiding H belong to one common component. The anchored-area argument of (46) still bounds a set outside that component by b^2 . At this equality threshold the only new issue is a horizontal matching exception: an occupied row could contain the entire majority parity U . In that case every one of the b minority columns has a selected or boundary vertex. Following each minority column toward an empty boundary-free anchor forces a boundary vertex in that same column. All b boundary vertices would therefore mark exactly V , contradicting the assumed adjacent unmarked pair. The exception is excluded. The original weighted row sum now gives area at most b^2 ; applied to the transposed complement it gives deficiency at most $b^2 + b$. Both contradict the strict-middle inequalities.

Thus no adjacent columns are unmarked. With only b marked columns on a path of $2b + 1$ positions, they must be exactly V , with one boundary vertex (v, c_v) on each. In phase zero every U fiber is a prefix, and directed closure gives

$$R_v \subseteq R_u, \quad R_u \setminus R_v \subseteq \{c_v\} \quad (uv \text{ an edge}).$$

If fiber $2j \in U$ were empty, summing the neighboring length bounds would give $|R| \leq j^2 + (b - j)^2 \leq b^2$. Applying the same argument to the transposed complement shows that a full U fiber forces deficiency at most B . Thus every U fiber is nonempty and nonfull. A V fiber cannot contain a

point above c_v : its forward chain would reach the missing far endpoint without another boundary vertex. Hence it too is a prefix, with the asserted adjacent lengths.

Its neighborhood leaves the U fibers unchanged, so misses their far endpoints. An opposite-phase strictly middle critical suffix contains all those endpoints, proving incompatibility. Physically the next belief omits an entire longitudinal endpoint row. Its next survivor lies in the minority class of the odd rectangle obtained by removing that row. Theorem 5.1 gives $L_1(s)$ for $s \leq b^2$; its opposite-corner term cannot improve the bound.

For the final assertion set $A = \overline{N(R)}$. Then $|A| = x - a$ and $N(A) \subseteq \overline{R}$, so its surplus is at most a . If the next survivor is strictly middle critical, it contains the favored endpoint row, and A omits it. In the smaller odd rectangle the minority capacity is $h - b - 1 \geq 2B$, whereas $|A| \leq B - a$. Its reflected profile term is at least $b + 1$. Thus surplus at most $a \leq b$ forces $|A| \leq a(a - 1)$, or $x \leq a^2$. $\square$

13.2 The physical construction

Lemma 11.2 with $Q = P_{2b+1}$ has threshold $H_Q = B$. A physical-color- p envelope A larger than B erodes to an opposite-phase pyramid E with

$$N(E) \subseteq A, \quad |E| = |A| - (b + 1 - p). \quad (57)$$

Choose a phase- p terminal survivor R of rank $2|R| + p = r + 1$: use $(|R|, p) = (r/2, 1)$ for even r , and $((r + 1)/2, 0)$ for odd r . Quadrant prefixes and the pyramid plateau bound provide such an R with $|N(R)| \leq J(r)$. For $r = 0$ use the empty phase-one survivor.

Enlarge R by m vertices inside the finite predecessor ideal, then apply (57). This backward step raises its rank by exactly D ; every enlarged set has at least $m > B$ vertices. After $q - 1$ steps the first survivor has rank $r + 1 + (q - 1)D = 2h - 2m$ and phase zero, since q and r have the same parity. One last enlargement by m gives the full color class. Reading the nested envelopes forwards gives q inspection days ending, after movement, in at most $J(r) \leq m$ rooms. All requested cardinalities are at most h , since the backward ranks increase to $2h - 2m$.

Reflection in the even longitudinal side lets either initial physical cohort use this favorable relative phase independently. A forward q -day half-search and a reversed half-search meet on one middle day; their endpoint envelopes have opposite colors and total size at most $2J(r)$. Inspect their union when $2J(r) \leq m$. Otherwise concatenate two $(q + 1)$ -day solo searches. For $r = 0$, two q -day searches suffice. Walk reversal justifies the backward half even when its terminal set is only a containing envelope. This proves all upper bounds in (56).

13.3 A corner potential with at most one lost inspection

Put $N = 2h$ and retain the function F of (36). Define

$$C = \min_{0 \leq x \leq B} \{x + F(2(h - x + \rho(x)))\}. \quad (58)$$

The following arithmetic records precisely how much the far corner can improve the potential.

Lemma 13.3 (Corner estimates). *One has $F(N) - 1 \leq C \leq F(N)$, with equality $C = F(N)$ if $r = 0$. For $0 \leq a \leq b$, $0 \leq x \leq a(a + 1)$, and $U \in \{N, N - 1\}$,*

$$x \leq a^2 \implies x + F(U - 2(x - a)) \geq F(U), \quad (59)$$

$$x + F(U - 2(x - a) + 1) \geq F(U). \quad (60)$$

In particular $C \geq F(N - 1)$.

Proof. Cases with $x \leq a$, or a source in the cardinality branch, follow directly from monotonicity and $F(v) \geq \lfloor v/2 \rfloor$. Otherwise put $t = x - a > 0$, $d = b - a$. Since $2t \leq 2b^2 < D$, a target crosses at most one period boundary. If it enters the final cardinality branch, replacing its actual value by the periodic expression only lowers it, as in Lemma 8.3. Write $\rho_p(v) = L_p(v) - v$, with $\rho_p(0) = 0$.

Without a crossing, an even remainder $2k$ in the uncorrected cost has possible discount $\rho_1(k) - \rho_1(k - t) - a$. If $k - t > 0$, square/pronic subadditivity makes this nonpositive. If $k = t$, it is at most one since $Q(t) \leq a + 1$; when $x \leq a^2$, the stronger $t \leq a(a - 1)$ makes it nonpositive. An odd remainder gives $\rho_0(k + 1) - \rho_0(k - t + 1) - a \leq 0$ by square-root subadditivity. For the corrected cost, put $v = k - t + 1 > 0$. The two discounts are

$$\rho_1(k) - a - 1 - \rho_0(v), \quad \rho_0(k + 1) - a - \rho_1(v),$$

both nonpositive: use $k \leq R(v)^2 + a^2 - 1$ and $k + 1 \leq Q(v)(Q(v) - 1) + a^2$, respectively. If a root is clipped, either the same estimate applies or the cap makes the claim immediate.

At a crossing set

$$\lambda = m - b + k - x + a \geq b^2 - a^2 + k + 1.$$

For even remainder $2k > 0$ put $c = \rho_1(k) \leq a + 1$; for odd remainder $2k + 1$ put $c = \rho_0(k + 1) \leq a$. The uncorrected discounts are respectively $d + c - \rho_0(\lambda)$ and $d + 1 + c - \rho_1(\lambda)$. They are at most one, because

$$\begin{aligned} \lambda - (d + c - 2)^2 &\geq 2d(a - c + 2) + c > 0, \\ \lambda - (d + c - 1)(d + c - 2) &\geq 2d(a - c) + 3d + c > 0. \end{aligned}$$

These follow from $k \geq (c - 1)(c - 2) + 1$ and $k \geq (c - 1)^2$, respectively. When $k = 0$ in the even case, $\lambda > d^2$ instead makes the discount negative; this also proves $C = F(N)$ for remainder zero.

Under $x \leq a^2$, the lower bound on λ gains a , and the stronger estimates are

$$\begin{aligned} \lambda - (d + c - 1)^2 &\geq 2d(a - c + 1) + a - c + 3 > 0, \\ \lambda - (d + c)(d + c - 1) &\geq 2d(a - c) + d + a - c + 2 > 0. \end{aligned}$$

Here $c \leq a$ in both rows. They prove (59). For the corrected target the discounts are $d + c - \rho_1(\lambda)$ and $d + c - \rho_0(\lambda + 1)$, controlled by

$$\begin{aligned} \lambda - (d + c - 1)(d + c - 2) &\geq 2d(a - c) + 3d + 2 > 0, \\ \lambda + 1 - (d + c - 1)^2 &\geq 2d(a - c + 1) + 2 > 0. \end{aligned}$$

The required root thresholds do not exceed their caps $b + 1$ and b . Again $k = 0$ in the even case follows directly from $\lambda > d(d - 1)$. A corrected target exactly at a period boundary uses $J(D) = m$, equal to the next period's zero remainder. This proves (60). Finally apply it at $U = N - 1$, with $a = \rho(x)$, to obtain $C \geq F(N - 1)$; $x = 0$ gives $C \leq F(N)$. $\square$

Initially mark both cohorts zero. Subsequently mark a belief by $z = 1$ exactly when its preceding survivor was strictly middle critical. Give a belief of size k potential

$$\Phi_z(k) = \min\{F(2k + z), \max(0, C - h + k)\}, \quad \Phi_z(0) = 0.$$

Every physical transition lowers this potential by at most its useful inspection allocation. For a survivor below $h - B$, this follows from Lemma 8.3 and Lemma 13.2: a critical middle move has marks $0 \rightarrow 1$, a larger surplus dominates the alternating rank step, and a small survivor uses L_0 or L_1 according to the source mark. The small target counts are at most $b^2 + b + 1 \leq m$, so their two ranks have the same cardinality potential. Sources of size at most m are handled directly by matching. For a co-small survivor with deficiency x , the definition of C handles the target's F branch. Matching makes deficits grow by at most the useful allocation, handling the linear branch. Thus

$$T_m \geq \lceil 2C/m \rceil. \quad (61)$$

Since $F(N) = qm + J(r)$, the corner estimates match the upper bound except possibly when

$$r > 0, \quad J(r) = \lfloor m/2 \rfloor + 1, \quad C = F(N) - 1. \quad (62)$$

13.4 Separating the two corners on longer boards

Suppose

$$h > m + 2b^2 + b. \quad (63)$$

It suffices to settle (62). Now mark a belief if its preceding survivor was either strictly middle critical, or co-small with deficiency x , surplus $a \leq b$, and $x > a^2$. Use the common-cap potential

$$\Psi_z(k) = \min\{F(N), F(2k + z)\}, \quad \Psi_z(0) = 0.$$

The new mark also forbids a critical next middle move by Lemma 13.2. For a co-small target, the corner estimates give

$$x + F(\text{target rank}) \geq F(N) :$$

use (59) when $x \leq a^2$ and (60) otherwise. Surplus greater than b already gives the interior rank inequality.

For a source deficiency d , monotonicity by two rank units gives $d + F(N - 2d) \leq F(N)$. A marked source has $d \geq 1$, and (62) together with $C \geq F(N - 1)$ gives

$$d + F(N - 2d + 1) \leq F(N - 1) + 1 \leq F(N).$$

Since the allocation is $x - d$, these estimates prove the potential inequality for every co-small target, including the common cap. For a small target, the old mark uses the conditional L_1 bound. A new co-small mark has source size at least $h - B$, so (63) prevents it from reaching size at most b^2 in one day. All other steps are the interior cases already checked. Both initial potentials are $F(N)$, proving $T_m \geq \lceil 2F(N)/m \rceil = 2q + 2$ in the sole possible gap.

Every $q \geq 4$ satisfies (63), since

$$h \geq 4m - 3b - 1, \quad h - m - 2b^2 - b \geq b^2 - b + 2 > 0.$$

Also $q \geq 1$ because $h > m$. For $q = 1$, write $r = 2s - 1$; then $h = m + s$ and $F(N) = h + \min(R(s), b)$. A corner departure with $t = x - a \leq a^2$ costs $h + a$ when $t \geq s$, and $h + a + \min(R(s - t), b)$ otherwise. Square-root subadditivity makes both at least $F(N)$, so $C = F(N)$ and there is no gap. For $q = 3$ in (62), write $r = 2k + 1$ and $c = \min(R(k + 1), b)$. Then $k + c = \lfloor m/2 \rfloor$, so $k \geq b(b - 1)/2$, while $h = 3m - 2b + k$. Hence

$$h - m - 2b^2 - b \geq k - b + 2 \geq (b^2 - 3b + 4)/2 > 0.$$

Only $q = 2$ still requires an argument.

13.5 A two-day bound closes the short-board case

Let $q = 2$, $r = 2k > 0$, and $c = \min(Q(k), b + 1)$. Thus $h = 2m - b + k$ and $J(r) = k + c$. We claim that after any two inspections and their following moves the *total* deficit of the two cohorts is at most

$$D_* = h - k - c. \quad (64)$$

Use the symmetric inverse profile

$$I(z) = z - \min\{\rho(z), b, R(h - z)\}, \quad 0 \leq z \leq h,$$

whose superadditivity was proved in Lemma 11.1. Since $h - m = m - b + k > b^2$, the first-day reflected term exceeds b . Unless all m useful first-day inspections go to one cohort and its survivor has surplus exactly b , the first-day total deficit is at most $m - b - 1$. Indeed a nontrivial split would otherwise have positive costs a, a' with $a + a' \leq b$, forcing $m \leq a(a + 1) + a'(a' + 1) \leq B$, a contradiction. Higher actual surplus only strengthens this conclusion.

In these cases superadditivity bounds the second-day total deficit by $I(Z)$, where $Z = 2m - b - 1 < h$. As $Z > B$ and $h - Z = k + 1$,

$$I(Z) = Z - \min(b, R(k + 1)) \leq h - k - c,$$

using $1 + \min(b, R(k + 1)) \geq c$.

In the exceptional pure first day, the survivor of size $h - m$ is strictly middle critical. The active next belief has $m + k$ rooms, and its mate is full. Allocate $m - u$ to the active cohort and u to its mate on day two, leaving active survivor size $s = k + u$. If $s < h - B$, Lemma 13.2 gives at least $L_1(s)$ neighbors, both in the small and middle ranges. The mate's deficit is at most u , so their total deficit is at most $h - L_1(k + u) + u \leq h - k - c$. If $s \geq h - B$, then $u \geq 2m - b - B > m - b \geq b^2 + 1$, and the inverse profile gives $I(u) = u - b$. The active survivor is nonempty and proper; the exact static profile therefore gives at least $s + 1$ neighbors. The total deficit is at most $h - k - b - 1 \leq h - k - c$. This proves (64). Unused quotas may be filled monotonically, so the argument includes all schedules with budget at most m .

Finally, a five-day capture would have two forward days, one middle inspection, and two reversed days. By (64), the two possible-position sets at the middle inspection intersect in at least $2h - 2D_* = 2J(r)$ cells. Every cell in this intersection supports an avoiding walk on each side; their concatenation must be intercepted by the middle inspection. If $2J(r) > m$, this is impossible, giving the required sixth day. All remaining cases already match (61). The one- and two-day regimes follow by inspecting whole color classes. This completes Theorem 13.1.

14 An exact isoperimetric order in every all-odd box

The preceding rectangle theorem extends to every dimension. The extension uses a weighted exchange argument: compression in lower-dimensional faces first makes neighborhood size a sum of local costs, and an order comparison then identifies exchanges that cannot increase that sum. Parity-restricted isoperimetry on hypercubes was studied by Körner and Wei [12], and local-to-global methods for Cartesian-product vertex boundaries by Bezrukov and Serra [4]. Here the neighborhood is the open neighborhood of a single color class. We give the face compressions and their exchange argument explicitly, rather than infer the result from a theorem about closed vertex boundaries.

Delete any factors of side length one, and write the remaining box as

$$B = \prod_{i=1}^d \{0, \dots, L_i\}, \quad 2 \leq L_1 \leq \dots \leq L_d, \quad L_i \text{ even.}$$

Thus coordinates are ordered from the shortest side to the longest. Let B_p be its color- p class, where the color of x is the parity of $|x| = \sum_i x_i$. Within either color, order vertices by increasing $|x|$ and then by *decreasing* lexicographic coordinates. Write $x \prec y$ for this order and $I_p(k)$ for its first k vertices. The direction of the lexicographic tie is part of the definition.

Theorem 14.1 (Compatible isoperimetric nesting in all-odd boxes). *For every $S \subseteq B_p$,*

$$|N(S)| \geq |N(I_p(|S|))|.$$

Moreover, $N(I_p(k))$ is a prefix in the opposite color. Consequently the minimum capture time for every inspection budget on B is given by an exact recurrence on $O(|B|^2)$ states, and optimal inspection sets can be recovered from that recurrence.

For $d = 1$ the minimizing order is the spacing-two prefix order of Lemma 3.4. For $d = 2$, putting the shorter coordinate first and ordering it decreasingly within a weight layer is equivalent to the increasing-long-coordinate convention of Theorem 5.1. These provide the induction bases. In particular, the higher-dimensional assertion does not assume a general Cartesian-product isoperimetric principle.

14.1 Prefix neighborhoods and a local cost identity

For a nonzero vertex z , let $j(z)$ be its last positive coordinate and put

$$\ell(z) = z - e_{j(z)}.$$

This is its earliest lower neighbor in decreasing lexicographic order.

Lemma 14.2 (Nesting of prefix neighborhoods). *For the stated weight and lexicographic order, the neighborhood of a prefix is a prefix of the opposite color. This statement holds even without the assumptions that the side lengths are odd and sorted.*

Proof. Within a fixed weight layer, ℓ preserves order, allowing ties. Indeed, suppose $z \prec w$ first differ in coordinate i , so $z_i > w_i$. Equal weights imply that w has a positive coordinate after i . If z does too, neither predecessor operation changes the first difference. Otherwise $j(z) = i$ and $z_i - 1 \geq w_i$. A strict inequality preserves the order. In the equality case the tail of w has total weight one; removing its unique unit gives $\ell(z) = \ell(w)$.

Consider a nonempty prefix whose last, possibly partial, layer has weight r . Every smaller opposite-color layer is completely covered: its nonzero vertices have lower neighbors in complete source layers. When weight zero belongs to the opposite color, it is covered by the first odd layer. No vertex above layer $r + 1$ is covered. On layer $r + 1$, a vertex is covered precisely when its earliest lower neighbor belongs to the selected prefix in layer r . The order preservation just proved makes these covered vertices an initial interval of layer $r + 1$. The empty prefix is immediate. $\square$

Call a set *pair-compressed* if, after fixing all but any two coordinates, its selected vertices in either free-coordinate parity form a prefix of the inherited order. Define

$$a(0) = d, \quad a(x) = d - j(x) + 1 - \mathbf{1}_{\{x_{j(x)} = L_{j(x)}\}} \quad (x \neq 0). \quad (65)$$

Lemma 14.3 (Neighborhood cost of a pair-compressed set). *If $d \geq 2$ and $\emptyset \neq S \subseteq B_p$ is pair-compressed, then*

$$|N(S)| = \sum_{x \in S} a(x) + \mathbf{1}_{\{p=1\}}. \quad (66)$$

Proof. For every nonzero z , we claim

$$z \in N(S) \iff \ell(z) \in S.$$

Only the forward implication needs proof. Choose a neighbor $x \in S$ of z . The vertices x and $\ell(z)$ differ in at most two coordinates. If x is a lower neighbor, then $\ell(z)$ is no later than x in their common weight layer. If x is an upper neighbor, $\ell(z)$ has weight two less than x . Pair compression therefore puts $\ell(z)$ in S . If only one coordinate differs, any second coordinate can be used; one exists because $d \geq 2$.

For $x \neq 0$ with last positive coordinate J , the preimages of x under ℓ are obtained by adding one in coordinate J , if it is not full, or by adding one in any coordinate after J . Their number is $a(x)$. The origin has the d unit vectors as preimages. Thus the sum in (66) counts every nonzero neighbor exactly once.

If $p = 1$, a nonempty pair-compressed set contains a unit vector. Starting with any of its vertices of weight greater than one, decrease by two in a coordinate that is at least two, or decrease by one in two positive coordinates. Each move stays in a pair fiber and goes earlier in its order, so it preserves membership. The process ends at weight one. The origin is therefore an additional neighbor when $p = 1$, and it can never be a neighbor when $p = 0$. $\square$

14.2 Face compression and its partial order

Suppose $d \geq 3$ and the isoperimetric theorem has been proved in dimension $d - 1$. On each fixed color define a partial order $\leq_P$ as the reflexive transitive closure of

$$x <_P y \quad \text{if } x \prec y \text{ and } x_i = y_i \text{ for at least one } i. \quad (67)$$

All generating relations increase the total order, so this is indeed a partial order. Its ideals are exactly the sets whose restriction to every codimension-one fiber is a prefix.

The inductive isoperimetric theorem and compatible neighborhood nesting make the prefix map on each $(d - 1)$ -dimensional box a strategy compression: it is monotone, preserves cardinality, and sends each color's neighborhood into the corresponding prefix of the original neighborhood's size. The remaining side lengths are still sorted, and their order is the restriction of the global order to a fixed-coordinate fiber. By Lemma 3.3, compressing all such fibers for one coordinate index is a strategy compression of B .

Cycle through the d indices. Every change strictly decreases the sum of the selected vertices' positions in the global weight and lexicographic order. This nonnegative integer is bounded by $|B|(|B| - 1)/2$. The simultaneous-compression argument in Section 3 therefore gives a uniform finite composition whose output is fixed by every face compression. In particular, every $S \subseteq B_p$ can be replaced by a P -ideal of the same size without increasing its neighborhood.

Every P -ideal is pair-compressed. A two-coordinate fiber lies inside a codimension-one fiber obtained by fixing a coordinate outside the pair; such a coordinate exists because $d \geq 3$. The restriction of a prefix to a subfiber is a prefix of its inherited order. Thus Lemma 14.3 applies to the resulting ideals.

14.3 The exchange comparison

The following elementary comparison is the step that links the local cost identity to the global order.

Lemma 14.4 (Terminal-coordinate comparison). *Let $d \geq 3$. If same-parity vertices satisfy $x \prec y$ and $x_d = 0 < y_d$, then $x \leq_P y$. Also, if $x \prec y$ and $x_d < y_d = L_d$, then $x \leq_P y$.*

Proof. First, if $u \leq v$ coordinatewise and their weights have the same parity, then $u \leq_P v$. Perform all required increments of size two within individual coordinates, and then pair the remaining increments of size one. Each step increases weight by two and changes at most two coordinates. Since $d \geq 3$, it leaves a coordinate unchanged and is a generator of (67).

For the first assertion, a coordinate agreement between x and y already gives a generator. Assume henceforth that all coordinates differ. If $|x| < |y|$ and some $i < d$ has $x_i > y_i$, transfer $t = x_i - y_i$ units from coordinate i to coordinate d of x , obtaining z . This is legal because $0 < t \leq L_i \leq L_d$ and $x_d = 0$. We have $x \prec z$ within their common weight layer, while $z \prec y$ follows from $|z| = |x| < |y|$. The vertex z retains $d - 2 \geq 1$ coordinates of x and agrees with y in coordinate i . Hence $x <_P z <_P y$. If there is no such excess coordinate, then $x \leq y$ coordinatewise, and the preceding observation applies.

It remains to consider $|x| = |y|$. Since all coordinates differ, $x_1 > y_1$. If a middle coordinate $2 \leq i \leq d - 1$ has $x_i > y_i$, use the same transfer to coordinate d . Again $x \prec z$, and $z \prec y$ because $z_1 = x_1 > y_1$. The same coordinate agreements give the two generators. Otherwise all middle coordinates increase: $x_i < y_i$ for $2 \leq i \leq d - 1$. Weight equality gives

$$x_1 - y_1 = y_d + \sum_{i=2}^{d-1} (y_i - x_i).$$

Transfer y_d units from the first to the last coordinate of x . The resulting vertex z is valid, since

$$z_1 = y_1 + \sum_{i=2}^{d-1} (y_i - x_i) > y_1 \geq 0, \quad z_d = y_d \leq L_d.$$

It satisfies $x \prec z \prec y$, shares the middle coordinates with x , and shares the last coordinate with y . This proves the first assertion.

The coordinate complement $\rho(x)_i = L_i - x_i$ reverses the weight and lexicographic order, and it preserves coordinate agreements. It therefore reverses the generated partial order. Applying the first assertion to $\rho(y) \prec \rho(x)$ proves the second assertion. $\square$

Lemma 14.5 (Cost comparison). *If $x \prec y$ have the same parity and $a(x) > a(y)$, then $x \leq_P y$.*

Proof. If $x_d = y_d$, they share a coordinate. If $x_d = 0 < y_d$, use the first part of Lemma 14.4. The case $y_d = 0 < x_d$ cannot give a strict cost decrease, since $a(x) \leq 1$ and $a(y) \geq 1$. Finally, if both last coordinates are positive, both costs are zero or one. A strict decrease forces $x_d < L_d = y_d$, so the second part of that lemma applies. $\square$

Proof of Theorem 14.1. The path and rectangle bases were noted above. In dimension $d \geq 3$, the inductive face compressions replace an arbitrary $S \subseteq B_p$ by a P -ideal of the same size and with no larger neighborhood. The empty set needs no further argument.

If a nonempty P -ideal is not a global prefix, let x be its earliest missing vertex and y its latest included vertex. Then $x \prec y$. They are incomparable in P : $x \leq_P y$ would contradict ideality, and $y \leq_P x$ would contradict the global order. By Lemma 14.5, $a(x) \leq a(y)$.

Replace y by x . Removing y preserves ideality because every proper successor of y is later in the global order and hence absent. Adding x preserves ideality because every proper predecessor of x is earlier and hence present; y is not such a predecessor. The new set has the same positive cardinality and color. The neighborhood cost identity (66) shows that its neighborhood does not increase. The sum of global positions strictly decreases, so iteration ends at the global prefix of that size. This proves the isoperimetric inequality. Lemma 14.2 gives compatible nesting, completing the induction. $\square$

14.4 Profiles, exact capture times, and optimal inspections

Write $V = |B|$, $E = (V + 1)/2$, and $M = (V - 1)/2$. The theorem supplies the complete profiles by a direct cumulative-sum construction:

$$g_p(0) = 0, \quad g_p(k) = \sum_{x \in I_p(k)} a(x) + \mathbf{1}_{\{p=1\}} \quad (k > 0). \quad (68)$$

For $d = 1$ these expressions also agree with the odd-path profiles. Enumerate the vertices, order each color as specified, evaluate (65), and take cumulative sums. Using a mixed-radix lexicographic index, sorting requires $O(V \log V)$ integer comparisons; computing weights, indices and costs directly requires $O(dV)$ additional operations. This is a constructive profile algorithm, without an isoperimetric optimization subroutine.

Corollary 14.6 (Feasibility by one profile maximum). *For every nontrivial all-odd box, the minimum feasible daily budget is*

$$h(B) = 1 + \max_{0 \leq k \leq E} (g_0(k) - k) = 1 + \max_{0 \leq k \leq E} \sum_{x \in I_0(k)} (a(x) - 1). \quad (69)$$

The one-room box has threshold one by direct inspection.

Proof. We verify the hypotheses of the general hunter-number criterion of Bolkema and Groothuis [5, Theorem 16]. That theorem states that a bipartite graph with compatible isoperimetric nesting has hunter number $1 + \min(u_0, u_1)$ when the maximum neighborhood surpluses $u_p = \max_k (g_p(k) - k)$ differ by at most one.

For completeness, the exact profile duality gives the required relation between these maxima. Define

$$J(z) = \max\{x : 0 \leq x \leq E, g_0(x) \leq z\}, \quad 0 \leq z \leq M.$$

Then

$$g_1(y) = E - J(M - y). \quad (70)$$

Indeed, a majority set of size x with at most $M - y$ neighbors leaves a minority y -set with no edges to it, giving $g_1(y) \leq E - x$. Conversely, the majority complement of the neighborhood of a minimizing minority y -set has size $E - g_1(y)$ and at most $M - y$ neighbors. These two choices prove (70). Since $E = M + 1$, it follows that

$$u_1 = 1 + \max_{0 \leq z \leq M} (z - J(z)).$$

This last maximum equals u_0 . If $J(z) < E$, let $x = J(z) + 1$. Then $g_0(x) > z$, so $z - J(z) \leq g_0(x) - x \leq u_0$. If $J(z) = E$, the absence of isolated vertices implies $g_0(E) = M$; thus $z = M$ and $z - J(z) = -1 \leq u_0$. For the reverse inequality choose a positive x attaining u_0 . Such an x exists even when $u_0 = 0$, because $g_0(1) \geq 1$ then forces $g_0(1) = 1$. For $z = g_0(x) - 1$ we have $0 \leq z < M$ and $J(z) \leq x - 1$, hence $z - J(z) \geq g_0(x) - x = u_0$. Therefore $u_1 = u_0 + 1$.

Theorem 14.1 supplies compatible nesting, so the cited criterion yields $h(B) = 1 + u_0$. Its general lower-bound theorem is being used here; profile duality alone is not a lower-bound proof. The cumulative-sum formula follows from (68). $\square$

The upper bound also admits a direct description. Spend all $m = u_0 + 1$ inspections on one current prefix cohort. If it is not yet captured, a majority step decreases its size by at least one, while a minority step does not increase it, because $u_1 = u_0 + 1$. Thus it clears after finitely many steps. Clear the other initial cohort next. This proves feasibility without asserting that serial allocation minimizes the time.

For example, the partial-sum scan gives

B	3×3	5×5	$3 \times 3 \times 3$	$3 \times 5 \times 5$	$5 \times 5 \times 5$	$3 \times 3 \times 3 \times 3$
$h(B)$	2	3	5	8	12	12

Feasibility therefore needs only the profile scan. Minimum capture time uses the following two-count recurrence.

Define

$$C(S) = I_0(|S \cap B_0|) \cup I_1(|S \cap B_1|).$$

Isoperimetry and nesting imply that C is an idempotent strategy compression. Theorem 3.2 therefore gives an optimal strategy from the full board in which both current-color parts and both survivor parts are prefixes.

A state is the pair (u, v) of counts in the *current* colors. Choosing survivor counts $0 \leq i \leq u$ and $0 \leq j \leq v$ costs $u + v - i - j$ inspections and gives the exact next state

$$(u, v) \longrightarrow (g_1(j), g_0(i)). \quad (71)$$

The inspection sets are the suffixes $I_0(u) \setminus I_0(i)$ and $I_1(v) \setminus I_1(j)$. Thus every transition is physically realized, and every unrestricted successful search is matched by a path in this state graph.

For a fixed daily budget m , the optimum is the shortest-path distance from (E, M) to $(0, 0)$, with value infinity when there is no such path. There are $(E + 1)(M + 1) = O(V^2)$ states. If $u + v \leq m$, one final day captures every possibility. Otherwise, additional inspections cannot hurt, so it suffices to allocate exactly m useful inspections. For

$$\max(0, m - v) \leq p \leq \min(m, u),$$

the corresponding transition is

$$(u, v) \longrightarrow (g_1(v - m + p), g_0(u - p)). \quad (72)$$

There are at most $m + 1$ transitions per state. Constructing the graph and performing breadth-first search therefore take $O((m + 1)V^2)$ operations after constructing the profiles. Recording a shortest path recovers an explicit optimal inspection sequence. A finite optimum is at most $(E + 1)(M + 1) - 1$, since a shortest path repeats no state. The bounds are polynomial in the number of rooms, rather than in the logarithms of the side lengths.

Equivalently, all budgets can be considered together. Let $W_t(u, v)$ be the smallest daily budget sufficient in at most t days. Then

$$\begin{aligned} W_0(0, 0) &= 0, & W_0(u, v) &= \infty \quad ((u, v) \neq (0, 0)), \\ W_{t+1}(u, v) &= \min_{\substack{0 \leq i \leq u \\ 0 \leq j \leq v}} \max\{u + v - i - j, W_t(g_1(j), g_0(i))\}. \end{aligned} \quad (73)$$

This exact recurrence includes arbitrary interleaving of the two initial parity cohorts; it makes no serial-allocation assumption.

Remark 14.7. Every weight and lexicographic prefix is a checkerboard corner ideal: a distinct coordinatewise smaller vertex of the same parity has smaller weight. Hence the theorem also proves a corner-ideal normal form for the full-board game on every all-odd box. It does not assert that a global prefix can replace an arbitrary prescribed partial starting set without altering its optimum, and it does not settle boxes with even side lengths. The all-odd-box result gives exact feasibility, minimum time, and strategies through a polynomial recurrence; a uniform closed arithmetic expression for that time is a further question.

15 Paths: matching geometry and a probe-count potential

We prove the path formula stated in (1). The geometric lower proof applies to arbitrary possible-position sets. An interval description is used only for the explicit strategy attaining the bound.




Figure 5: An optimal four-day search of eight rooms with two daily probes. The rows show every position still possible before that day's inspection.

15.1 Neighborhood inequalities

Number the vertices $0, \dots, n-1$. Each set below is contained in one color class. If n is even, the edges $(0, 1), (2, 3), \dots$ give a perfect matching and therefore $|N(R)| \geq |R|$. We need the following information about equality.

Lemma 15.1. *On an even path, if $0 < |R| < n/2$ and $|N(R)| = |R|$, then $N(R)$ contains neither endpoint. Every nonempty set containing neither endpoint has strictly more neighbors than vertices. On an odd path with larger class E and smaller class O ,*

$$|N(R)| \geq \min\{|R|, |O|\} \quad (R \subseteq E), \quad |N(R)| \geq |R| + 1 \quad (\emptyset \neq R \subseteq O).$$

Proof. Let M send each vertex of an even path to its matching partner. Equality implies $N(R) = M(R)$. If R consists of even vertices and contains $2i > 0$, its neighbor $2i - 1$ forces $2i - 2 \in R$. Thus equality makes R an initial segment of the even vertices. Similarly an odd equality set is a final segment of the odd vertices. These statements prove both assertions about endpoints: a proper initial even segment has an odd neighborhood missing the far endpoint, and a proper final odd segment has an even neighborhood missing the other endpoint. Conversely, a nonempty equality set must itself contain an endpoint.

When n is odd, for every even vertex g there is a matching covering all vertices except g : match consecutive vertices separately to its left and right. If $R \subsetneq E$, choose $g \notin R$ to inject R into $N(R)$. For $R = E$ its neighborhood is O . Finally, for nonempty $R \subseteq O$, the matching exposing vertex 0 injects R into $N(R)$, and equality is impossible: the least occupied odd vertex has a left neighbor that would force a smaller occupied odd vertex, or the unmatched vertex 0. Thus there is at least one extra neighbor. $\square$

Track separately the targets that started in the two colors. Their current colors are opposite each day, so their allocations p, q always satisfy $p + q \leq m$. Write b for a cohort's size before inspection. On an even path retain a bit z : it is one just after a nonempty proper zero-growth neighborhood, and zero otherwise. Such a set is endpointless by the lemma. A surviving subset of it therefore expands strictly. The rank $u = 2b + z$, with rank zero for the empty set, is at most n and each physical step dominates

$$F_C(u, p) = \begin{cases} 0, & \lfloor u/2 \rfloor \leq p, \\ \min\{C, u - 2p + c\}, & \lfloor u/2 \rfloor > p, \end{cases} \quad (74)$$

with $C = n, c = 1$. Indeed, a zero-growth step raises the bit from zero to one; a step following bit one gains a vertex and may discard the bit. Both changes give at least one unit after subtracting $2p$.

For odd n , assign rank $2b$ in E , rank $2b + 1$ in nonempty O , and zero to the empty set. The same comparison holds with $c = 1$ and alternating new caps $n + 1$ in E and n in O . Initially the ranks are $n + 1, n$. Replacing both caps and initial ranks by n gives a valid lower comparison, since F_C is nondecreasing in C and u .

15.2 A general potential for the counting argument

Lemma 15.2 (Affine-rank potential). *Let $0 \leq c < 2m$, $D = 2m - c$, and $u \leq C$. Define*

$$J(u) = \left\lfloor \frac{\max\{u - c - 2, 0\}}{D} \right\rfloor, \quad \Phi(u) = \left\lfloor \frac{u + cJ(u)}{2} \right\rfloor.$$

For $0 \leq p \leq m$, the constant-cap transition (74) satisfies $\Phi(u) \leq p + \Phi(F_C(u, p))$. Consequently two such ranks sharing budget m need at least $\lceil (\Phi(u_0) + \Phi(v_0))/m \rceil$ rounds to become zero.

Proof. Both J and Φ are nondecreasing. If capture occurs, then $u \leq 2p + 1 \leq 2m + 1$, whence $J(u) = 0$ and $\Phi(u) \leq p$. If $2p \leq c$ and capture does not occur, the next rank is at least u , because $u \leq C$. Otherwise put $v = u - 2p + c$. It is positive and at most $u \leq C$, so the cap does not act. Since $2p - c \leq D$, $J(u) \leq J(v) + 1$. Therefore

$$u + cJ(u) \leq v + 2p + cJ(v).$$

Divide by two and take floors. Summing the resulting inequalities over both ranks and all rounds proves the last assertion. $\square$

For paths $c = 1$. Suppose $n > m \geq 2$, put $D = 2m - 1$, and use positive remainder coordinates

$$n = qD + r + 2, \quad 1 \leq r \leq D.$$

Here $n \geq 3$ and $q \geq 0$. Direct substitution gives

$$\Phi(n) = qm + \lfloor (r + 2)/2 \rfloor.$$

The resulting lower bound is $2q + 1$ when $r \leq m - 2$, or when $r = m - 1$ and m is even; otherwise it is $2q + 2$. For even n this is the required answer. For odd n it misses just the case m even and $r = m - 1$, which we now settle.

15.3 The equality obstruction on odd paths

Here $n = q(2m - 1) + m + 1$ with m, q even. Assume a strategy succeeds in $T = 2q + 1$ days. Use the exact alternating-cap comparison counters. For each counter, select its first zero and its last full state before that zero. If this interval starts at time a , has length L , and uses P probes, all intermediate transitions are unsaturated. Telescoping the affine transitions, including the final capture inequality, gives

$$C(a) + L \leq 2P + 2, \quad P \leq Lm.$$

Since $C(a) \geq n$, these imply $L \geq q + 1$ and $P \geq A := qm + m/2$. The two cohorts together have only $Tm = 2A$ probes, so equality holds for both. In particular $L = q + 1$ and $C(a) = n$: both intervals start when their cohort is in the smaller color class.

Their starting times have opposite parity and are at most q . The odd one, a , satisfies $1 \leq a \leq q - 1$ because q is even. If the even starting time is zero, both intervals lie in the first $2q$ days; otherwise both lie in the last $2q$ days. Their total probe count $2A = (2q + 1)m$ exceeds the $2qm$ available in either window. This contradiction proves that one extra day is necessary. The argument permits overlapping active intervals and arbitrary wasted inspections outside them.

For $m = 1$, the same potential has $D = 1$ and $\Phi(n) = n - 2$ for $n \geq 3$, giving $2n - 4$ days. On P_2 one probe cannot inspect both initial possibilities, and two consecutive inspections of one endpoint suffice. A single room is inspected directly.

15.4 An explicit strategy attaining the bound

For this construction number rooms $1, \dots, n$. Put

$$A(r) = \{r, r - 2, r - 4, \dots\} \cap \{1, \dots, n\}.$$

For $1 \leq r < n$, $N(A(r)) = A(r+1)$. Inspecting the b largest vertices of $A(r)$ leaves $A(r-2b)$; if nonempty, movement changes the frontier to $r-2b+1$. Reflection in the middle of the path gives the same rule in the opposite direction.

Assume $n \geq 3$ and $n > m$. Set $d = 2m - 1$ and

$$e = \left\lceil \frac{n-2}{d} \right\rceil, \quad \rho = n - 2 - (e-1)d, \quad c = \lfloor \rho/2 \rfloor + 1, \quad \ell = m - c.$$

Thus $1 \leq \rho \leq d$. Start with the cohort occupying $A(n-1)$ and inspect its largest m possible rooms each day. Its frontier before day t is $n-1-(t-1)d$, so it finishes on day e , using only c inspections that day. If $\ell > 0$, assign the remaining ℓ inspections to the other cohort immediately; if $\ell = 0$, start that cohort on day $e+1$.

Before its first inspection the second cohort is a whole color class. For odd n , use the same orientation. Its frontier is n on an odd starting day and $n-1$ on an even starting day. For even n , choose the orientation whose frontier is $n-1$: reflect when the starting day is odd. If its first batch has size b , the surviving frontier is $R_0 - 2b + 1$. Continue with batches of m until capture. A nonempty frontier $A(R)$ with $R \geq 2$ requires $\lceil (R-1)/d \rceil$ such days; a singleton is inspected once.

For completeness the arithmetic determining the joining day is as follows. If $\ell = 0$ the total is $2e$. If $\ell > 0$, the totals are

$$2e - 2 + \left\lceil \frac{2\rho + 2}{d} \right\rceil \quad (n \text{ odd}), \quad 2e - 2 + \left\lceil \frac{\rho + 2c}{d} \right\rceil \quad (n \text{ even}).$$

They equal $2e - 1$ exactly when $\rho \leq m - 2$, or when $\rho = m - 1$ and n, m are both even; otherwise they equal $2e$. Comparing with $\lceil (2n-4)/d \rceil$ proves (1). For $m = 1$ it gives $2n - 4$. If $n \leq m$, inspect every vertex on the first day. This completes both the construction and the arbitrary-strategy lower bound.

16 Two rows: an exact pair of counters

A two-row rectangle $P_2 \square P_n$ is also called a *ladder*: the two horizontal paths are joined by one vertical edge in each column. Its graph terminology does not impose any restriction on the searcher's movements between inspections.

For a cohort of fixed initial color, at any one time exactly one room per column has the right color. Identify its possible rooms with their column set $B \subseteq \{1, \dots, n\}$. A vertical move keeps the column, and a horizontal move changes it by one. Consequently the next column set after inspections is the *closed* path neighborhood

$$N_{P_n}[B \setminus S] = (B \setminus S) \cup N_{P_n}(B \setminus S).$$

Every nonempty proper column set gains at least one column. An initial or final interval attains the bound and stays an interval. Therefore the exact one-cohort size transition with p probes is

$$F(b, p) = \begin{cases} 0, & p \geq b, \\ \min\{n, b - p + 1\}, & p < b. \end{cases}$$

The two counters start at (n, n) and share the budget m . This is an exact lower and upper model: the inequality holds for arbitrary supports, and two boundary intervals attain every chosen sequence of allocations.

Theorem 16.1. *For $n \geq 2$ the minimum-time function is (2), with value ∞ for $m \leq 1$ and value one for $m \geq 2n$.*

Proof. For $2 \leq m < 2n$, put $N = n - 1$ and $M = m - 1$. Assign a counter of size b the weight $w(b) = \max\{b - 1, 0\}$. With a positive allocation p , its weight drops by at most $p - 1$; with no allocation its weight cannot drop. Thus the combined daily drop is at most M , and $T \geq \lceil 2N/M \rceil$.

Suppose $2N = TM$. Equality in this bound requires every day to allocate all m probes to exactly one cohort and to reduce its weight by exactly M . Each initial weight N must consequently be divisible by M . If $2N$ is divisible by M but N is not, one extra day is necessary. This is exactly the exceptional condition that M is even and $N \equiv M/2 \pmod{M}$.

To attain the answer, write $N = qM + s$, $0 \leq s < M$. If $s = 0$, clear the first interval in q full days and the second in q more. Suppose $s > 0$. After q full days on the first interval its size is $s + 1$. On the next day finish it with $s + 1$ probes and assign the remaining $M - s$ to the other, still full, interval. Its next size is

$$b = n - (M - s) + 1 = (q - 1)M + 2s + 2.$$

It then needs $\lceil (b - 1)/M \rceil$ days. The total is

$$2q + \left\lceil \frac{2s + 1}{M} \right\rceil,$$

which is the claimed ceiling plus precisely the equality correction. The second interval always survives the shared day. If $q = 0$, then $s = n - 1$ and its allocation is $M - s = m - n < n$, because $m < 2n$. If $q \geq 1$, then $M \leq n - 1$ and $M - s < n$ as well. Thus $b \geq 2$ in both cases, and the displayed solo-day formula counts at least one remaining day. The algebra also holds for $q = 0$, by the integer translation rule for ceilings.

For one probe, an uninspected cohort remains full. Inspecting one room of a full color class also leaves the other color class full after movement, since every room on a ladder with $n \geq 2$ has at least two neighbors. Hence capture is impossible. The one-day assertion is immediate from the number of initial rooms. For $n = 1$ the graph is P_2 and the path result applies. $\square$

17 Three rows: a historical rank and the exact time

We prove the three-row classification in (3), including optimal inspection schedules. Put $G = P_3 \square P_n$. The length $n = 1$ is a three-vertex path. Henceforth $n \geq 2$; a four-cycle contained in G precludes capture with one inspection per day. Indeed, a target may stay on that cycle, where a nonempty parity cohort always has two possible vertices after movement. The constructions below show that two inspections suffice.

17.1 Even lengths: the efficient survivor sets

Let $n = 2s$ and $h = 3s$, so both parity classes have h vertices. In coordinates (column, row), label one parity by

$$a_i = (2i, 0), \quad b_i = (2i, 2), \quad c_i = (2i + 1, 1) \quad (0 \leq i < s).$$

Match these vertices respectively to $(2i + 1, 0)$, $(2i + 1, 2)$, $(2i, 1)$ in the other parity. Contracting an edge followed by the inverse matching map gives a directed graph D on the labelled parity. Besides

loops, its arcs are

$$a_i \leftrightarrow c_i \leftrightarrow b_i, \quad a_i \rightarrow a_{i-1}, \quad b_i \rightarrow b_{i-1}, \quad c_i \rightarrow c_{i+1},$$

with out-of-range subscripts omitted. Consequently, for every survivor set R in this parity,

$$|N_G(R)| - |R| = |N_D^+(R) \setminus R|. \quad (75)$$

Each triple is strongly connected; its middle chain moves right and its outer chains move left. Thus D is strongly connected, and every nonempty proper parity set has surplus at least one.

Lemma 17.1. *If R is nonempty, $|R| \leq h - 2$, and $|N(R)| = |R| + 1$, then*

$$R = \{a_i, b_i, c_i : 0 \leq i < j\} \cup U, \quad U \subseteq \{a_j, b_j\}, \quad 0 \leq j < s.$$

For the opposite parity the analogous sets start at the other end. Moreover, if $R' \subseteq N(R)$ is nonempty, also has surplus one, and $|R'| \leq h - 2$, then $|R| = h - 2$ and $|R'| = 1$.

Proof. By (75), the outside boundary is a single vertex v , and R is closed under outgoing arcs of $D - v$. Deleting an outer vertex leaves a strongly connected graph: the other outer chain still moves left, the middle chain moves right, and all remaining outer vertices connect to their middle vertices. A nonempty closed set would then have all $h - 1$ remaining vertices, contrary to the hypothesis.

Deleting c_j leaves the prefix triples, the two singletons a_j, b_j , and the suffix triples as strongly connected components, omitting empty components. Arcs go from the suffix to each singleton and from each singleton to the prefix. A nonempty outgoing-closed set of size at most $h - 2$ is therefore precisely the stated prefix with some of the two singletons. Reflection in the long coordinate exchanges the physical parities and gives the opposite description.

Every nonempty classified set contains an outer corner at its starting end, and a set of size at least two contains both outer corners there. The neighborhood of R contains no outer corner at the opposite end unless $j = s - 1$. In that case the size bound permits at most one of a_j, b_j . Containing an opposite corner therefore forces $|R| = h - 2$ and supplies exactly one such corner. Hence R' must be a singleton. $\square$

17.2 A rank valid for arbitrary supports

We first assume $m < h - 2$. Track the two initial-parity cohorts separately. For a physical support A , set $z = 1$ if the preceding survivor set was nonempty, had surplus one, and had size at most $h - 2$; otherwise set $z = 0$. In particular full and empty supports have label zero. Define $\rho(A) = 2|A| + z$ for nonempty A , and $\rho(\emptyset) = 0$. This label records history and imposes no shape restriction on A .

For a numerical rank $0 \leq u \leq 2h$ and an allocation p , put

$$F_h(u, p) = \begin{cases} 0, & \lfloor u/2 \rfloor \leq p, \\ \min\{2h, u - 2p + 3\}, & \lfloor u/2 \rfloor > p. \end{cases} \quad (76)$$

It is nondecreasing in u and nonincreasing in p .

Lemma 17.2. *If $p \leq m < h - 2$ useful probes leave $R \subseteq A$, then $\rho(N(R)) \geq F_h(\rho(A), p)$.*

Proof. Write $r = |R| = |A| - p$. Empty survivors are immediate. If $r = h$ or $r = h - 1$, the neighborhood is full: every opposite-parity vertex has degree at least two. Thus its rank is the cap $2h$.

Suppose $1 \leq r \leq h - 2$. If the surplus is one, the next label is one. The old label cannot be one: Lemma 17.1 would require a support of size $h - 1$ to be reduced to one vertex, using $h - 2$ probes. Hence the new rank is $2(r + 1) + 1 = \rho(A) - 2p + 3$. If the surplus is at least two, the new rank is at least $2r + 4 \geq \rho(A) - 2p + 3$, since the old label is at most one. Capping can only weaken these inequalities. $\square$

Starting from $(2h, 2h)$, compare any actual schedule to the two-counter process

$$(u, v) \mapsto (F_h(u, p), F_h(v, m - p)), \quad 0 \leq p \leq m.$$

Its terminal inspection condition is $\lfloor u/2 \rfloor + \lfloor v/2 \rfloor \leq m$. Use the actual two useful allocations; assign unused budget arbitrarily. Monotonicity and Lemma 17.2 keep the abstract ranks below the actual ranks throughout, so the numerical process captures no later than the actual schedule.

Conversely this process is physically realizable. Order each parity by increasing coordinate sum, breaking ties by increasing column. The two orders are

$$\begin{aligned} E &:: (0, 0), \quad ((2i - 2, 2), (2i - 1, 1), (2i, 0))_{i=1}^{s-1}, \quad (2s - 2, 2), (2s - 1, 1), \\ O &:: (0, 1), (1, 0), \quad ((2i - 1, 2), (2i, 1), (2i + 1, 0))_{i=1}^{s-1}, \quad (2s - 1, 2). \end{aligned}$$

Empty repeated blocks are omitted. Reading the neighbors of successive vertices in these lists gives, for every $k > 0$,

$$N(E_k) = O_{\min(h, k+1)}, \quad N(O_k) = E_{\min(h, k+2)}.$$

Here E_k, O_k denote prefixes of size k ; empty prefixes have empty neighborhood. Deleting a suffix of p vertices and moving therefore realizes (76), with phases $z = 0, 1$ respectively. A full cohort may choose either reflected orientation freely; reflection in the long coordinate exchanges physical parities. The two cohorts choose these orientations independently. Proper cohorts retain their orientation until they fill a parity class again or disappear. Thus the numerical process gives the exact time for $m < h - 2$.

17.3 The even case: one critical budget and a specialization

At $m = 2$ and $n \geq 4$, put $V(u) = (u - 4)_+$. Directly from (76), allocating zero or one probe cannot lower V , and allocating two lowers it by at most one, including capture. The shared two-probe budget therefore lowers the sum of the two potentials by at most one per day. Its initial value is $2(2h - 4) = 6n - 8$, proving $T_2 \geq 6n - 8$. A full two-probe solo sweep decreases its rank by one until the final rank is at most five, then captures; its duration is $2h - 4$. Two independently oriented sweeps attain $6n - 8$. For $n = 2$, the same value is four by the two-row theorem (2).

For every $m \geq 3$ and even $n \geq 4$, specialize Theorem 13.1 to $b = 1$. Its numerator is $3n - 4$, its denominator is $D = 2m - 3$, and its corner cost is $J(r) = \lfloor r/2 \rfloor + 2$ for $r > 0$. Since $3n - 4 = qD + r$ is even, $r \equiv q \pmod{2}$, so

$$2J(r) = r + 4 - (q \bmod 2).$$

This is exactly the even-length condition in (3). The same theorem includes all boundary budgets and the one- and two-day cases; $n = 2$ is already covered by the two-row theorem.

The preceding numerical model also retains a useful interval bound. For $m < h - 2$, put $W = 2h - 4$ and $D = 2m - 3$. In a cohort's interval from its last full state before first capture to that capture, let ℓ be the number of days and P its total inspections. No intermediate step

saturates; the $\ell - 1$ nonterminal moves each add three rank units. The final capture inequality and $P \leq m\ell$ give

$$P \geq \left\lceil \frac{W + 3\ell}{2} \right\rceil, \quad \ell D \geq W. \quad (77)$$

These two cohort intervals may overlap arbitrarily, and their inspection totals still sum to at most the whole schedule's budget.

17.4 Odd lengths as a specialization of the general theorem

For odd $n \geq 3$, Theorem 8.1 applies with $b = 1$ at every feasible budget $m \geq 2$. For a positive argument, its corner functions are $L_0(k) = k + 1$ and $L_1(k) = k + 2$. Consequently

$$J(r) = \lfloor r/2 \rfloor + 2, \quad \delta(r) = r \bmod 2, \quad 2J(r) + \delta(r) = r + 4 \quad (r > 0).$$

Since $N - w - 1 = 3n - 4$, its quotient, remainder, and shared-day condition are exactly the odd-length branch of (3). The same theorem gives the one- and two-day thresholds. Its lower bound applies to arbitrary strategies, and its two sweeps give the attaining schedules. Thus no separate odd-length rank or equality argument is required. Together with the even-length proof and the already treated degenerate boards, this completes every three-row parameter case.

18 Four rows as a uniform-theorem corollary

For $G = P_4 \square P_n$, $n \geq 4$, the improved root threshold makes every feasible budget a case of Theorem 11.3. No separate productive-inspection argument is needed.

Corollary 18.1 (Every four-row budget). *Formula (4) holds, with physical attaining strategies and lower bounds against arbitrary searches. In particular*

$$T_3(G) = 4n - 4, \quad T_4(G) = 2n - 2.$$

Proof. Put $b = 2$, $h = 2n$. The uniform budget threshold is $\max\{3, 2(2 - 1) + 1\} = 3$. Its division is $h - 2 = q(m - 2) + r$. Since $J(1) = 2$ and $J(r) = r + 2$ for $r \geq 2$, the central-day condition is automatic for $r = 1$ and otherwise is $m \geq 2r + 4$. This gives (4), including its extreme budgets. For $m = 3, 4$ the remainder is zero, giving the two displayed specializations. The rectangular feasibility threshold excludes $m \leq 2$. Smaller lengths reduce by exchanging coordinates to the path, two-row, or three-row cases. $\square$

19 Five rows

Five rows admit a complete minimum-time classification at every budget. The two parities of the long side require different lower arguments.

Theorem 19.1 (Five rows, three inspections). *For every $n \geq 5$,*

$$T_3(P_5 \square P_n) = 10n - 20. \quad (78)$$

The odd-length proof uses the profiles of Theorem 5.1. The even-length proof uses an exhaustive finite symbolic certificate for arbitrary supports. Both arithmetic potential inequalities are verified in Lean; the complete five-row geometric argument is not presently formalized.

19.1 Odd lengths: exact profiles and a parity potential

Suppose $n \geq 5$ is odd, and put $h = (5n - 1)/2 \geq 12$. The even physical parity has $h + 1$ vertices and the odd physical parity has h . Denote these parities by $z = 0, 1$, respectively. Specializing Theorem 5.1 to short-side parameter $b = 2$ gives $g_0(0) = g_1(0) = 0$ and

$$g_0(k) = \begin{cases} h, & h \leq k \leq h + 1, \\ k + 1, & k = 1 \text{ or } h - 3 \leq k \leq h - 1, \\ k + 2, & 2 \leq k \leq h - 4, \end{cases} \quad (79)$$

$$g_1(k) = \begin{cases} h + 1, & k = h, \\ k + 2, & 1 \leq k \leq 2 \text{ or } h - 2 \leq k \leq h - 1, \\ k + 3, & 3 \leq k \leq h - 3. \end{cases} \quad (80)$$

For example, the complementary corner term in the majority profile is at most one exactly when $h + 1 - k \leq 4$. The minority corner term $Q(k)$ is at most two exactly when $k \leq 2$, and $Q(0) = 1$ handles its full set.

Define

$$W_h(b, z) = \begin{cases} 0, & b = 0, \\ 1, & 1 \leq b \leq 2, \\ b - 2, & 3 \leq b \leq 5, \\ \min\{2b + z - 8, b + h + z - 10\}, & b \geq 6. \end{cases} \quad (81)$$

A reachable nonempty support has at least three vertices in parity zero and at least two in parity one: these are the minimum degrees of vertices in the opposite parities. Thus its size satisfies

$$b = 0 \quad \text{or} \quad 3 - z \leq b \leq h + 1 - z. \quad (82)$$

Lemma 19.2 (Odd-length potential inequality). *For $h \geq 12$, $z \in \{0, 1\}$, a size satisfying (82), and $0 \leq p \leq 3$,*

$$W_h(b, z) \leq W_h(g_z((b - p)_+), 1 - z) + \mathbf{1}_{\{p \geq 2\}}. \quad (83)$$

For fixed h, z , the function $W_h(b, z)$ is nondecreasing in b .

Proof. Substitute (79)–(81), split source and target sizes at 0, 2, 5, select the smaller affine expression for sizes above five, and split the survivor size at the displayed profile endpoints. Every resulting branch is a linear integer inequality under $h \geq 12$, $p \leq 3$, and (82). This complete case split, including truncated subtraction, is checked by `potential_step` and `value_monotone` in `FiveRowOddRankArithmetic.lean`. These proofs use Lean’s ordinary kernel-checked integer-arithmetic tactic and have no admitted cases. $\square$

Take p to be the number of useful probes actually meeting one cohort. Its survivor support has size $b - p$, and its actual next support has at least $g_z(b - p)$ vertices. Monotonicity makes (83) valid for this actual transition. The cohorts occupy disjoint parities each day. Their useful allocations sum to at most three, so at most one receives two or more inspections. The sum of their potentials decreases by at most one per day. Initially

$$W_h(h + 1, 0) = W_h(h, 1) = 2h - 9;$$

at capture both vanish. Thus every strategy takes at least $4h - 18 = 10n - 20$ days.

For an attaining strategy, sweep the initially minority cohort first. Use the compatible diagonal prefixes of Theorem 5.1; each day delete its last three vertices, or all of it if fewer remain. The first two transitions are

$$(h, 1) \longrightarrow (h, 0) \longrightarrow (h - 2, 1).$$

For every minority size $6 \leq b \leq h - 2$, two rounds give

$$(b, 1) \longrightarrow (b, 0) \longrightarrow (b - 1, 1).$$

After $h - 7$ such pairs the support has size five in the minority parity. The final three transitions are

$$(5, 1) \longrightarrow (4, 0) \longrightarrow (2, 1) \longrightarrow 0.$$

The solo duration is $2 + 2(h - 7) + 3 = 2h - 9$, which is odd. The untouched cohort stays the full alternating parity, and is consequently minority when its own sweep begins. Repeat the same sweep, for a total of $4h - 18$ days. For $n = 5$, each solo sweep takes fifteen days.

19.2 Even lengths: a finite symbolic boundary classification

Suppose $n = 2q \geq 6$, and put $h = 5q$. Match horizontal pairs of columns. In row y and matched column j , the physical parity-zero vertex is $(2j + (y \bmod 2), y)$. Use the matching to identify the opposite parity with the same $5 \times q$ array. The resulting directed adjacency has a loop at every vertex, bidirectional vertical edges, leftward horizontal edges on rows 0, 2, 4, and rightward edges on rows 1, 3. The other physical parity reverses the horizontal arrows. For a support R , write $\sigma(R) = |N(R)| - |R|$. The matching identifies this surplus with its directed outside-boundary size. The directed graph is strongly connected: vertical travel reaches a row of either horizontal orientation. Thus a nonempty proper support has positive surplus.

Represent a column by a subset $B \subseteq \{0, 1, 2, 3, 4\}$, and let $V(B)$ be its vertical neighbors. For successive column masks A, B, C , define

$$c(A, B, C) = \left| \overline{B} \cap (V(B) \cup (A \cap \{1, 3\}) \cup (C \cap \{0, 2, 4\})) \right|. \quad (84)$$

The complement is within the five rows. With zero sentinel columns, the surplus of a word $B_1 \cdots B_q$ is exactly

$$\sum_{j=1}^q c(B_{j-1}, B_j, B_{j+1}).$$

For a bound $s \leq 2$, use states (A, B, e) , $0 \leq e \leq s$. Every $(0, B, 0)$ is initial. For each of the 32 choices of C , retain

$$(A, B, e) \longrightarrow (B, C, e + c(A, B, C))$$

if the new weight is at most s . A state is terminal for boundary exactly s when $e + c(A, B, 0) = s$. Explore every reachable state. Remove only self-loops, verifying that each repeats an empty or full column at zero additional weight. The remaining graph is acyclic. Its complete enumeration gives

	boundary one	boundary two
Reduced terminal paths	59	1057
Maximum reduced word length	5	8

At bound two the reduced reachable graph has 1124 states and 3874 edges. These counts follow from the stated finite rule; they are not a cutoff on physical grid length.

Every accepted word reduces to one of these terminal paths by deleting loop repetitions. Conversely, arbitrary nonnegative repetitions at the recorded positions preserve its boundary. A reduced word with k present and a absent cells therefore represents counts $k + 5u$ and $a + 5v$, where $u, v \geq 0$ count allowed full and empty repetitions. A parameter is zero if its loop type is absent. Repetition preserves whether each row is a prefix.

Lemma 19.3 (Five-row small-boundary geometry). *For $q \geq 3$, in the directed orientation just specified:*

1. *A surplus-one support has size 1, $h - 2$, or $h - 1$.*
2. *A surplus-two support of size $3 \leq k \leq h - 5$ has left-prefix rows.*
3. *A surplus-two support with prefix rows and an empty row has at most five vertices.*
4. *A surplus-two prefix support of size four, five, or six has no neighbor in the last column.*

Finite symbolic verification. Enumerate every terminal path in the specified finite graph, recording every removed-loop position. For a reduced word let k, a be its present and absent counts, and let F, E indicate a full or empty repetition position. Discard only words shorter than three columns with no repetition positions. Every remaining word satisfies these checks:

1. At boundary one, either $k = 1$ and F is false, or $a \in \{1, 2\}$ and E is false.
2. At boundary two, F implies $k \geq 10$. If $(k \geq 3$ or $F)$ and $(a \geq 5$ or $E)$, all rows are prefixes.
3. At boundary two, prefix rows and an empty row imply that F is false and $k \leq 5$.
4. At boundary two, prefix rows and $4 \leq k \leq 6$ imply that F is false and

$$B_q \cup V(B_q) \cup (B_{q-1} \cap \{1, 3\}) = \emptyset.$$

These checks cover 59 and 1057 reduced words, respectively. They are implemented directly from (84) in `five_row_boundary_automaton.py`; the full words and repetition positions are included in the accompanying certificate. No interval assumption is made about arbitrary supports.

The first three conclusions persist for all repetition counts by $k \mapsto k + 5u$ and $a \mapsto a + 5v$. In the fourth size range, no full repetition is possible. Empty repetitions occur after the nonempty part of a prefix support, and cannot create a neighbor in the new last column. This proves the assertions for every $q \geq 3$. $\square$

Lemma 19.4 (Incompatible efficient transitions). *Suppose $\sigma(R) = 2$, $4 \leq k = |R| \leq h - 5$, and $R' \subseteq N(R)$ belongs to the opposite physical parity. If*

$$k - 1 \leq |R'| \leq k + 2, \quad |R'| \leq h - 5,$$

then $\sigma(R') \neq 2$.

Proof. If both surpluses were two, the preceding lemma would give left-prefix rows for R and right-prefix rows for R' , since the second parity reverses horizontal arrows. The neighborhood of a left-prefix support also has left-prefix rows. Its size is $k + 2 \leq h - 3$, so some row misses its last cell. The corresponding right-prefix row of R' must be empty. Part 3 gives $|R'| \leq 5$, hence $k \leq 6$. Part 4 now says $N(R)$ misses the entire last column, forcing every right-prefix row of R' to be empty, a contradiction. $\square$

19.3 The even-length historical potential and attaining sweeps

Mark a current cohort by $z = 1$ exactly when its immediately preceding survivor support had surplus two and size in $[4, h - 5]$; otherwise use $z = 0$. A marked support has $6 \leq b \leq h - 3$. Every nonempty reachable support has size at least two. Define

$$V_h(b, z) = \begin{cases} 0, & b = 0, \\ 1, & 1 \leq b \leq 2, \\ b - 2, & 3 \leq b \leq 5, \\ \min\{2b + z - 8, b + h - 10\}, & b \geq 6. \end{cases} \quad (85)$$

Here z records history, unlike the physical-parity variable in (81).

For $p \leq 3$ useful inspections put $r = b - p$, and denote the next state by (b', z') . These alternatives exhaust every physical transition:

1. $r = 0$, giving $(b', z') = (0, 0)$;
2. $r = h$, necessarily $b = h, p = 0$, giving $(h, 0)$;
3. $\sigma = 1$, giving $r \in \{1, h - 2, h - 1\}$ and $(b', z') = (r + 1, 0)$;
4. $\sigma = 2$, giving $b' = r + 2$ and $z' = \mathbf{1}_{\{4 \leq r \leq h-5\}}$; if $z = 1$ and $3 \leq r \leq h - 5$, this is forbidden by Lemma 19.4;
5. $\sigma \geq 3$, giving $b' \geq r + 3$ and $z' = 0$.

In the forbidden case, the previous survivor size is $k = b - 2$ and $r = k + 2 - p$, so $k - 1 \leq r \leq k + 2$ follows from $0 \leq p \leq 3$.

Substitution into (85) gives

$$V_h(b, z) \leq V_h(b', z') + \mathbf{1}_{\{p \geq 2\}}. \quad (86)$$

This is a finite split of linear integer inequalities under $h \geq 15$, $b \leq h$, $b = 0$ or $b \geq 2$, and $z = 1 \Rightarrow 6 \leq b \leq h - 3$. Its universal proof is `FiveRowRankArithmetic.potential_step` in the Lean artifact; the geometric alternatives remain the ordinary argument just given. The two-cohort potential decreases by at most one per day. Initially it is $2V_h(h, 0) = 4h - 20 = 10n - 20$, proving the lower bound.

For the upper bound use ascending $(x + y, x)$ prefixes, choosing a fixed reflected orientation independently for each initial cohort. Their surpluses are as follows; these describe the stated orders, rather than claiming that they simultaneously minimize all even-length profiles:

Phase	survivor sizes	surplus
0	$1, h - 2, h - 1$	1
0	all other proper nonempty sizes	2
1	$h - 1$	1
1	$1, 2, h - 4, h - 3, h - 2$	2
1	$3, \dots, h - 5$	3

Empty and full supports have empty and full neighborhoods. Neighborhoods of prefixes are opposite-phase prefixes. To derive the tables, let $L(d)$ be the length of diagonal $x + y = d$, and set $A(0) = 0$, $A(d + 1) = L(d) - A(d)$. For a prefix with j cells in its final diagonal, the surplus is

$$A(d) + 1 - \mathbf{1}_{\{d \geq 4\}} - \mathbf{1}_{\{\max(0, d-4) + j - 1 = n-1\}}.$$

The diagonal lengths are 1, 2, 3, 4, 5, then five through diagonal $n - 1$, then 4, 3, 2, 1. The recurrence gives $A(d) = \lceil d/2 \rceil$ for $d \leq 4$, $A(d) = 2 + (d \bmod 2)$ for $4 \leq d \leq n - 1$, and final values 2, 2, 1, 1. Substitution gives the tables and the neighborhood-prefix assertion.

Starting in phase zero, three rounds give $h \rightarrow h - 1 \rightarrow h - 2 \rightarrow h - 3$. Then $h - 8$ pairs of a stationary round and a one-cell decrease reach size five. The last three rounds are $5 \rightarrow 4 \rightarrow 2 \rightarrow 0$. The solo duration is

$$3 + 2(h - 8) + 3 = 2h - 10.$$

Reflection $x \mapsto n - 1 - x$ swaps physical parities because n is even, so each initial cohort can start its sweep in favorable phase zero. The untouched cohort remains full until its sweep begins. Two consecutive sweeps take $4h - 20$ days, completing Theorem 19.1.

19.4 Even lengths at every budget

The one-bit obstruction can be strengthened enough to resolve every budget on the even-length boards. We treat these first, then solve the odd-length recurrence explicitly.

Theorem 19.5 (Five rows, even length, every budget). *Let $n \geq 6$ be even and $h = 5n/2$. Budgets one and two are impossible; $T_3 = 4h - 20$; and*

$$T_4 = 2 \left\lceil \frac{2h - 8}{3} \right\rceil. \quad (87)$$

For $5 \leq m < h$, write $2h - 6 = q(2m - 5) + r$, $0 \leq r < 2m - 5$, and set

$$J(0) = 0, \quad J(1) = 2, \quad J(2) = 3, \quad J(3) = J(4) = 4, \quad J(r) = \left\lceil \frac{r + 5}{2} \right\rceil \quad (r \geq 5).$$

Then

$$T_m = \begin{cases} 2q, & r = 0, \\ 2q + 1, & r > 0 \text{ and } 2J(r) \leq m, \\ 2q + 2, & \text{otherwise.} \end{cases} \quad (88)$$

Finally, $T_m = 2$ for $h \leq m < 2h$, and $T_m = 1$ for $m \geq 2h$.

For $m \geq 7$, this is already Theorem 13.1 with $b = 2$: its numerator is $2h - 6$, and its corner function is exactly the displayed J . We retain the stronger boundary relation below, then use it to settle the remaining budgets four, five, and six.

We first strengthen the geometry in Lemma 19.3. Call the two endpoints in the last matched column *outward corners* when considering the reversed directed orientation.

Lemma 19.6 (Endpoint compatibility). *Suppose $\sigma(R) = 2$ and $3 \leq k = |R| \leq h - 5$. If $N(R)$ contains an outward corner, then $k \geq h - 6$. If $R' \subseteq N(R)$ has opposite physical parity, surplus two, and $3 \leq \ell = |R'| \leq h - 5$, then*

$$(k, \ell) \in \{(h - 6, 3), (h - 5, 3), (h - 5, 4)\}. \quad (89)$$

Proof. Write $n = 2s$. For $k \geq 6$, the prefix and empty-row conclusions of Lemma 19.3 imply that every row of R is a nonempty left prefix. Let its lengths be t_i and set $d_i = s - t_i$. The missing lengths in the rows of $N(R)$ are exactly

$$\begin{aligned} e_0 &= \min(d_0, d_1), & e_1 &= \min(d_0, (d_1 - 1)_+, d_2), & e_2 &= \min(d_1, d_2, d_3), \\ e_3 &= \min(d_2, (d_3 - 1)_+, d_4), & e_4 &= \min(d_3, d_4). \end{aligned}$$

The surplus identity is $\sum d_i = \sum e_i + 2$. Splitting the minima gives these three elementary integer consequences:

$$\begin{aligned} e_0 e_4 = 0 &\implies \sum d_i \leq 6, \\ (e_0 = e_1 = e_2 = 0) \text{ or } (e_2 = e_3 = e_4 = 0) &\implies \sum d_i \leq 5, \\ e_0 = e_2 = e_4 = 0 &\implies \sum d_i \leq 2. \end{aligned}$$

Their complete universal arithmetic proofs are the theorems `endpoint`, `three_rows`, and `alternating_rows` in `FiveRowDeficitArithmetic.lean`. The first proves the corner claim.

The rows of R' are right prefixes. Every occupied row must therefore be full in $N(R)$, that is, have $e_i = 0$. Since $|N(R)| \leq h - 3$, some row is not full, so the empty-row conclusion implies $\ell \leq 5$. The surplus-two prefix row vectors of sizes three through five, up to vertical reflection, are

$$\begin{array}{l|l} 3 & (2, 1, 0, 0, 0), (1, 1, 1, 0, 0), (1, 0, 1, 0, 1) \\ 4 & (2, 1, 1, 0, 0), (1, 1, 1, 0, 1) \\ 5 & (2, 1, 1, 0, 1), (1, 1, 1, 1, 1). \end{array}$$

This finite list follows by substituting the row-neighborhood maxima for vectors with total at most five. For $s \geq 6$ the far boundary cannot affect the check; direct substitution for $s = 3, 4, 5$ gives the same list. A size-three vector forces a zero endpoint deficit; a size-four vector forces three consecutive zeros at an end or the stronger alternating zeros; a size-five vector forces alternating zeros. The three inequalities above now give respectively $k \geq h - 6$, $k \geq h - 5$, and the contradiction $k \geq h - 2$.

For the source sizes $3 \leq k \leq 5$, use the same finite list directly. Every neighborhood row has length at most two, hence misses the last column because $s \geq 3$. It contains neither an outward corner nor a nonempty right prefix. This separate check is needed because the displayed deficit formula assumes that every source row is nonempty. $\square$

For this all-budget argument, mark a cohort by $z = 1$ when its previous survivor had surplus two and size in $[3, h - 5]$. Thus $z = 1$ implies $5 \leq b \leq h - 3$, and the previous survivor size was $b - 2$. Every physical transition, for p useful inspections and $a = b - p$, belongs to the following relation:

1. $a = 0$ gives $(0, 0)$, and $a = h$ gives $(h, 0)$;
2. surplus at least three gives any $(b', 0)$ with $a + 3 \leq b' \leq h$;
3. surplus one gives $(a + 1, 0)$ for $a \in \{1, h - 2, h - 1\}$, with $a = 1$ forbidden when $z = 1$ and $b < h - 4$;
4. surplus two gives $(a + 2, \mathbf{1}_{\{3 \leq a \leq h - 5\}})$; when $z = 1$ and $3 \leq a \leq h - 5$, the pair $(b - 2, a)$ must satisfy (89).

The last two restrictions are precisely Lemma 19.6. In particular, the large-surplus branch includes every possible target size, so this is a lower relaxation for arbitrary supports.

19.4.1 A boundary potential and the budgets five and six

For $m \geq 5$, let F_m be the general potential (36) with $b = 2$. An explicit residue form is

$$\begin{aligned} D &= 2m - 5, & j(u) &= \left\lfloor \frac{(u-7)_+}{D} \right\rfloor, & s(u) &= ((u-7)_+ \bmod D) + 1, \\ F_m(u) &= \left\lfloor \frac{u + 5j(u)}{2} \right\rfloor - \mathbf{1}_{\{j(u) \geq 1, s(u) \in \{1,2,4\}\}}. \end{aligned} \quad (90)$$

The common interior lemma gives the following properties:

1. F_m is nondecreasing and $F_m(u+2) - F_m(u) \leq 3$;
2. if $\lfloor u/2 \rfloor \leq m$, then $F_m(u) = \lfloor u/2 \rfloor$;
3. if $p \leq m$ and $\lfloor u/2 \rfloor > p$, then $F_m(u) \leq p + F_m(u - 2p + 5)$.

Indeed, Lemma 8.3 applies because $m \geq b^2 + 1 = 5$. Writing $u = 2k + z$, its target rank is at most $u - 2p + 5$, since $L_0(s) \leq s + 2$ and $L_1(s) \leq s + 3$. Monotonicity gives property 3, while property 2 is the defining cardinality branch. The only additional assertion, the upper bound on a two-unit increment, follows from the displayed J values and $J(D) = m$, including the first period and the wrap between periods.

For $5 \leq m \leq h - 8$, set

$$C = \min\{F_m(2h), 2 + F_m(2h - 2), 4 + F_m(2h - 4)\}, \quad \Phi(b, z) = \min\{C, F_m(2b + z)\}.$$

Every permitted transition lowers Φ by at most its useful allocation p . Here are all exceptional cases in that verification. A bulk surplus-two transition from $z = 0$ has target rank $u - 2p + 5$; from $z = 1$ it would require $p \geq h - 7 > m$ by (89). Surplus at least three has target rank at least $u - 2p + 5$. The three properties of F_m therefore handle these branches; a common cap preserves the inequality.

Capture follows from property 2. A surplus-one singleton is impossible from a marked state because it would require $p \geq h - 5$. From an unmarked state, $b \leq m + 1$; the only extra endpoint value is $F_m(2m + 2) = m + 2$, paid by m inspections and the target pair. For surplus two at survivor sizes one or two, states $b \leq m$ follow from cardinality. The four extra source values at $(b, z) = (m + 1, 0), (m + 1, 1), (m + 2, 0), (m + 2, 1)$ are $m + 2, m + 3, m + 4, m + 4$, respectively, and the target size is at most four.

Finally consider survivor sizes $h - 4, \dots, h - 1$. Rank increases need only monotonicity. The sole proper-source rank decrease is unmarked $h - 1 \rightarrow h - 2$ with three inspections, paid by property 1. The improvements from the full state to $h - 1$ with two inspections and to $h - 2$ with four inspections are exactly the last two terms defining C . The full, zero-inspection transition is tautological. This exhausts the relation.

Both initial cohorts have potential C and both final potentials vanish, so this boundary relation gives $T_m \geq \lceil 2C/m \rceil$ throughout $5 \leq m \leq h - 8$. Direct substitution in the quotient and remainder of (88) gives, for $m \geq 6$,

$$C = qm + J(r) - \mathbf{1}_{\{r=2\}}.$$

Only $m = 5, 6$ remain after the general theorem, and both satisfy $m \leq h - 8$ because $h \geq 15$. For $m = 5$, the physical value h is divisible by five: $q = n - 2$, $r = 4$, and $C = 5q + 4$, giving $T_5 \geq 2q + 2 = 2n - 2$. For $m = 6$, the values of $C - 6q$ for $r = 0, \dots, 6$ are

$$0, 2, 2, 4, 4, 5, 6.$$

The resulting lower bounds $\lceil 2C/6 \rceil$ are $2q$, $2q + 1$ for $r = 1, 2$, and $2q + 2$ for $r \geq 3$, as required.

19.4.2 The four-inspection lower bound

Define

$$H(b, z) = \begin{cases} b, & b \leq 4, \\ 6, & (b, z) = (5, 0), \\ 4 \lfloor (2b + z)/3 \rfloor - 8 + 3\mathbf{1}_{\{(2b+z) \bmod 3=2\}}, & \text{otherwise,} \end{cases}$$

and

$$A_1 = \min\{H(h-1, 0), 3 + H(h-2, 0)\}, \quad C_4 = \min\{H(h, 0), 2 + A_1, 4 + H(h-2, 0)\}.$$

Use potential C_4 at the full state, A_1 at the unmarked state $h-1$, and H elsewhere. Allocating $p \leq 4$ lowers it by at most p . This assertion is a finite residue split with the unbounded parameter $h \geq 15$: the theorem `potential_step` in `FiveRowFourProbeArithmetic.lean` checks every branch of the physical relaxation above. In this budget range a marked source forbids all surplus-one transitions and all bulk surplus-two transitions; the endpoint restrictions prove these exclusions. The checker includes all larger targets in the surplus-at-least-three branch.

Writing $h = 3t, 3t + 1, 3t + 2$ gives respectively

$$(A_1, C_4) = (8t - 12, 8t - 10), \quad (8t - 9, 8t - 8), \quad (8t - 5, 8t - 4).$$

For $L = \lceil (2h - 8)/3 \rceil$, this means $C_4 = 4L$ except when $h \equiv 0 \pmod{3}$, where $C_4 = 4L - 2$. The total-ammunition inequality already proves $T_4 \geq 2L$ outside that residue.

In the exceptional residue, $2L - 1$ days would use exactly $2C_4$ probes, so every cohort transition must be tight and every day must use four useful probes. From a full state, the only positive tight allocation is two inspections, producing size $h - 1$. Thus the first day must split $2 + 2$. From an unmarked state of size $h - 1$, every tight allocation requires four inspections. Two such transitions cannot share day two. This proves $T_4 \geq 2L$. The two local tightness assertions are universally checked by `full_tight` and `penultimate_tight` in `FiveRowFourProbeEquality.lean`.

19.4.3 Explicit attaining schedules for the remaining budgets

All budgets $m \geq 7$, including the boundary cases, are covered by Theorem 13.1. The independent finite lower certificates at $(h, m) = (15, 8), (15, 11), (20, 14)$ remain kernel checked in `FiveRowHighBudgetArithmetic.lean`.

For the matching upper bounds, use the previously displayed physical prefix tables. With four inspections, start each solo sweep in phase one. Its first round gives $(h - 2, 0)$. At phase-zero size $b \geq 9$, two rounds give $(b - 2, 1)$ and $(b - 3, 0)$. Repeat until size 7, 8, 6, for h modulo three equal to 0, 1, 2, respectively. The final tails are $7 \rightarrow 5 \rightarrow 3 \rightarrow 0$, $8 \rightarrow 6 \rightarrow 4 \rightarrow 0$, and $6 \rightarrow 4 \rightarrow 0$. The solo duration is L , so two independently oriented sweeps attain $2L$.

For $m = 5, 6$, full allocations decrease the bulk rank by $D = 2m - 5$. If $r = 0$, each solo sweep lasts q rounds. If $r > 0$, after q full allocations its remaining size is $J(r)$; substitution in the last two prefix transitions gives the special residues one, two, and four. At $m = 5$, only $r = 4$ occurs, so two separate sweeps suffice. At $m = 6$, a shared day is needed only for $r = 1, 2$, where $J(r) = 2, 3$. The corresponding initial allocations to a suitably oriented fresh cohort decrease its rank by one and two, respectively. Its new rank is therefore at most $qD + 6$, which clears within q further

full allocations by the same prefix table. Finish the first cohort and start the second with $J(r)$ inspections each on that shared day. The remaining residues use two separate sweeps, attaining (88).

The impossibility at budgets below three follows from the 4×4 subgrid obstruction. At budget at least h , inspect one full parity then the other cohort's full parity; one day is possible exactly at budget $2h$. Together with Theorem 19.1, these facts complete the proof of Theorem 19.5. The geometry here is an ordinary proof with a finite symbolic certificate; the cited Lean modules verify its universal scalar consequences, not the complete physical five-row classification.

19.5 Odd lengths at every budget

Theorem 19.7 (Five rows, odd length, every budget). *Let $n \geq 5$ be odd and $h = (5n - 1)/2$. Budgets at most two are impossible, $T_3 = 10n - 20$, and*

$$T_4 = 2 \left\lceil \frac{5n - 9}{3} \right\rceil + 2\mathbf{1}_{\{3|n\}}, \quad T_5 = 2n - 2.$$

For $5 \leq m < h$, write $5n - 6 = q(2m - 5) + r$, $0 \leq r < 2m - 5$. Use the function J from Theorem 19.5, and put

$$\delta(r) = \mathbf{1}_{\{r \in \{1, 2, 4\} \text{ or } (r \geq 5 \text{ and } r \text{ odd})\}}.$$

Then

$$T_m = \begin{cases} 2q, & r = 0, \\ 2q + 1, & r > 0 \text{ and } 2J(r) + \delta(r) \leq m, \\ 2q + 2, & \text{otherwise.} \end{cases} \quad (91)$$

Finally, $T_m = 2$ for $h \leq m < 2h + 1$, and $T_m = 1$ for $m \geq 2h + 1$.

All transitions below use the exact profiles (79)–(80). An arbitrary physical search dominates the corresponding count search, and compatible prefixes attain it. Thus a lower proof may either use a nondecreasing potential directly on arbitrary supports, or rule out an equally fast prefix search. No serial-search assumption is needed.

19.5.1 Four inspections

Use the interior function $H(b, z)$ from the four-inspection even-length proof. For $h = 3t, 3t+1, 3t+2$, respectively, put $C(h) = 8t - 8, 8t - 5, 8t - 4$, and define

$$V_h(b, z) = \begin{cases} C(h) - 1, & z = 0, b = h, \\ \min\{H(b, 0), C(h) - 2\}, & z = 0, b = h - 1, \\ \min\{C(h), H(b, z)\}, & \text{otherwise.} \end{cases}$$

For $h \geq 12$, $b = 0$ or $3 - z \leq b \leq h + 1 - z$, and $0 \leq p \leq 4$,

$$V_h(b, z) \leq p + V_h(g_z((b - p)_+), 1 - z). \quad (92)$$

Here is a finite reduction that verifies the assertion for unbounded h . The identities $C(h + 3) = C(h) + 8$ and $H(b + 3, z) = H(b, z) + 8$ for $b \geq 6$ give the same translation for both exceptional caps. Also $g_{h+3,z}(k + 3) = g_{h,z}(k) + 3$ for $k \geq 3$ in the domain. For $h \geq 24, b \geq 13$,

reduce both h, b by three; survivors are at least six before reduction, and both sides of (92) decrease by eight. For $b \leq 12$, reduce only h by three: target sizes are at most fifteen, no high-end clause is reached, and the caps are inactive since $C(h-3) \geq 48 > H(b', z')$ for $b' \leq 15$. Only $12 \leq h \leq 23$ remain. All states and all five allocations in this finite range are checked by `RecurrenceBridgeOddFiveFour.finite_certificate` in Lean. The translations are ordinary proofs, separate from this finite check.

The initial potential sum is $2C(h)$, so summation gives $T_4 \geq \lceil C(h)/2 \rceil$. For attainment start the majority cohort. The first two rounds give $(h+1, 0) \rightarrow (h-2, 1) \rightarrow (h-3, 0)$. From majority size $b \geq 9$, two rounds give $(b-2, 1) \rightarrow (b-3, 0)$. The final majority sizes 6, 7, 8 have tails $6 \rightarrow 4 \rightarrow 0$, $7 \rightarrow 5 \rightarrow 3 \rightarrow 0$, and $8 \rightarrow 6 \rightarrow 4 \rightarrow 0$. The solo durations are $L = 2t-2, 2t-1, 2t-1$, respectively. When L is odd the second cohort starts in majority parity. When L is even, its minority solo has the same duration: the first round reaches majority size $h-1$, followed by the same pairs and the size-eight tail. Thus $2L = \lceil C(h)/2 \rceil$ days attain the formula.

19.5.2 Larger budgets as a specialization of the general theorem

For $m \geq 5$, apply Theorem 8.1 with $b = 2$. Its parameter $B = b^2$ is four, so its full budget range applies. The two uncapped corner functions simplify to

$$L_0(k) = \begin{cases} 0, & k = 0, \\ 2, & k = 1, \\ k+2, & k \geq 2, \end{cases} \quad L_1(k) = \begin{cases} 0, & k = 0, \\ k+2, & 1 \leq k \leq 2, \\ k+3, & k \geq 3. \end{cases}$$

Substituting in (33) gives precisely the function J in Theorem 19.5. Substitution in (34) gives

$$\delta(r) = \mathbf{1}_{\{r \in \{1,2,4\} \text{ or } (r \geq 5 \text{ and } r \text{ odd})\}}.$$

Finally $N - w - 1 = 5n - 6$. Thus the general theorem is exactly (91), including its one- and two-day thresholds and its explicit attaining schedules. At $m = 5$, the quotient and remainder are $q = n - 2, r = 4$, giving $T_5 = 2n - 2$. The proof of the general theorem establishes this boundary case directly, so this specialization has no circular dependency.

The independent audits and direct room-coordinate constructions remain in the research archive. The finite Lean arithmetic bases do not amount to a complete physical Lean proof of the five-row theorem.

20 Seven rows with five inspections per day

The following low-budget family illustrates how a periodic potential and a short equality obstruction can solve the two-count recurrence explicitly. The proof is an ordinary argument with finite rational certificates; it is not presently a complete physical Lean theorem.

Theorem 20.1. *For every odd $n \geq 7$,*

$$T_5(P_7 \square P_n) = 2 \left\lceil \frac{7n-16}{3} \right\rceil.$$

Put $H = (7n+1)/2$, so the two checkerboard classes have sizes $H, H-1$. Theorem 5.1, with $b = 3$, supplies exact profiles g_0, g_1 and compatible prefixes. Thus $g_p(0) = 0$, and

$$g_0(k) = k + \min\{R(k), 3, R(H-k) - 1\}, \quad g_1(k) = k + \min\{Q(k), 4, Q(H-1-k)\}$$

on the respective nonempty domains. We use the profiles for the lower bound on arbitrary supports and the compatible prefixes for construction.

Finite potentials and the exceptional residue. Assign charges

$$c_0 = (0, 0, \frac{1}{2}, \frac{3}{4}, 1, 1), \quad c_1 = (0, 0, \frac{1}{4}, \frac{1}{2}, 1, 1).$$

They satisfy $c_0(r) + c_1(s) \leq 1$ whenever $r + s \leq 5$. For each finite base below, define $f_p(k)$ as the least total charge of a path from (p, k) to an empty state in the one-cohort graph

$$(p, k) \longrightarrow (1 - p, g_p(\max\{k - r, 0\})), \quad \text{edge charge } c_p(r), \quad 0 \leq r \leq 5.$$

All quantities are multiples of $1/4$. The accompanying certificate computes these finite shortest-path distances exactly and verifies their monotonicity, zero terminal values, and every inequality

$$f_p(k) \leq c_p(r) + f_{1-p}(g_p(\max\{k - r, 0\})). \quad (93)$$

The complete base data give the following initial sums $V = f_0(H) + f_1(H - 1)$ and attaining serial times U :

n	7	9	11	13	15	17	19	21	23
V	22	63/2	41	50	119/2	69	78	175/2	97
U	22	32	42	50	60	70	78	88	98

There are 5,778 one-cohort inequalities in these nine certificates. Summing (93) across the current physical colors bounds the total potential loss by one per day. Monotonicity handles neighborhoods larger than the profile minimum. Hence $T_5 \geq \lceil V \rceil$.

For $n = 11, 17, 23$, this leaves one day. If capture occurred in $V = U - 1$ days, every day would have to lose exactly one potential unit. In these three bases the only tight first allocation is three inspections in color zero and two in color one. Its minimum successor counts are $(H - 1, H - 2)$. The second-day losses from that state, for allocations $r = 0, \dots, 5$ in color zero, are

$$\left(\frac{3}{4}, -\frac{1}{4}, \frac{1}{2}, 0, 0, \frac{1}{4}\right). \quad (94)$$

None is one, a contradiction. An actual first successor with larger counts cannot evade this calculation: if its first day is tight, its potential equals that of the minimum successor, while monotonicity makes its next potential at least as large. Filling unused quota is harmless. This proves all nine base lower bounds.

A periodic extension valid at every larger length. Use the three base triples

$$(n_0, H_0, c) = (19, 67, 33), (21, 74, 36), (23, 81, 39).$$

Their exact tables satisfy, for both colors,

$$\begin{aligned} f_p(k + 3) &= f_p(k) + 2, & c - 10 &\leq k \leq c + 10, \\ g_p(k) &= k + 3 + p, & c - 15 &\leq k \leq c + 15. \end{aligned} \quad (95)$$

For $n = n_0 + 6j$, write $\Delta = 21j$, so $H = H_0 + \Delta$. Keep f_p unchanged through c . On the inserted interval define

$$\tilde{f}_p(c + 3\ell + r) = f_p(c + r) + 2\ell, \quad 1 \leq r \leq 3, \quad 0 < 3\ell + r \leq \Delta.$$

Above $c + \Delta$, put $\tilde{f}_p(k) = f_p(k - \Delta) + 2\Delta/3$. The period identity makes this a nondecreasing extension.

We check (93) by three exhaustive source ranges. For $k \leq c - 5$, every output is at most $c - 1$ and the distant upper corner has no effect; the old inequality is unchanged. For $k \geq c + \Delta + 5$, subtract Δ from source and output. The lower corner is saturated, and both profile and potential translate exactly. In the remaining range $c - 4 \leq k \leq c + \Delta + 4$, residuals are at least $c - 9$ and outputs are at most $c + \Delta + 8$. All relevant profiles are on their plateau. Reduce the source modulo three to a representative in $\{c, c + 1, c + 2\}$. The output changes by the same multiple of three, and both potentials change by the same multiple of two. The ten-place collar in (95) contains all needed old values, so this too is a verified base inequality. The explicit profile formulas justify the corner invariance and translation at every Δ .

The initial sum increases by $4\Delta/3 = 28j$. For the residue $n \equiv 5 \pmod{6}$, the first two transitions remain in the translated upper part, so the tight first allocation and the loss table (94) are unchanged. The extra-day obstruction therefore persists. Since the stated formula also increases by $28j$ when n increases by $6j$, all lower bounds follow. The six smaller bases cover the lengths below these three starting points.

An attaining physical strategy. Use the compatible prefixes from Theorem 5.1. Allocate five inspections to one cohort until it is captured, assigning any spare quota on its last day to the other cohort, and then finish that cohort. The base times in the table are direct evaluations of this strategy. At bases $n_0 = 19, 21, 23$, successful first colors are respectively 1, 0, 0. The two active middle visits, recorded as (count, current color), are

n_0	first cohort	second cohort
19	(34, 1)	(34, 1)
21	(37, 1)	(38, 1)
23	(41, 1)	(41, 1)

All cuts exceed five. The companion is consequently globally full at the first visit and empty at the second; the finite certificates check these conditions directly.

On the common plateau, two consecutive five-inspection days reduce an active count by three and restore its color, from either starting color. Insert $2\Delta/3 = 14j$ days at each visit to reduce the translated count from $k + \Delta$ to k . Before a visit, the profile trajectory translates by Δ ; afterwards its lower part is unchanged. The active count decreases under five inspections, so the old segments stay on their respective sides. The inserted segment lies wholly in the plateau. Its even duration preserves the companion's color and the original last-day spare allocation. Both insertions are therefore compatible physical prefix strategies. Their total added time is $28j$, matching the lower bound and proving Theorem 20.1.

The finite certificates and their independent replay include the full rational base tables, all inequalities, the strict second-day losses, the collar identities, and actual coordinate-neighborhood replays of the nine base searches. The three-range proof and the physical insertion establish every larger length; the result does not rely on extrapolating the finite time table.

21 Minimum-budget searches on every odd width

At the minimum feasible budget, one width-dependent corner clock gives the exact time on every odd rectangle. The common-ancestry lemma removes the last central-day ambiguity uniformly,

including widths for which many exceptional mixed states survive near the final corner.

Fix $w = 2b + 1$, $m = b + 1$, and set $C = b(b - 1)$. Define the width-only recurrence

$$\begin{aligned} A_0 &= B_0 = 0, \\ A_{t+1} &= B_t + m - \min\{q(B_t + m), b\}, \\ B_{t+1} &= A_t + m - \min\{R(A_t + m), b + 1\}. \end{aligned} \tag{96}$$

Here $q(z)$ is the least nonnegative integer r with $z \leq r(r + 1)$. Let L_b be the first time $B_t = C$, with $L_1 = 0$.

Theorem 21.1. *For every odd $n \geq w = 2b + 1 \geq 3$,*

$$T_{b+1}(P_w \square P_n) = 2wn - C_w, \quad C_w = 8b^2 - 4b - 4L_b + 4. \tag{97}$$

The budget $b + 1$ is the minimum feasible budget, and the recurrence (96) terminates independently of n . Its clock can be evaluated in $O(b)$ arithmetic stages using Proposition 6.4. In particular,

w	3	5	7	9	11	13	15
C_w	8	20	44	84	136	208	288.

Proof. The two scalar maps in (96) are at least $z + 1$ and z , respectively; each two-step composition therefore increases its argument by at least one. The recurrence is monotone from zero. At input $C - 1$ both maps equal C , while at C they equal $C + 1$ and C . Before B first reaches C , A cannot exceed C : producing $A \geq C + 1$ requires the preceding $B \geq C$. Thus the first arrival is $(C - 1, C)$ or (C, C) , followed by $(C + 1, C)$ and $(C + 1, C + 1)$. For $b = 1$ this follows directly from $(A_0, B_0) = (0, 0)$; expressions at $C - 1$ are needed only for $b \geq 2$. These observations also prove finite termination and exact arrival at C rather than an overshoot.

This bottom history through time $L_b + 2$ is the physical solo history on every allowed rectangle. Its inverse inputs are at most $b^2 + 2$. A reflected branch can first improve ι_0 at input $M - C$, at least $b^2 + 3b$, and the corresponding ι_1 threshold is farther away. The $b = 1$ endpoints satisfy the same exclusion directly. At time $L_b + 2$ both solo deficits are $C + 1$. Total concentration makes every mixed state low-total; Lemma 7.2 therefore replaces the retained frontier by the two pure endpoints exactly.

In the common affine interval the solo ties evolve as

$$(s, s) \mapsto (s + 1, s) \mapsto (s + 1, s + 1).$$

At the first step every mixed successor has total at most s , by both inverse concentration inequalities. The second step is a solo tie and again has no exceptional states. Hence the retained frontier stays pure. Choose

$$a_* = M - b^2 - 1, \quad t_* = L_b + 2(a_* - C).$$

The inverse formulas give these affine transitions through the tie (a_*, a_*) . It lies after $C + 1$ since $a_* - (C + 1) \geq 3b - 2 > 0$. At time t_* the two active solo belief sizes are $b^2 + 2$ and $b^2 + 1$.

For a canonical monochromatic belief of current color p and size k , greedy capture within T days is equivalent to

$$k \leq m + D_p(T - 1), \tag{98}$$

where D is the fresh inverse recursion from zero. Indeed, apply the integer inverse profile successively backward through the $T - 1$ nonterminal moves, with the final inspection threshold m .

For the two displayed belief sizes the required fresh deficits are $C + 1$ in color zero and C in color one. Their exact upper-tail solo durations are consequently $L_b + 2$ and $L_b + 1$. The full-board solo durations are consecutive, the faster being

$$\tau = t_* + L_b + 1 = wn - 4b^2 + 2b + 2L_b - 2.$$

It remains to exclude $2\tau - 1$ days. At the last pure tie, every subsequent secondary ancestry starts from zero. The precentral upper tail lasts exactly L_b steps, so monotonicity bounds every secondary by the fresh solo deficits up to time L_b , hence by C . Lemma 7.3 says all final exceptional pairs have the same primary ancestry. Any two such pairs leave at least

$$M - 2C \geq 4b > b + 1 = m$$

rooms in their other color, so cannot win centrally. If either pair is nonexceptional, their combined total is at most $D_0 + D_1$ and they cannot improve on the solo endpoints. The slower solo needs $\tau + 1$ days; its remaining count at time $\tau - 1$ therefore exceeds m , or it could be captured on day τ . Thus the solo central sum also exceeds m . The exact midpoint theorem forces $T = 2\tau$, proving (97).

Finally the majority profile has maximum surplus b , so Corollary 14.6 gives minimum budget $b + 1$. The clock also satisfies $L_b = \lceil \mu_1(C)/(b + 1) \rceil$ by Proposition 7.6. Equivalently, (98) identifies $L_b + 1$ with the solo clearing time of a minority prefix of size $b^2 + 1$ on the width- w square. Proposition 6.4 evaluates this time in $O(b)$ stages. Evaluating the width-only recurrence gives the seven displayed constants. $\square$

The proof is uniform in the width. It does not assert serial optimality from arbitrary partial states. Its geometric and corner-clock arguments are ordinary proofs; the abstract inverse, persistence and ancestry implications have separate Lean verification with their hypotheses explicit.

21.1 Reusable lower certificates and affine insertion

The following certificate method has a broader purpose than the numerical corollaries above: its insertion lemmas also apply to other compatible profile families, including the three-dimensional cylinder below.

Set $w = 2b + 1$, $m = b + 1$, $D = m + 1$, and let the two color classes have sizes H and $H - 1$. Use the profiles g_0, g_1 from Theorem 5.1. Choose nonnegative rational numbers $c_p(r)$, for $p \in \{0, 1\}$ and $0 \leq r \leq m$, with

$$c_0(r) + c_1(m - r) \leq 1. \tag{99}$$

Suppose nonnegative functions F_p , with $F_p(0) = 0$, satisfy

$$F_p(k) - F_{1-p}(g_p((k - r)_+)) \leq c_p(r) \quad (0 \leq k \leq H - p, \ 0 \leq r \leq m). \tag{100}$$

Here $(x)_+ = \max(x, 0)$, as elsewhere in the search recurrences.

At a current state (u, v) the sum $F_0(u) + F_1(v)$ decreases by at most one per day, by (99). It suffices to consider strategies using the full quota whenever more than m possibilities remain: additional suffix inspections preserve the prefix normal form and cannot hurt. On the final day one can pad the two quotas to sum to m , with the positive-part convention making both survivors empty. Consequently

$$T_m \geq \lceil F_0(H) + F_1(H - 1) \rceil. \tag{101}$$

The functions are indexed by current color, so this argument includes arbitrary interleaving of the two initial cohorts.

For each finite certificate below, $F_p(k)$ is computed as the shortest total charge of a path from (p, k) to an empty cohort, where a quota r costs $c_p(r)$ and sends the cohort to $(1-p, g_p((k-r)_+))$. Multiplying by the common denominator gives a finite graph with nonnegative integer edge weights. After computing the functions, the verifier checks every inequality (99)–(100) directly with exact fractions.

21.2 Inserting an arbitrarily long affine interval

Lemma 21.2 (Affine insertion). *Let $A = b(b+1)$. Suppose a certificate at majority size H has an integer cut c such that*

$$c \geq A + 2D, \quad H - c \geq A + 2D + 2, \quad (102)$$

$$F_p(k) = 2k + \alpha_p \quad (c - 2D \leq k \leq c + 2D, \ p \in \{0, 1\}). \quad (103)$$

For every integer $\Delta \geq 0$, replace H by $H + \Delta$ in the profile formulas and define

$$F'_p(k) = \begin{cases} F_p(k), & k \leq c, \\ F_p(c) + 2(k - c), & c \leq k \leq c + \Delta, \\ F_p(k - \Delta) + 2\Delta, & k \geq c + \Delta. \end{cases} \quad (104)$$

Then F'_p satisfies the same charge inequalities. The initial lower bound (101) increases by exactly 4Δ .

Proof. The pieces of (104) agree at their endpoints. For either the base or enlarged profiles, every successor $j = g_p((k-r)_+)$ satisfies

$$|j - k| \leq D. \quad (105)$$

For a nonempty survivor, its surplus lies between -1 and $b+1$, and $r \leq m$. For an empty survivor, $k \leq m$ and $j = 0$.

If $k < c - D$, the enlarged profile agrees with the base profile at $q = (k-r)_+$. Indeed, the far-corner terms in both profile formulas exceed their plateau caps by (102). The successor is unchanged and is less than c by (105). Both potentials are unchanged, so the base inequality applies.

If $k > c + \Delta + D$, write $k_0 = k - \Delta$ and $q_0 = k_0 - r$. Since $q_0 > c + D - m = c + 1$, the low-corner terms have reached their plateau caps. The profile formulas therefore give

$$g'_p(q_0 + \Delta) = g_p(q_0) + \Delta.$$

Writing j_0 for the base successor of k_0 , we have $j = j_0 + \Delta$ and $j_0 > c$. Both potentials gain 2Δ , which cancels in their difference.

Finally suppose $c - D \leq k \leq c + \Delta + D$. The survivor $q = k - r$ is positive, with

$$q \geq c - D - m \geq A + 1, \quad H + \Delta - q \geq H - c - D \geq A + D + 2.$$

Both profiles are therefore on their affine plateaus: $j = k - r + \delta_p$, where $\delta_0 = b$ and $\delta_1 = b + 1$. By (105), both k and j lie in $[c - 2D, c + \Delta + 2D]$. Throughout that interval the corresponding potentials are $2k + \alpha_p$ and $2j + \alpha_{1-p}$. Their difference is

$$2r - 2\delta_p + \alpha_p - \alpha_{1-p}.$$

This is exactly the base potential difference at source $k = c$ with the same current color and quota r . Its survivor and successor lie in the same affine profile and potential collar, so the verified base inequality bounds the expression by $c_p(r)$.

These three cases exhaust the enlarged state space. At each full color class (104) adds 2Δ to the potential, giving the stated increase of 4Δ in the integer lower bound. $\square$

The matching upper construction has the same insertion property. A *serial strategy* directs the full budget to one initial cohort until it can be eliminated, uses any unused quota on that finishing day on the other cohort, and then directs the full budget to the other.

Lemma 21.3 (Insertion in a serial strategy). *Under (102), suppose a base serial strategy visits count c in each active cohort before an inspection. An enlarged rectangle with majority size $H + \Delta$ has a strategy taking exactly 4Δ more days.*

Proof. Under a full-budget inspection an active cohort's size never increases, because every neighborhood surplus is at most m . The certified visit to c is therefore reached from above, and the subsequent base trajectory stays at or below c . Above the cut, the enlarged active trajectory is the translation by Δ of the base trajectory, by the high profile translation identity in the preceding proof. The uninspected cohort remains its full current color class. When the first active cohort reaches $c + \Delta$, insert 2Δ full-budget days. Throughout this inserted interval the profile plateaus give

$$g_0(k - m) = k - 1, \quad g_1(k - m) = k.$$

Every two days reduce its size by one and restore its current color. The inserted block therefore ends at c with the phase of the base schedule unchanged. Follow the base trajectory below the cut.

The first cohort's finishing-day unused quota is unchanged. The second cohort consequently begins its active phase at the translated base state. This state is above the cut: it lies within $m + 1$ of a full color class, whereas (102) places c farther from that edge. Follow its translated trajectory until $c + \Delta$, insert another 2Δ full-budget days, and then finish along the unchanged lower trajectory. Both inserted blocks have even length, so the later phases of the schedule agree with those in the base strategy. $\square$

The two insertion proofs also apply to any other compatible profile family with color-class sizes $H, H - 1$ and a fixed budget m , provided that a corner radius A has the following properties: every surplus is between -1 and m ; the central profiles are $k + m - 1$ and $k + m$; the low profiles are unchanged when the remaining tail exceeds A ; and the high profiles translate when the source count exceeds A . These are exactly the properties used in the displacement, collar, and trajectory arguments. In that formulation one sets $D = m + 1$ and uses the same margin and affine-collar conditions. This observation will also give an exact three-dimensional family below.

The original exact-rational certificates for widths three through fifteen remain in the companion artifact. Their verifier is `src/odd_minimum_budget_all_lengths.py`; its receipt is `research/odd-minimum-budget-all-lengths.json`. They check every charge inequality, base potential, affine collar and attaining serial sweep; no numerical linear-programming output is trusted by the verifier. Those separate width calculations are now consequences of Theorem 21.1, so their charge tables are not needed for its proof. The generic certificate and insertion statements above retain their full conditional scope.

22 Eventual affine periods in odd cylinders

The interior-crossing argument extends to every fixed all-odd cross-section, in every dimension. The first step is an exact finite description of its two end regions; an affine plateau alone would not justify changing the cylinder's length.

Let Q be a fixed nontrivial box with odd side lengths. Put

$$W = |Q|, \quad S = \sum_i (\text{side}_i - 1), \quad Z = W(S + 1), \quad \beta = (W - 1)/2.$$

Consider $Q \square P_n$ with odd longitudinal length n . The one-point cross-section gives paths, already classified separately. For $n \geq S + 3$, the longitudinal coordinate is longest and is placed last in the order of Theorem 14.1. Write $H = (Wn + 1)/2$ for its majority class size and g_p^n for its exact color- p profile.

22.1 Two finite tables describe both ends

For $v \in Q_p$, let $r_p(v)$ be the rank of $(v, 0)$ in the parity order of the infinite prism $Q \times \{0, 1, \dots\}$. Every such point has weight at most S , so every preceding point has longitudinal coordinate at most S . Its rank is therefore independent of n for $n \geq S + 1$ and can be computed in the single finite slab $Q \square P_{S+1}$. At most $Z = W(S + 1)$ vertices have weight at most S , so $r_p(v) \leq Z$.

Using the transverse cost a_Q from (65), define

$$U_p(k) = \sum_{\substack{v \in Q_p \\ r_p(v) \leq k}} a_Q(v), \quad B_p(k) = |\{v \in Q_p : r_p(v) \leq k\}|.$$

Both tables are constant for $k \geq Z$.

Lemma 22.1 (Stable corner decomposition). *For $n \geq S + 3$ and $k > 0$,*

$$g_p^n(k) = k + \mathbf{1}_{\{p=1\}} + U_p(k) - |Q_p| + B_p(H - p - k), \quad g_p^n(0) = 0. \quad (106)$$

In particular, every surplus is at most $\beta + p$, and

$$k > Z, \quad H - p - k > Z \quad \implies \quad g_p^n(k) = k + \beta + p. \quad (107)$$

Proof. The ambient cost $a(x) - 1$ equals $a_Q(v)$ on the bottom slice, zero on an interior slice, and -1 on the top slice. This includes the origin: its ambient cost minus one is the transverse origin cost. The prefix-cost identity therefore counts the selected bottom weights through $U_p(k)$ and subtracts the number of selected top vertices.

Coordinate complement preserves parity and reverses both weight and lexicographic order. It carries the top slice to the bottom slice. Thus selected top vertices correspond to bottom vertices outside the prefix of size $H - p - k$, and their number is $|Q_p| - B_p(H - p - k)$. This proves (106). The empty prefix is separate because the extra origin-neighbor term requires a nonempty odd prefix.

The transverse full-class cost identity gives $U_p(\infty) = |Q_{1-p}| - \mathbf{1}_{\{p=1\}}$ and $B_p(\infty) = |Q_p|$. Nonnegativity of a_Q gives the asserted upper bound on surplus. Stabilizing both tables gives (107). $\square$

First slice	Every interior slice	Last slice																											
<table> <tr><td>2</td><td>1</td><td>0</td></tr> <tr><td>2</td><td>1</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>0</td></tr> </table>	2	1	0	2	1	0	1	1	0	<table> <tr><td>0</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>0</td></tr> </table>	0	0	0	0	0	0	0	0	0	<table> <tr><td>-1</td><td>-1</td><td>-1</td></tr> <tr><td>-1</td><td>-1</td><td>-1</td></tr> <tr><td>-1</td><td>-1</td><td>-1</td></tr> </table>	-1	-1	-1	-1	-1	-1	-1	-1	-1
2	1	0																											
2	1	0																											
1	1	0																											
0	0	0																											
0	0	0																											
0	0	0																											
-1	-1	-1																											
-1	-1	-1																											
-1	-1	-1																											

Local contribution to neighborhood surplus in a $3 \times 3 \times n$ odd cylinder.

Figure 6: The local contribution $a(x) - 1$ in a cylinder with 3×3 cross-section: transverse costs on the first slice, zero in interior slices, and -1 on the last slice. The surplus of a nonempty odd prefix also has the additional 1 in (106); that term is outside the local sum.

The decomposition gives the stronger identities needed to change length. For two valid majority sizes H and $H + \delta$, with both longitudinal lengths at least $S + 3$, we have

$$H - p - k > Z \implies g_p^{H+\delta}(k) = g_p^H(k), \quad (108)$$

$$k > Z \implies g_p^{H+\delta}(k + \delta) = g_p^H(k) + \delta. \quad (109)$$

In the first identity the tail table is full; in the second the bottom weight table is full and the tail argument is unchanged. The first also holds at $k = 0$, since both profiles are zero there. Superscripts now indicate majority size rather than longitudinal length.

An all-odd box has a spanning path whose two endpoints lie in its majority color: traverse successive slices alternately forward and backward and induct on dimension. On that odd path a proper majority k -set has at least k neighbors, by omitting an unselected majority vertex and matching selected vertices toward it from both sides. A nonempty minority k -set has at least $k + 1$ neighbors, by its consecutive blocks in the spacing-two path order. The box contains these path edges, so

$$g_0(k) \geq k \ (k < H), \quad g_1(k) \geq k + 1 \ (k > 0), \quad g_0(H) = H - 1. \quad (110)$$

Corollary 22.2 (Eventual feasibility threshold). *If $H \geq 2Z + 3$, then $h(Q \square P_n) = (W + 1)/2$.*

Proof. The size condition implies $n > S + 3$. The majority surplus is at most β and attains β at $k = Z + 1$ by (107). Apply Corollary 14.6. $\square$

This is an eventual threshold; shorter cylinders can need fewer inspections. For example, a $5 \times 5 \times 5$ box needs twelve, whereas a sufficiently long cylinder with 5×5 cross-section needs thirteen.

22.2 The period and an explicit threshold

Fix a feasible eventual budget $m \geq \beta + 1$ and set

$$d = 2m - W > 0, \quad g = \gcd(W, d), \quad \Delta = \text{lcm}(W, d) = Wd/g.$$

For the explicit threshold define

$$\begin{aligned} J &= m + 1, & A &= Z + 1, & R &= \Delta + 2J, \\ B &= 8Z + 6d + 2, & L &= Z + m + 1, & \mathcal{E} &= B + 4L(B + 1), \\ H_* &= R + 2A + 4J + 2 + 2\mathcal{E}(R + 2J + 1). \end{aligned} \tag{111}$$

Theorem 22.3 (Eventual affine period in every odd cylinder). *For every odd n with $(Wn + 1)/2 \geq H_*$,*

$$T_m(Q \square P_{n+2d/g}) = T_m(Q \square P_n) + 4W/g. \tag{112}$$

At the eventual minimum budget $m = (W + 1)/2$, increasing a sufficiently large odd length by two increases the optimal time by exactly $4W$.

The threshold is not intended to be sharp. The proof uses the profile bounds, plateau and stable translations just proved, together with connectedness.

22.3 A potential with a bounded total deficit

For $n \geq S + 3$, write $H = (Wn + 1)/2$ and $M = H - 1$. Define

$$K = 2Z + 2m + 1, \quad \Gamma = (2H - 4Z - 2m - 5)_+, \quad F_p(k) = \min\{\Gamma, (2k + p - K)_+\}.$$

Lemma 22.4 (Uniform potential and time bounds). *For $0 \leq r \leq m$,*

$$F_p(k) - F_{1-p}(g_p((k - r)_+)) \leq (2r - W)_+. \tag{113}$$

Consequently

$$\lceil 2\Gamma/d \rceil \leq T_m(Q \square P_n) \leq 4 \lceil (M - m)_+/d \rceil + 2. \tag{114}$$

Proof. The potential inequality is trivial if $\Gamma = 0$. Otherwise let $q = (k - r)_+$ and $j = g_p(q)$. If $q \leq Z$, then $k \leq Z + m$ and $2k + p - K \leq 0$, so the current potential is zero. If $q \geq H - Z - 1$, then $j \geq q - 1 \geq H - Z - 2$, and

$$2j + (1 - p) - K \geq 2H - 2Z - 4 - K = \Gamma.$$

The successor potential is therefore Γ . In the remaining region both profiles are on their plateaus, so $j = k - r + \beta + p$. The difference of the unclipped affine expressions is $2r - W$. Common monotone, 1-Lipschitz clipping proves (113).

For two quota shares $r, m - r$, the right sides sum to at most $d = 2m - W$. Both full color classes have potential Γ , giving the lower bound by summing daily decreases. For the upper bound, start with the minority cohort of size M . Each pair of full-budget days reduces an uncleared cohort by at least d . Thus $2 \lceil (M - m)_+/d \rceil + 1$ days suffice. Pad early clearing to this fixed odd phase length. The other initial cohort remains its full current color class and, after the odd phase, is also in the minority color with size M . Repeat the phase to obtain the upper bound. $\square$

For a rectangle $Q = P_w$, $w = 2b + 1$, the explicit profiles of Theorem 5.1 permit the sharper corner radius $Z = b^2$ in the preceding argument. At minimum budget, strengthen its potential slightly as follows, for every odd $n \geq w$. Take $K = 2b^2 + 2b + 2$ and $\Gamma = (2H - 4b^2 - 2b - 6)_+$. In the low region the current potential is zero for $r < m$ and at most one for $r = m$; the high-region and plateau arguments are unchanged. This gives the useful uniform estimate

$$\max\{0, 2wn - 2w^2 + 2w - 10\} \leq T_{(w+1)/2}(P_w \square P_n) \leq 2wn - 2w - 2. \quad (115)$$

Thus its leading term is $2wn$ for every fixed odd width, independently of whether a sharper correction has been determined.

Return to the general cylinder, with $Z = W(S + 1)$ and the potential of Lemma 22.4. Assume $H \geq H_*$, so $\Gamma > 0$ and $M > m$. Fix an optimal prefix strategy, using padded full quotas as in Section 21. Let Φ_t be the sum of its two current-color potentials and put

$$\eta_t = d - (\Phi_t - \Phi_{t+1}) \geq 0.$$

These are integers, and the upper bound in (114) gives

$$\sum_t \eta_t = dT_m - 2\Gamma \leq 8Z + 6d + 2 = B. \quad (116)$$

Here $d \lceil (M - m)/d \rceil \leq M - m + d - 1$ accounts for the rounding. Call a day *efficient* if $\eta_t = 0$, and *bad* otherwise. There are at most B bad days.

An efficient day gives all m inspections to one cohort. Indeed, if both positive-part losses in (113) are positive, their sum is $d - W < d$. If only one is positive but neither share is m , it is at most $d - 2$. Thus on an efficient day the active cohort loses exactly d potential and the inactive cohort's potential remains exactly constant.

22.4 From flat potentials to empty or full cohorts

Zero inspections strictly increase every intermediate potential:

$$0 < F_p(k) < \Gamma \implies F_{1-p}(g_p(k)) > F_p(k). \quad (117)$$

For a proper majority prefix, $g_0(k) \geq k$, so its unclipped expression increases by at least one after the color reversal. A full majority prefix has potential Γ and is excluded. For a nonempty minority prefix, $g_1(k) \geq k + 1$, giving the same conclusion. Clipping preserves strict growth while the old value is below Γ .

After an active cohort loses d on an efficient day, its potential is strictly below Γ . If it is still positive, (117) forces it to stay active on the next efficient day. A switch is possible only when its potential reaches zero. It cannot become active again in the same efficient run, since zero cannot lose d and its inactive value must stay constant. Hence each consecutive efficient run has at most two focused blocks, and there are at most $2(B + 1)$ such blocks in total.

Call a day *clean* when all m inspections target one cohort and the other cohort is actually empty or actually its full current color class. Flat potential alone is insufficient for this conclusion. We now bound the transient days needed to reach actual emptiness or fullness.

If a nonempty proper support X lies in one color of a connected bipartite graph without isolated vertices, then $X \subsetneq N^2(X)$. Inclusion follows by backtracking along an edge. Equality would make X closed under two-step paths; connectivity makes the two-step graph connected within each color

and would force X to be full. Therefore an uninspected nonempty proper prefix grows by at least one every two days until it is full.

During an efficient focused block the inactive potential is constantly zero or Γ . At zero, its count is at most $Z + m = L - 1$; if nonempty, it cannot remain in that range for more than $2L$ days. At Γ , its count is at least $H - Z - 2$, hence within L of its full class; it becomes full within $2L$ days. Empty and full cohorts remain so under further uninspected moves. Each efficient block therefore contributes at most $2L$ nonclean days. Counting all bad days as well, the entire optimal search has at most

$$\mathcal{E} = B + 4L(B + 1) \quad (118)$$

nonclean days.

22.5 A protected band and its two individual cuts

Every one-day count change has absolute value at most $J = m + 1$: profile surplus is between -1 and $\beta + 1$, and each quota is at most m . Record both source counts from each nonclean day. A count k forbids integer cuts c with $k \in [c - J, c + R + J]$, at most $R + 2J + 1$ cuts. There are at most $2\mathcal{E}$ records. Restrict candidate cuts to

$$A + 2J \leq c \leq H - R - A - 2J - 2. \quad (119)$$

There are $H - R - 2A - 4J - 1$ candidates. By (111), this exceeds $2\mathcal{E}(R + 2J + 1)$. Choose a cut forbidden by no record.

No nonclean transition can touch the protected interval $[c, c + R]$. A cohort in that interval is therefore active on every day, receiving all m inspections; an inactive clean cohort is empty or full, both outside the interval. Its count never increases. Nor can it cross the interval upwards: a nonclean transition cannot touch it, and a clean active transition does not increase. Both cohorts begin above and finish below the interval.

For each *initial* cohort i , choose its first count x_i at most $c + \Delta + J$. The previous count is larger, so the displacement bound gives

$$c + \Delta < x_i \leq c + \Delta + J. \quad (120)$$

The entire protected interval is inside the profile plateaus, including all relevant survivors. A full-budget step there sends k to $k - m + \beta + p$, so every pair of days reduces its count by exactly d and restores its color. The next

$$\tau = 2\Delta/d$$

days therefore take x_i to $x_i - \Delta \in (c, c + J]$. This duration is even, and the trajectory stays in the protected band. Throughout the block its inactive companion remains empty or full. The two crossing blocks cannot overlap.

Define the individual endpoints

$$\ell_i = x_i - \Delta, \quad u_i = x_i.$$

They need not agree for the two cohorts. This avoids an otherwise real residue problem: when $d > 1$, a trajectory may skip a prescribed count, and the two cohorts need not have matching residues. One common protected band supplies two legitimate individual cuts. Every visit to (ℓ_i, u_i) belongs to the chosen crossing block, because active clean trajectories are nonincreasing and nonclean transitions cannot touch the larger band. Boundary stalls at the minimum budget are harmless; the even block duration preserves the needed phase.

22.6 Removing and enlarging the crossing blocks

Proof of Theorem 22.3. First contract the majority size from H to $H - \Delta$ and delete both crossing blocks of length τ . At each remaining state transform initial cohort i by

$$f_i(k) = \begin{cases} k, & k \leq \ell_i, \\ k - \Delta, & k \geq u_i. \end{cases}$$

No retained state lies strictly between these cases. At the ends of a deleted block, u_i and ℓ_i both map to ℓ_i . Its inactive companion is empty or full at both ends; a full class maps to the contracted full class. Since τ is even, the color labels also match. The two state paths therefore glue.

For a low retained source, its survivor is at most ℓ_i . The contracted far-distance satisfies

$$(H - \Delta) - \ell_i = H - u_i \geq H - (c + R) \geq A + 2J + 2,$$

so the low profile identity (108) applies. Its successor cannot enter the removed interval, by the protected-band argument. For a high retained source, the contracted survivor is at least

$$\ell_i - m > c - m > Z.$$

The stable upper profile identity (109) therefore gives

$$g_p^H(q + \Delta) = g_p^{H-\Delta}(q) + \Delta.$$

These identities verify every retained transition. Quotas remain within the same budget, and unused quota may be wasted. The result is a physical prefix search on length $n - 2\Delta/W$. Its majority size is at least $H_* - \Delta \geq 2Z + 6J + 4$, so its length exceeds $4(S + 1)$ and the longitudinal coordinate remains longest. Thus all stable profile identities still apply. The resulting search is completed in $T_m(Q \square P_n) - 4\Delta/d$ days. Hence

$$T_m(Q \square P_{n-2\Delta/W}) \leq T_m(Q \square P_n) - 4\Delta/d. \quad (121)$$

Conversely, start with an optimal search on any $H \geq H_*$ and its protected blocks. Enlarge H to $H + \Delta$. Keep lower counts at or below ℓ_i unchanged and increase upper counts at or above u_i by Δ . Replace each block from u_i to ℓ_i by a clean block from $u_i + \Delta$ to ℓ_i , lasting $4\Delta/d$ days. The affine profiles realize this trajectory explicitly, with the inactive companion empty or the enlarged full class. The extra τ days per block are even. The same low and high identities verify every other transition. Thus

$$T_m(Q \square P_{n+2\Delta/W}) \leq T_m(Q \square P_n) + 4\Delta/d.$$

Apply (121) at $H + \Delta$ for the reverse inequality. Since $\Delta/W = d/g$ and $\Delta/d = W/g$, this proves (112). $\square$

Let n_0 be the least odd length at or beyond the threshold. For each of the d/g odd residue classes modulo $2d/g$, the exact two-count recurrence determines the optimum and an optimal strategy at one of

$$n_0, n_0 + 2, \dots, n_0 + 2d/g - 2.$$

The theorem then determines every later time in that class and constructs its optimal strategy by insertion. In particular,

$$T_m(Q \square P_n) = \frac{2W}{2m - W} n + O_{Q,m}(1).$$

At the eventual minimum budget there is one offset $C(Q)$, with $T_{(W+1)/2}(Q \square P_n) = 2Wn - C(Q)$ for every sufficiently large odd n . The offset can depend on the transverse shape, not only on its number of rooms. For rectangular cross-sections $Q = P_w$, the stronger bounds in (115) give $2w + 2 \leq C(w) \leq 2w^2 - 2w + 10$. Finding smaller corner recurrences or closed expressions for the eventual offsets remains a further problem.

22.7 An exact three-dimensional family

For one cross-section the finite corner calculation gives the optimum from the shortest nontrivial cylinder, without the conservative threshold in Theorem 22.3.

Theorem 22.5. *For every odd $n \geq 3$, five inspections per day are necessary and sufficient on $P_3 \square P_3 \square P_n$, and*

$$T_5(P_3 \square P_3 \square P_n) = 18n - 36.$$

Proof. Put $H = (9n + 1)/2$. For $Q = P_3 \square P_3$, the ranks of bottom-slice vertices in their respective parity orders, and their transverse costs, are

p	ranks	costs
0	(1, 2, 3, 5, 8)	(2, 1, 1, 0, 0)
1	(1, 2, 4, 6)	(2, 1, 1, 0).

These ranks already stabilize at $n = 3$. The last even bottom point $(2, 2, 0)$ is first in weight layer four, and its earlier even layers have longitudinal coordinate at most two. The last odd bottom point $(1, 2, 0)$ is sixth in its parity order: its predecessors are the three unit vectors and $(2, 1, 0)$, $(2, 0, 1)$. Enlarging the longitudinal path therefore introduces no earlier point in either list.

Let $U_p(k)$ be the cumulative listed cost through rank k , and let $B_p(k)$ count the listed ranks at most k . The cost and complement argument of Lemma 22.1 gives, for $k > 0$,

$$g_0(k) = k + U_0(k) - 5 + B_0(H - k), \quad g_1(k) = k + 1 + U_1(k) - 4 + B_1(H - 1 - k),$$

with $g_0(0) = g_1(0) = 0$. Thus the corner radius is eight, and the central profiles are $k + 4$ and $k + 5$. The majority surplus is at most four and equals four at $k = 3$: $U_0(3) = 4$ and $B_0(H - 3) = 5$ for every $H \geq 14$. Corollary 14.6 gives the minimum budget five.

For the exact time use the identical charge sequences

$$c_0 = c_1 = (0, 0, \frac{1}{3}, \frac{2}{3}, 1, 1).$$

They satisfy the daily charge bound (99). The rational shortest-charge potentials described in the preceding section are checked against every transition

$$F_p(k) - F_{1-p}(g_p(\max\{k - r, 0\})) \leq c_p(r), \quad 0 \leq r \leq 5.$$

Together with a minority-first serial strategy, the five finite certificates give

n	3	5	7	9	11
$F_0(H) + F_1(H - 1)$	18	54	90	126	162
strategy length	18	54	90	126	162.

Here the strategy spends all available inspections on the first cohort until it is captured, uses any remaining inspections on its finishing day on the second cohort, and then finishes that cohort. The

finite verifier checks the physical room neighborhoods and replays these inspection sets, as well as checking all 1,950 rational inequalities.

At the base $n = 11$, take $H = 50$, $c = 25$, and $D = 6$. The exact arrays satisfy

$$F_p(k) = 2k - 14 + p \quad (13 \leq k \leq 37).$$

This is the full collar $[c - 2D, c + 2D]$, and $c \geq 8 + 2D$, $H - c \geq 8 + 2D + 2$. Both active cohorts visit count 25. At the first visit the other cohort is full, since $25 > 5$ places the visit before the first cohort's finishing day. At the second visit the first cohort is empty.

Consequently Lemmas 21.2 and 21.3, in their compatible-profile formulation, apply with corner radius eight. Increasing H by Δ inserts an affine interval in each potential and raises their initial sum by 4Δ . Inserting 2Δ full-budget days at each serial cut realizes the same increase in time. Each pair of inserted days reduces the active count by one and restores its color, while the other cohort stays full or empty. All remaining transitions follow the stable low and translated high profiles above.

For every odd $n \geq 11$, choose $\Delta = 9(n - 11)/2$. The matching bounds are then $162 + 4\Delta = 18n - 36$. The finite certificates cover the smaller odd lengths, completing the proof. $\square$

The exact verifier and receipt are included in the companion artifact:

```
src/three_by_three_cylinder_time.py
research/three-by-three-cylinder-time.json
```

The receipt stores the corner tables, charges, finite values, and complete base potential arrays. An independent verifier also checked actual room neighborhoods and schedules at seven lengths, including 13 and 31. These checks establish the finite premises of the insertion proof. The variable-length theorem is an ordinary proof with exact rational certificates; it is not claimed here as a complete physical Lean theorem.

23 Higher-dimensional boxes

The preceding compression theorems work in arbitrary dimension, but do not by themselves provide a closed time formula for every box. We give two unbounded applications and a fully certified three-dimensional example.

23.1 A side of length two

Let H be bipartite, with color function χ , and put $G = H \square P_2$. In either color class of G , there is exactly one vertex above each vertex v of H : its second coordinate is $s - \chi(v)$ modulo two. Under this identification,

$$\pi(N_G(R)) = \pi(R) \cup N_H(\pi(R)).$$

The vertical move preserves the projection, and an H -move changes it to an H -neighbor. This proves the equality for arbitrary sets, including boundary and empty sets.

Suppose H has an order with prefixes I_k such that every k -set has at least $c(k)$ vertices in its closed neighborhood and $I_k \cup N_H(I_k) = I_{c(k)}$. Put $A = |H|$. The two initial-color cohorts then have the exact state transition

$$(a, b) \longmapsto (c(\max\{a - p, 0\}), c(\max\{b - m + p, 0\})), \quad 0 \leq p \leq m, \quad (122)$$

starting at (A, A) . A state is capturable on the current day exactly when $a+b \leq m$. Tail inspections of the prefixes attain each step. Conversely the cardinalities in every unrestricted search dominate these counters when the same allocations are used. Thus shortest-path search on $(A+1)^2$ states returns both the exact time and an actual optimal strategy, using $O((m+1)(A+1)^2)$ transitions after the profile is known.

Theorem 23.1. *Under the preceding closed-neighborhood nesting hypothesis,*

$$h(H \square P_2) = 1 + \max_{0 \leq k \leq A} (c(k) - k).$$

Both this feasibility formula and recurrence (122) apply to every Cartesian box having a side of length two.

Proof. Write $B = \max_k (c(k) - k)$. If $m \leq B$, select k with $c(k) \geq k + m$. A cohort of size at least $k + m$ retains at least k possible positions after inspection and has at least $k + m$ after movement. Initially $A \geq k + m$. Since $k > 0$, the invariant prevents capture forever. If $m > B$, a canonical cohort of size $b > m$ has next size at most $b - m + B < b$. Clear it, then clear the other cohort in the same way; the untouched projected cohort stays full.

For Cartesian H , omit length-one factors, list the other side lengths in increasing order, and order vertices by increasing coordinate sum, breaking a tie by putting the larger first differing coordinate first. The classical simplicial isoperimetric theorem gives exactly the closed-neighborhood nesting hypothesis. We use the statement and definitions in Otachi–Suda [13, Theorem 2.5], where the result is attributed to Moghadam and Bollobás–Leader. For the one-vertex empty product take $c(0) = 0, c(1) = 1$. $\square$

For $a, b \geq 2$, the vertex boundary width of $P_a \square P_b$ is $\min(a, b)$, so $h(P_a \square P_b \square P_2) = \min(a, b) + 1$. If exactly one of a, b is one, the remaining nontrivial ladder has threshold two; if both are one, it is a single edge with threshold one. The theorem also covers all hypercubes. Its input is a classical *closed*-neighborhood theorem; it does not assert open-neighborhood nesting on arbitrary even rectangles.

23.2 Exact profile transfer for either longitudinal parity

There is a useful further recurrence even when global minimizing prefixes are unavailable. Let a bipartite transverse graph Q have compatible minimizing parity orders with profiles g_0, g_1 . For a one-color support in $Q \square P_n$, let k_j be its size in column $j = 0, \dots, n-1$, and let p be its global color.

Proposition 23.2. *For each prescribed column-count vector, the exact minimum neighborhood size is*

$$\sum_{j=0}^{n-1} \max\{g_{(p+j) \bmod 2}(k_j), k_{j-1}, k_{j+1}\}, \quad k_{-1} = k_n = 0.$$

This holds for every positive n , including even n .

Proof. In column j , the neighborhood is the union of the transverse neighborhood of that column and the supports in the two adjacent columns. Its size is at least the displayed maximum. Replace every column support by the corresponding transverse prefix. All three sets are now prefixes of the same opposite color order, and their union has size exactly the maximum. These replacements attain all column minima simultaneously. $\square$

Minimize the sum over vectors with $\sum k_j = k$ to obtain the exact global cardinality profile. A transfer state remembers the previous count a , current count b , and cumulative count. Appending c charges $\max\{g(b), a, c\}$ and replaces the pair by (b, c) . Zero counts at the two outside columns give the boundary conditions. Writing $W = |Q|$, this direct recurrence uses $O(n^2 W^4)$ arithmetic operations and $O(n W^3)$ memory once the transverse profiles are known. It applies in particular to every all-odd cross-section.

This computes exact neighborhood minima, not compatible global orders. The minimizing vector can depend on k without being nested. A search using only these cardinality profiles therefore supplies a lower time bound; it is not asserted to attain that bound on arbitrary even cylinders. The accompanying code checks the recurrence against all 524 288 one-color supports of the $3 \times 3 \times 4$ box.

23.3 The finite $3 \times 3 \times 4$ check as a corollary

Proposition 23.3. *On $P_3 \square P_3 \square P_4$, the optimal capture time with five inspections per day is 36.*

Proof. This is $n = 4$ in the all-length theorem 23.5: $18n - 36 = 36$. □

An independent finite verification remains available. Slice compression reduces the physical game to eight counts; exhaustive breadth-first search first leaves at most five rooms after 35 movements. It checks 1,109,364 transitions and discovers 2,017 states. A separately implemented physical-neighborhood search obtains the same layers and optimum, and its saved 36-day schedule replays on the actual room graph. The complete recurrence and verifier are in `src/three_by_three_by_four_exact.py`, with receipt and witness in `research/three-by-three-by-four-exact.json`. This supplies an independent finite check of the general theorem; repeating its full finite proof is unnecessary. It is not a complete physical Lean theorem.

23.4 Two complete examples with even sides

The distinction between a minimum neighborhood size and a sequence of compatible minimizing shapes is visible even on small cubes. The following two tables are exact ordinary computer-assisted theorems. Their lower certificates and physical schedules have been checked by a second implementation, independently of the search that found them.

Theorem 23.4. *The minimum capture times on P_4^3 and $P_3 \square P_4^2$ are as follows.*

$4 \times 4 \times 4$		$3 \times 4 \times 4$	
Daily budget	Minimum days	Daily budget	Minimum days
0–7	∞	0–6	∞
8	40	7	20
9	20	8	14
10	16	9	11
11	12	10	8
12	10	11	7
13–14	8	12–13	6
15	7	14	5
16–17	6	15–19	4
18	5	20–23	3
19–25	4	24–47	2
26–31	3	48 or more	1
32–63	2		
64 or more	1		

Proof. Apply Theorem 3.6 to every equal-length pair, and also compress the length-three coordinate in the second box. The one-color fixed families contain respectively 292 and 1936 sets per color. These families are enumerated without a geometric guess: order vertices by increasing $\sum_i ix_i$, and either omit each vertex or include it when all its legal predecessors are already present. The predecessor moves are (3.6), together with $x \mapsto x - 2e_i$ on an odd axis. This recursively enumerates every ideal exactly once. Direct physical neighborhoods preserve the fixed families.

Minimizing $|N(R)|$ at each cardinality in these families gives the exact unrestricted one-color neighborhood profiles, by compression. They give an impossibility trap below budgets eight and seven, respectively. On the four-cube the profile is the same in both colors and is

$$g = (0, 3, 6, 7, 9, 11, 12, 13, 15, 16, 17, 18, 19, 19, 21, 22, 23, \\ 24, 25, 25, 26, 27, 28, 28, 29, 29, 30, 31, 31, 31, 32, 32, 32).$$

For example, a cohort of size at least 21 retains at least 14 rooms after seven inspections, and then has at least $g(14) = 21$ positions again.

The two-count lower relaxation formed from the exact profiles gives every listed finite lower bound except at budget eight in either box. To verify this statement, enumerate all count pairs, all quota splits, and successive sets of pairs that can reach capture. This calculation is the same monotone finite recurrence used for the three-cube above. Actual room-coordinate schedules attain every listed bound; the verifier updates the complete physical belief set after each inspection and movement. Budget monotonicity extends endpoint schedules across each displayed range.

The two exceptional lower bounds have short certificate descriptions. For each fixed one-color set A let $F(A)$ be the nonnegative integer in the accompanying table. For every fixed survivor $R \subseteq A$, put $p = |A| - |R|$. At budget eight the tables satisfy

$$F(A) \leq c(p) + F(N(R)), \quad 0 \leq p \leq 8, \quad F(\emptyset) = 0. \quad (123)$$

On P_4^3 , the charges at $p = 0, \dots, 8$ are $(0, 0, 0, 0, 1, 2, 2, 2, 2)$; both full-color potentials are 40. Two allocations totaling at most eight have total charge at most two. Telescoping (123) therefore gives $T_8 \geq (40 + 40)/2 = 40$. There are 29 686 inequalities to check. An attaining prefix sweep for each

cohort has sizes

32, 29, 27, 25, 24, 23, 22, 22, 21, 21, 19, 19,
18, 18, 17, 16, 15, 13, 11, 8, 0.

Each sweep takes twenty days; reflection in an even axis gives the favorable starting phase for the second cohort.

On $P_3 \square P_4^2$, take $c(p) = p$ and full-color potential 54. The 654 570 inequalities give $T_8 \geq \lceil 108/8 \rceil = 14$, attained by the stored physical schedule. These certificate checks require only integer inequalities and direct finite neighborhoods. They do not trust the shortest-path procedure that generated the potential tables.

Finally, both boxes have a perfect matching. Its disjoint alternating trajectories require at least half the volume in two days; two successive inspections of one full color attain that threshold. One-day capture requires the entire volume. This proves the remaining endpoint ranges. $\square$

The independent verifier is `src/even_cube_resumed_independent_review.py`; complete certificates, room schedules, and review receipts are in the corresponding `research/even-resume-*` files in the research archive. At eight probes the four-cube’s cardinality relaxation predicts only 32 days, compared with the true 40; on $3 \times 4 \times 4$ it predicts 12 instead of 14. Thus exact neighborhood profiles alone need not determine exact capture times. The shape information retained by the compression is mathematically necessary for these lower arguments.

23.5 The complete $3 \times 3 \times 3$ example

Encode a room $(x, y, z) \in \{0, 1, 2\}^3$ by $9x + 3y + z$. Its color is $x + y + z$ modulo two. The two classes have sizes 14, 13. Their minimum open-neighborhood sizes are the following; a dash indicates a source cardinality larger than the class.

Layer $x = 0$	Layer $x = 1$	Layer $x = 2$																											
<table border="1"> <tr><td>0</td><td>1</td><td>2</td></tr> <tr><td>3</td><td>4</td><td>5</td></tr> <tr><td>6</td><td>7</td><td>8</td></tr> </table>	0	1	2	3	4	5	6	7	8	<table border="1"> <tr><td>9</td><td>10</td><td>11</td></tr> <tr><td>12</td><td>13</td><td>14</td></tr> <tr><td>15</td><td>16</td><td>17</td></tr> </table>	9	10	11	12	13	14	15	16	17	<table border="1"> <tr><td>18</td><td>19</td><td>20</td></tr> <tr><td>21</td><td>22</td><td>23</td></tr> <tr><td>24</td><td>25</td><td>26</td></tr> </table>	18	19	20	21	22	23	24	25	26
0	1	2																											
3	4	5																											
6	7	8																											
9	10	11																											
12	13	14																											
15	16	17																											
18	19	20																											
21	22	23																											
24	25	26																											

First day of the optimal 18-day, five-probe schedule.

Figure 7: The three-cube drawn as three layers. Adjacent rooms in a layer are connected; rooms at the same position in consecutive layers are also connected. Shading marks the first five inspections of the schedule.

k	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
$g_E(k)$	0	3	5	7	8	9	10	10	11	12	12	13	13	13	13
$g_O(k)$	0	4	6	7	9	10	11	12	12	13	13	14	14	14	–

Here is a finite certificate procedure for the table. For each color, list its vertices; recursively either include or exclude the next vertex, maintaining the selected count k and the union U of its physical neighbor sets. At every leaf check $|U| \geq g(k)$. This has $2^{14} + 2^{13} = 24576$ leaves. The generic recursion is proved sound in Lean before the concrete finite check is kernel evaluated. Prefixes in increasing coordinate sum with decreasing lexicographic tie breaking attain the values. Neither the lower bound nor the formal classification assumes that an arbitrary search uses these prefixes.

For a fixed budget define the relaxed count successor

$$f_p(a, b) = (g_O(\max\{b - m + p, 0\}), g_E(\max\{a - p, 0\})).$$

All arbitrary physical successors dominate one of these. The finite calculation below independently specifies the time lower bounds. It uses just $15 \cdot 14 = 210$ count pairs; W_t denotes relaxed states that can reach zero in at most t inspection rounds.

```

W = {(0, 0)}
for t = 1, 2, ...:
  Wnext = W union {(a,b): 0<=a<=14, 0<=b<=13,
                        f_p(a,b) belongs to W for some 0<=p<=m}
  if (14,13) belongs to Wnext: return t
  if Wnext == W: return infinity
W = Wnext

```

The required lower endpoint checks are 18, 10, 6, 4, 3 at budgets 5, 6, 8, 11, 12, respectively. In Lean the corresponding distance potentials are stored as finite integers and checked to be nondecreasing in both counts, zero at zero, and to drop by at most one under every relaxed transition. This proves the lower bound without trusting the search routine that found the potentials. For budgets at most four, there is a shorter proof: every same-color triple has at least seven neighbors, so a cohort of at least seven positions can never fall below seven after four inspections and movement.

For clarity, explicit upper schedules are given next. Each row is a sequence of daily inspection sets; a superscript $\times 2$ means repeat the entire listed sequence twice. The physical update $B \leftarrow N(B \setminus S)$, starting from all 27 rooms, verifies every schedule directly.

m	Inspection sets in order
5	$(\{1, 3, 5, 9, 11\}, \{4, 6, 10, 12, 18\}, \{5, 7, 11, 13, 15\},$ $\{8, 10, 12, 14, 16\}, \{9, 11, 13, 15, 17\}, \{10, 12, 14, 16, 18\},$ $\{11, 13, 15, 19, 21\}, \{8, 14, 16, 20, 22\}, \{15, 17, 21, 23, 25\})^{\times 2}$
6	$(\{8, 14, 16, 20, 22, 26\}, \{5, 7, 11, 13, 19, 25\}, \{2, 4, 10, 16, 22, 24\}, \{1, 7, 13, 15, 19, 21\},$ $\{0, 4, 6, 10, 12, 18\})^{\times 2}$
7	$(\{8, 14, 16, 20, 22, 24, 26\}, \{5, 7, 11, 13, 15, 19, 21\}, \{0, 2, 4, 6, 10, 12, 18\})^{\times 2}$
9	$\{2, 4, 8, 14, 16, 20, 22, 24, 26\}, \{1, 3, 7, 9, 11, 13, 15, 19, 21\},$ $\{5, 7, 11, 13, 15, 17, 19, 23, 25\}, \{0, 2, 4, 6, 10, 12, 18, 22, 24\}$
12	$\{2, 4, 6, 8, 10, 12, 14, 16, 20, 22, 24, 26\}, \{1, 3, 8, 9, 14, 16, 19, 20, 21, 22, 24, 26\},$ $\{1, 3, 5, 7, 9, 11, 13, 15, 19, 21\}$
13	All odd-numbered rooms, on each of two consecutive days.

Larger budgets inherit these schedules. Fewer than 27 probes cannot capture every possible initial room in one day, and 27 can. We have therefore proved the complete table stated in the introduction.

Its geometry, finite inequalities, physical schedules, and interpretation as capture of every actual target walk are all checked in Lean.

23.6 The exact five-inspection time on every $3 \times 3 \times n$ cylinder

Theorem 23.5. *For every $n \geq 3$,*

$$T_5(P_3 \square P_3 \square P_n) = 18n - 36.$$

The proof combines an explicit physical sweep with an ordinary lower certificate for arbitrary supports. The new even-length argument is not presently a complete physical Lean theorem.

The odd case is Theorem 22.5. Assume $n = 2L \geq 4$ and put $h = 9L$. Both checkerboard classes have h rooms. The compatible transverse 3×3 profiles are

$$G_0 = (0, 2, 3, 4, 4, 4), \quad G_1 = (0, 3, 4, 5, 5).$$

Compress every transverse slice to its parity prefix. This operator C is monotone and cardinality preserving, and satisfies $N(CS) \subseteq C(NS)$ by the fiber-compression theorem. For a compressed phase- p set with slice counts k_j , its exact neighborhood count in slice j is

$$\max\{G_{(p+j) \bmod 2}(k_j), k_{j-1}, k_{j+1}\}, \quad k_{-1} = k_n = 0.$$

We first establish a sharp obstruction to consecutive small boundaries.

23.6.1 A finite word certificate for every even length

For phase zero, pair consecutive counts into letters (a, b) with $0 \leq a \leq 5$, $0 \leq b \leq 4$. Give a letter the cost

$$v(a, b) = \max\{G_0(a), b\} - a + \max\{G_1(b), a\} - b,$$

and consecutive letters $s = (a, b)$, $t = (c, d)$ the edge cost

$$\begin{aligned} e(s, t) = & \max\{G_1(b), a, c\} - \max\{G_1(b), a\} \\ & + \max\{G_0(c), d, b\} - \max\{G_0(c), d\}. \end{aligned}$$

The neighborhood surplus is exactly $\sum_i v(s_i) + \sum_{i < L-1} e(s_i, s_{i+1})$; no outside edge is added. Every edge cost is nonnegative. The vertex costs are the following nonnegative matrix, with rows $a = 0, \dots, 5$ and columns $b = 0, \dots, 4$:

$$\begin{pmatrix} 0 & 3 & 4 & 5 & 5 \\ 2 & 3 & 3 & 4 & 4 \\ 3 & 3 & 3 & 3 & 3 \\ 4 & 3 & 3 & 3 & 2 \\ 4 & 3 & 2 & 2 & 1 \\ 4 & 3 & 2 & 1 & 0 \end{pmatrix}.$$

Thus a word of surplus at most four has no prefix of larger cost.

Here is the entire finite certificate specification. A state records its last letter, cost at most four, occupied count capped at four, omitted count capped at eight, whether its first $a \geq 3$, and length capped at two. Initialize with every one-letter word of cost at most four. Appending $t = (c, d)$

adds $v(t) + e(s, t)$ to the cost, $c + d$ to the occupied count, and $9 - c - d$ to the omitted count; saturate the two counts, preserve the first-letter flag, and cap length at two. Discard costs above four. Finite closure gives exactly 113 states and 346 edges. Every terminal state of capped length two satisfies:

1. Its cost is at least zero if either capped count is zero, and otherwise at least $\min\{4, K + 1, \lfloor (D + 2)/2 \rfloor\}$, where K, D are its two capped counts.
2. If its cost is four, $K = 4$, and $D = 8$, its first $a \geq 3$ and its last letter is $(a, 0)$ with $a \leq 2$.

These are finite integer checks of the stated initialization and transition rule, supplied in the accompanying certificate. Induction on word length makes them valid for every $L \geq 2$; costs above four trivially satisfy the first bound. Reflection in the even longitudinal side exchanges the colors, giving phase one too.

Consequently every monochromatic set of size k has at least $g_h(k)$ neighbors, where

$$g_h(0) = 0, \quad g_h(h) = h, \quad g_h(k) = k + \min\{4, k + 1, \lfloor (h - k + 2)/2 \rfloor\} \quad (0 < k < h). \quad (124)$$

Compression transports this bound to arbitrary supports. Call a survivor *critical* if its surplus is four and $4 \leq k \leq h - 8$. A compressed phase-zero critical set has at most two neighbors in the last slice, whereas every compressed phase-one critical set contains at least three rooms in that slice. Hence consecutive compressed critical survivors are impossible. For arbitrary critical R , (124) forces $|N(CR)| = |N(R)| = |R| + 4$. The compression inclusion therefore becomes $N(CR) = C(NR)$. If a second critical S lay in $N(R)$, monotonicity would give $CS \subseteq C(NR) = N(CR)$, the same contradiction. This proves the obstruction for arbitrary consecutive survivors.

23.6.2 A two-component integer potential

Let $e \in \{0, 1\}$ record whether the previous survivor was critical, starting with $e = 0$. A marked count satisfies $8 \leq k \leq h - 4$, and the obstruction forbids $e = e' = 1$. Use charges

$$c(0), \dots, c(5) = (0, 0, 1, 2, 3, 3), \quad c(p) + c(q) \leq 3 \quad \text{if } p + q \leq 5.$$

Define

$$\begin{aligned} P_0(k) &= (0, 0, 1, 2, 3, 3, 5, 6, 9)[k], & 0 \leq k \leq 8, \\ P_0(k) &= \min\{6k - 42, 3k + 3h - 51, 6h - 54\}, & k \geq 9, \\ P_1(k) &= 6k - 39, & 8 \leq k \leq h - 4. \end{aligned}$$

Both functions are nondecreasing in their valid count domains. For every one-cohort step with useful allocation $p \leq 5$,

$$P_e(k) \leq c(p) + P_{e'}(k'). \quad (125)$$

To specify its complete check, put $s = k - p$. At $s = 0$ the only output is $(0, 0)$. If $4 \leq s \leq h - 8$, the minimum marked output is $(s + 4, 1)$ and the minimum unmarked output is $(s + 5, 0)$; the former is forbidden when $e = 1$. Outside this interval the output is unmarked with count at least $g_h(s)$. Monotonicity handles every larger actual neighborhood. These cases include all geometric outputs.

For $h = 18, 27, 36, 45$, substitution checks respectively 183, 345, 507, and 669 inequalities, using integers only. The following three cases prove every larger parameter, so this finite verification is

not an extrapolation. Write $h = 45 + \Delta$, $\Delta \geq 0$. If $k \leq 17$, every minimum output is at most 22 and the base inequality is unchanged. If $k \geq 27 + \Delta$, subtract Δ from source and output counts: the residual is at least $22 + \Delta$, and all profiles, history endpoints, and potentials translate exactly, the latter by 6Δ . Finally, if $18 \leq k \leq 26 + \Delta$, the residual lies in $[13, 26 + \Delta] \subseteq [4, h - 8]$, and source and minimum output are in the affine region $P_e(k) = 6k - 42 + 3e$. The possible drops are $6p - 27$ for $0 \rightarrow 1$ or $1 \rightarrow 0$, and $6p - 30$ for $0 \rightarrow 0$; each is at most $c(p)$. This proves (125) for every $h \geq 45$. The four bases cover all remaining physical half-sizes $h = 9n/2$.

Both initial potentials are $P_0(h) = 6h - 54$, and both terminal potentials vanish. The combined potential decreases by at most three per day. Therefore

$$T_5 \geq \frac{2(6h - 54)}{3} = 4h - 36 = 18n - 36.$$

23.6.3 An explicit sweep attaining the bound

By Lemma 14.2, weight-then-reverse-lexicographic prefixes have prefix neighborhoods even when a box has an even side. We use this nesting only for construction. Let $U_p(k)$ be the cumulative cost through the bottom-slice ranks at most k , and $B_p(k)$ the number of such ranks. The stable bottom data are

p	ranks	costs
0	(1, 2, 3, 5, 8)	(2, 1, 1, 0, 0)
1	(1, 2, 4, 6)	(2, 1, 1, 0).

They stabilize already at $n = 3$, as in the odd-cylinder proof. Reflection now exchanges the colors. Counting the top slice gives the exact physical prefix profiles, zero at zero,

$$f_p(k) = k + U_p(k) - 4 + B_{1-p}(h - k), \quad k > 0.$$

Indeed the top slice has $4 + p$ vertices of source color p , its omitted vertices have the bottom ranks of color $1 - p$, and the origin correction is p ; the resulting constant is $p - (4 + p) = -4$.

Start one cohort in phase zero of this order and inspect its last five rooms each day. The full count follows

$$h \longrightarrow h - 2 \longrightarrow h - 3,$$

returning to phase zero. For every phase-zero count $10 \leq k \leq h - 3$, the next two days are $k \rightarrow k - 1 \rightarrow k - 1$. After $h - 12$ such pairs, the count is nine in phase zero. The final four days are

$$9 \longrightarrow 8 \longrightarrow 7 \longrightarrow 5 \longrightarrow 0.$$

All these transitions follow by direct substitution in f_p : U_0 is saturated by rank three and U_1 by rank four, while the opposite tails supply the displayed upper range. Thus the solo duration is $2 + 2(h - 12) + 4 = 2h - 18 = 9n - 18$ for every $h \geq 18$.

This duration is even. The untouched cohort is then in physical phase one. Reflect the order in the longitudinal coordinate for its sweep; reflection exchanges colors, so it starts in virtual phase zero and has the same duration. The first cohort is empty and remains empty. The two physical sweeps take $18n - 36$ days, proving the theorem. The independently reflected second order matters: the same orientation would take one additional day.

The accompanying integer certificate and an independent implementation check the finite boundary assertions and all base inequalities. Separate coordinate-neighborhood replays check the two complete sweeps at $n = 4, 6, 8, 10, 12, 30$. The displayed word induction, three-interval potential argument, and explicit sweep establish the unbounded theorem.

23.7 Six inspections on even $3 \times 3 \times n$ cylinders

Theorem 23.6. *For every even $n \geq 4$,*

$$T_6(P_3 \square P_3 \square P_n) = 6n - 8.$$

Proof. Put $n = 2L$ and $h = 9L$. We reuse the arbitrary-support geometry proved in Section 23.6; those geometric statements do not depend on the inspection budget. For a cohort of size k , let e record whether the preceding survivor had surplus four and size in $[4, h - 8]$. Initially $(k, e) = (h, 0)$. The valid domains are $0 \leq k \leq h$ for $e = 0$, and $8 \leq k \leq h - 4$ for $e = 1$. After p useful inspections, put $s = k - p$. An empty survivor gives $(0, 0)$. Otherwise the next size y satisfies $y \geq g_h(s)$, where g_h is (124). The next bit is one exactly when $4 \leq s \leq h - 8$ and $y = s + 4$. Consecutive bits equal to one are impossible. This relaxes every physical search, including searches whose supports are not global prefixes.

Assign charges $(c(0), \dots, c(6)) = (0, 0, 0, 1, 1, 2, 2)$. They satisfy $c(p) + c(q) \leq 2$ whenever $p + q \leq 6$. Define

$$\begin{aligned} (F_0(0), \dots, F_0(8)) &= (0, 0, 0, 1, 1, 2, 2, 3, 4), \\ F_0(k) &= \min \{ \lfloor 4k/3 \rfloor - 7, 4h/3 - 8 \} & (k \geq 9), \\ F_1(k) &= \lfloor (4k + 2)/3 \rfloor - 7 & (8 \leq k \leq h - 4). \end{aligned}$$

Both functions are nondecreasing. Every allowed transition satisfies

$$F_e(k) \leq c(p) + F_{e'}(y). \tag{126}$$

Here is a finite verification with an explicit extension to all lengths. For $h = 18, 27$, enumerate valid (k, e) , $0 \leq p \leq \min(k, 6)$, and every $g_h(k - p) \leq y \leq h$. Set $e' = 1$ precisely in the critical case above, reject $e = e' = 1$, and check (126); when $p = k$, check just $(y, e') = (0, 0)$. These are respectively 1,095 and 3,174 integer inequalities, all satisfied. This specification and the formulas fully determine the finite certificate; the companion artifact includes an exact enumerator.

For $h = 27 + \Delta$, where $\Delta \in 9\mathbb{N}$, it suffices by monotonicity to check the least output in each next-bit class. Split the source range:

- If $k \leq 14$, the profile and critical status are unchanged from $h = 27$, and every minimum output is at most 19. The potentials are also unchanged, so the base inequalities apply.
- If $k \geq 20 + \Delta$, subtract Δ from source, survivor, and output. Since $s - \Delta \geq 14$, the lower profile and critical upper endpoint depend only on $h - s$; the full case $s = h$ remains full. The transformed transition is therefore an allowed $h = 27$ transition. Both potentials decrease by $4\Delta/3$, including the full-state cap.
- If $15 \leq k \leq 19 + \Delta$, then $9 \leq s \leq h - 8$. The minimum branches are $0 \rightarrow 1$ with $y = s + 4$, and $0 \rightarrow 0$ or $1 \rightarrow 0$ with $y = s + 5$. The uncapped formula $\lfloor (4k + 2e)/3 \rfloor - 7$ applies throughout. The maximum drops over $k \bmod 3$, for $p = 0, \dots, 6$, are

	0	1	2	3	4	5	6
$0 \rightarrow 1, 1 \rightarrow 0$	-6	-4	-3	-2	0	1	2
$0 \rightarrow 0$	-6	-5	-4	-2	-1	0	2

Every entry is at most $c(p)$.

This proves (126) for every even length. The two initial potentials sum to $2(4h/3 - 8)$ and both vanish at capture. Their combined daily drop is at most two, giving $T_6 \geq 4h/3 - 8 = 6n - 8$.

For the upper bound, the first two moves depart slightly from a global prefix sweep. Start with the phase-one cohort, whose full slice counts are $(4, 5)^L$. Retain the transverse slice prefixes with counts

$$R_1 = ((4, 5)^{L-1}, 2, 1), \quad R_2 = ((5, 4)^{L-1}, 1, 0)$$

on the first two days. Their successive neighborhoods, calculated from the transverse profiles, are

$$((5, 4)^{L-1}, 5, 2), \quad A = ((4, 5)^{L-1}, 4, 1).$$

Each move inspects six rooms, and $|A| = h - 4$. The global weightlex prefix of size $h - 10$ is contained in A . Indeed, its only missing rooms are four final-slice cells of coordinate weight at least $n + 1$. Reflection in all three coordinates maps rooms of weight at least $n + 1$ to the seven phase-zero rooms of weight at most two. Thus all four missing cells lie among the seven largest-weight rooms and are omitted by this prefix.

Now repeatedly retain the global prefix of size $\max(k - 6, 0)$. The compatibility and exact prefix profiles established in the preceding subsection justify every later move. Starting in phase one, each pair with $k \in \{17, 20, \dots, h - 4\}$ gives $k \rightarrow k - 1 \rightarrow k - 3$. After $(h - 18)/3$ pairs the count is 14, and the final six days are

$$14 \longrightarrow 13 \longrightarrow 11 \longrightarrow 10 \longrightarrow 8 \longrightarrow 6 \longrightarrow 0.$$

One cohort therefore takes $2 + 2(h - 18)/3 + 6 = 3n - 4$ days, an even number. The other cohort stays full during these moves. Reflecting the long coordinate converts its phase zero into the virtual starting phase one, so the same construction clears it in a further $3n - 4$ days. $\square$

The finite arithmetic certificate and its all-length extension are ordinary proofs. This additional even-cylinder result is not presently a complete physical Lean theorem.

23.8 A uniform height strategy for cylinders

Theorem 23.7. *If H is any finite bipartite graph with $A \geq 1$ vertices and $n \geq 2$, then*

$$h(H \square P_n) \leq \lfloor A/2 \rfloor + 1.$$

For $m > A/2$, at most $2 \lceil A(n + 2)/(2m - A) \rceil$ days suffice.

Proof. We first clear one initial-color cohort. Write the path coordinate as $0 \leq z < n$. Bound the possible positions in each fiber by a cutoff a_v , of the current required parity, such that $|a_u - a_v| = 1$ whenever uv is an edge of H . Initially take $a_v = n - 1$ or $n - 2$ according to its parity; this bounds the full cohort. Cutoffs may later be negative or exceed the board.

At a local maximum v , lower a_v by two. Record the room at the old height if it lies on the board. Every adjacent difference stays one. A sequence of m such virtual operations records at most m distinct rooms: repeated operations in a fiber use strictly decreasing heights. Inspect those rooms simultaneously. The survivors lie below the new cutoffs a'_v . After movement their cutoffs are bounded by $a'_v + 1$: a path move increases height by at most one, and an H -move arrives from u with $a'_u \leq a'_v + 1$.

Choose the virtual operations in a fixed cyclic order. List first the initially higher color class of H , then the other class, and repeat this A -letter word. Every operation is at a local maximum;

after a whole color class has been lowered, the other is higher. Adding one to all cutoffs after a day changes no comparison. Isolated vertices of H cause no difficulty.

After t days every vertex has been lowered at least $\lfloor mt/A \rfloor$ times, so

$$\max_v a_v(t) \leq n - 1 + t - 2 \lfloor mt/A \rfloor < n + 1 - t(2m - A)/A.$$

The displayed number of days per cohort makes every cutoff negative. Since $H \square P_n$ has no isolated vertices, an empty post-movement belief certifies capture during inspection. Repeat for the other initial cohort, starting with the full bound of its then-current parity. $\square$

Theorem 23.8 (Exact feasibility on sufficiently long boxes). *Let H be a finite bipartite graph with $A \geq 2$ vertices and a Hamiltonian path. If $n \geq 2 \lfloor A/2 \rfloor$, then*

$$h(H \square P_n) = \lfloor A/2 \rfloor + 1.$$

In particular, this applies to every Cartesian box used as the transverse graph, with arbitrary even or odd side lengths.

Proof. The Hamiltonian path supplies a spanning subgraph P_A of H , so $P_A \square P_n$ is a subgraph of $H \square P_n$. An evader can restrict its moves to this subgraph; the subgraph need not be induced. The classical rectangle theorem [1, Theorem 2] therefore gives

$$h(H \square P_n) \geq \lfloor \min(A, n)/2 \rfloor + 1 = \lfloor A/2 \rfloor + 1.$$

Theorem 23.7 gives the matching upper bound. A Cartesian box has a Hamiltonian path by the usual snake construction: traverse successive slices alternately forwards and backwards along a Hamiltonian path in the lower-dimensional box. The joining endpoints differ only in the new coordinate. $\square$

Thus, for example, $4 \times 4 \times n$ requires exactly nine probes for every $n \geq 16$, and $3 \times 4 \times n$ requires exactly seven for every $n \geq 12$. The length threshold is sufficient; no claim is made that it is the first length attaining the eventual budget. This theorem settles feasibility in an unbounded family containing all side parities, without asserting an optimal-time formula.

Corollary 23.9. *For every $n \geq 3$, $h(P_3 \square P_3 \square P_n) = 5$. The remaining values are two at $n = 1$ and four at $n = 2$.*

Proof. For $n \geq 3$ the board contains the three-cube. An evader may choose to stay in that subgraph, so its hunting number is at least five. Theorem 23.7 gives the upper bound with $A = 9$. At $n = 1$ use the rectangle theorem; at $n = 2$ use Theorem 23.1 with the 3×3 base. $\square$

For a general box this construction can be applied along a longest axis, giving $\lfloor (\prod_{i \neq j} n_i)/2 \rfloor + 1$ as a sufficient budget. It can be loose. The five-inspection time of every $3 \times 3 \times n$ cylinder is determined above; the full classification at larger budgets remains open here.

23.9 The optimal growth rate for every transverse box

The exact offsets remain sensitive to geometry, but the leading term of the optimal time has a uniform answer across all side parities.

Theorem 23.10. *Let H be a finite bipartite graph with $A \geq 2$ vertices and a Hamiltonian path. For each fixed integer $m > A/2$, as $n \rightarrow \infty$ through either parity,*

$$T_m(H \square P_n) = \frac{2An}{2m - A} + O_{A,m}(1). \quad (127)$$

The height strategy in Theorem 23.7 has a bounded additive excess over the optimal time. When A is even, the following bounds hold explicitly for every $n \geq 2$:

$$\frac{2(An/2 - A^2 - m)_+}{m - A/2} \leq T_m(H \square P_n) \leq 2 \left\lceil \frac{A(n+2)}{2m - A} \right\rceil. \quad (128)$$

In particular the theorem applies to every Cartesian transverse box.

Lemma 23.11 (An interior expansion bound for even-width rectangles). *On $P_{2b} \square P_n$, $b \geq 1, n \geq 2$, put $h = bn$ and $K = 2b^2$. Every one-color set R with $K < |R| < h - K$ satisfies $|N(R)| \geq |R| + b$.*

Proof. Match consecutive short-axis rows in pairs. Contract an edge followed by the inverse matching to obtain a directed $b \times n$ grid D . It has loops, bidirectional horizontal edges, and vertical rungs directed one way in even columns and the other way in odd columns. The opposite physical color reverses all rung directions. In either case

$$|N(R)| - |R| = |N_D^+(R) \setminus R|.$$

Suppose the outside boundary $B = N_D^+(R) \setminus R$ has size $s < b$. In $D - B$, at least one entire horizontal row is undeleted. Its vertices lie in a single strongly connected component C .

Call a column bad if it contains a vertex of B . Every adjacent pair of good columns is strongly connected: the horizontal edges go both ways, and the opposite rung directions let a walk go both ways between adjacent rows. The pair meets the undeleted row, hence belongs to C . Only bad columns and isolated good columns can contain vertices outside C . There are at most s of the former and $s + 1$ of the latter, so

$$|V(D) \setminus C| \leq b(2s + 1) < 2b^2 = K.$$

This count includes deleted vertices and also covers the case with no adjacent good columns. The set R is outgoing-closed in $D - B$. It therefore either avoids C or contains C entirely. In the first case $|R| \leq K$; in the second $h - |R| \leq K$, proving the contrapositive. $\square$

Proof of Theorem 23.10. First let $A = 2b$ and consider the spanning rectangle $P_A \square P_n$. Put $C = (h - 2K - m)_+$, $L = K + m$, and

$$\psi(k) = \min\{C, (k - L)_+\}.$$

If a cohort of size k receives $p \leq m$ useful inspections, write $r = k - p$ and $k' = |N(R)|$ for its survivor size and actual next size. The matching gives $k' \geq r$ for every survivor set. If $r \leq K$, the source potential is zero. If $r \geq h - K$, the target potential is C whenever $C > 0$, while $C = 0$ is trivial. In the remaining case Lemma 23.11 gives $k' \geq k - p + b$. Since clipping is monotone and has Lipschitz constant one, all cases give

$$\psi(k) - \psi(k') \leq \max\{0, p - b\}.$$

The two cohorts' useful allocations sum to at most m , so their combined potential decreases by at most $m - b$. Their initial sum is $2C$ and their final sum zero. Thus $T_m \geq 2C/(m - b)$ on the rectangle. The rectangle is a subgraph of $H \square P_n$, giving the same lower bound there. Theorem 23.7 gives the displayed upper bound, with the same coefficient of n .

For odd $A \geq 3$, take the largest odd $\ell \leq n$. The spanning rectangle $P_A \square P_\ell$ has, by the fixed-budget odd-rectangle theorem, an eventual affine period with length increment $2(2m - A)/\gcd(A, 2m - A)$ and time increment $4A/\gcd(A, 2m - A)$. Its time is consequently $2A\ell/(2m - A) + O_{A,m}(1)$. Since $n - \ell \leq 1$, this gives the required lower bound on the whole cylinder. The same height upper bound completes the proof. $\square$

For example, at budgets nine and ten respectively, $T_9(P_4^2 \square P_n) = 16n + O(1)$ and $T_{10}(P_4^2 \square P_n) = 8n + O(1)$. At seven probes, $T_7(P_3 \square P_4 \square P_n) = 12n + O(1)$. Here the cross-section and budget are fixed while the longitudinal side grows. The next theorem strengthens this leading-term result to an exact eventual affine period for every fixed transverse box. Exact constant terms for all short boxes remain a separate question.

24 Eventual affine periods for every fixed transverse box

The optimal growth rate in Theorem 23.10 can be strengthened to an exact eventual recurrence for every fixed transverse box, without a parity restriction on its remaining side.

Theorem 24.1 (Every fixed transverse box). *Let Q be a fixed Cartesian product of finite paths, with $A = |Q| \geq 1$. For each fixed integer $m > A/2$, put*

$$D = 2m - A, \quad g = \gcd(A, D), \quad \Delta = 2D/g.$$

There is an effectively specified positive integer N , bounded by a polynomial in A and m , such that

$$T_m(Q \square P_{n+\Delta}) = T_m(Q \square P_n) + 4A/g, \quad n \geq N. \quad (129)$$

Both longitudinal parities are covered, and Δ is an even integer. For each integer $0 \leq m \leq A/2$, all sufficiently long cylinders are impossible to search with that budget.

We prove the time statement first for even-order Hamiltonian cross-sections, then for odd-order Hamiltonian cross-sections at even longitudinal lengths. The earlier odd-box theorem closes the remaining case for boxes. The argument applies to arbitrary strategies: optimality forces all but a bounded number of days to move a narrow physical boundary. A common interval untouched by exceptional behavior lets us shorten or lengthen the two boundary passages through it.

Theorem 24.2. *Let Q be a finite bipartite graph with a Hamiltonian path and $2b$ vertices, where $b \geq 1$. Fix an integer $m > b$, and put $d = m - b$. There is an effective positive integer N such that, with the even period $\Delta = 2d/\gcd(b, d)$,*

$$T_m(Q \square P_{n+\Delta}) = T_m(Q \square P_n) + \frac{2b\Delta}{d}, \quad n \geq N. \quad (130)$$

The assertion holds for both parities of n . Thus the optimal time is eventually affine on every residue class modulo Δ . The displayed period is valid but need not be the smallest; the effective threshold below is deliberately conservative.

Fix a Hamiltonian order on Q and match consecutive vertices in that order. Each checkerboard class in the cylinder identifies with a $b \times n$ array of matched pairs. Write $h = bn$. As in Lemma 23.11, following an edge by the inverse matching gives a directed graph containing loops, bidirectional edges along every longitudinal row, and alternating vertical rungs between consecutive rows. Extra edges of Q add directed edges inside columns. For a one-color support R , its ordinary neighborhood, expressed using the opposite-color matching, is $R \cup B$, where B is the directed outside boundary. In particular $|N(R)| - |R| = |B|$.

24.1 Equality in the interior expansion bound

Lemma 24.3 (A finite family of minimum-surplus frontiers). *Put $K = 4b^2$. Suppose a one-color set R in $Q \square P_n$ satisfies*

$$K < |R| < h - K, \quad |N(R)| = |R| + b.$$

In the matched array, its rows are either all prefixes or all suffixes: there are integers $1 \leq c_i \leq n-2$, $0 \leq i < b$, such that either

$$R_i = \{0, \dots, c_i - 1\} \quad \text{for every } i, \quad \text{or} \quad R_i = \{c_i + 1, \dots, n - 1\} \quad \text{for every } i.$$

Adjacent cuts satisfy $|c_{i+1} - c_i| \leq 2$. If their difference is two, the rung at the sole intermediate column points from outside R into R . The outside boundary consists exactly of the b cut points. Additional edges of Q may exclude some such patterns, but introduce no additional forms.

Proof. First use only the spanning Hamiltonian rectangle. Its boundary has size at most b , and Lemma 23.11 shows that its size is at least b . Hence its boundary B has exactly b points and also equals the boundary for the whole cylinder.

If one row misses B , the strongly connected component argument in Lemma 23.11 still applies with b deleted points. One component contains all but at most $b(2b+1) \leq K$ vertices, including the deleted points in the exceptional count. The outgoing-closed set R either contains or avoids that component, contrary to the two strict size bounds. Thus every row contains exactly one boundary point, say c_i . Horizontal bidirectionality shows that a row of R is empty, a prefix, a suffix, or the entire row except its boundary point.

For two adjacent rows, the difference of their row sets is a union of at most two integer intervals. At a column where a rung points from the first row to the second, such a difference can contain only the second row's boundary point. It therefore has at most one point of that column parity. A union of at most two intervals with this property has at most four points. Applying this in both directions gives

$$|R_i \setminus R_{i+1}| \leq 4, \quad |R_{i+1} \setminus R_i| \leq 4.$$

An empty row would now imply $|R| \leq 4b^2$, by propagating the bound through all rows. A row full except at its boundary similarly implies $h - |R| \leq b(1 + 4(b-1)) \leq 4b^2$. Both are excluded.

Suppose an adjacent prefix and suffix have lengths u and v . Their two differences have sizes $\min(u, n-v)$ and $\min(v, n-u)$, both at most four. The hypotheses give $n > 8$. If either length exceeds four, both lengths are at least $n-4$; otherwise both are at most four. Propagation through all rows again makes either R or its complement have size at most K . Hence all rows have the same orientation, and their cuts are interior.

Every column strictly between two adjacent cuts has one rung endpoint inside R and its other endpoint outside both R and B . Its rung must therefore point into R . Alternating directions

permit at most one such column, proving the cut bound and parity condition. Conversely, for the spanning rectangle these conditions give exactly the indicated boundary: horizontal edges reach every cut, and no rung reaches another outside point. Extra transverse edges impose only further restrictions on these same bounded-width patterns. $\square$

The cuts have total range at most $2(b-1)$; for subsequent spatial margins use the larger value $v = 2b + 2$. For a prefix survivor with cuts (c_i) , the next matched belief consists of prefixes of lengths $c_i + 1$. A following prefix survivor with cuts (e_i) is possible precisely when $e_i \leq c_i + 1$, and its useful inspection count is

$$p = \sum_{i=0}^{b-1} (c_i + 1 - e_i). \quad (131)$$

At $p = m$ its sum of cuts decreases by d . The suffix case is reflected and its sum increases by d . An orientation cannot change between two such bulk survivors: the next belief of a prefix survivor omits the last column in every row, whereas every nonempty suffix contains it.

24.2 Only boundedly many days can behave differently

Consider an optimal strategy on a sufficiently long cylinder, and set

$$C = h - 2K - m > d, \quad \psi(k) = \min\{C, (k - K - m)_+\}.$$

The argument proving Theorem 23.10, with the larger cutoff $K = 4b^2$, gives for a cohort receiving p useful inspections and changing from count k to k' ,

$$\psi(k) - \psi(k') \leq \max\{0, p - b\}. \quad (132)$$

The two cohorts start with total potential $2C$ and end with zero. Their combined daily decrease is at most d . The height upper bound $T_m \leq 2 \lceil b(n+2)/d \rceil$ consequently gives

$$dT_m - 2C \leq B, \quad B = 4K + 4m + 2b - 2. \quad (133)$$

Indeed the integer ceiling bound gives $dT_m \leq 2b(n+2) + 2d - 2$.

Call a day *efficient* when its total potential loss is exactly d . Every other day contributes at least one to the nonnegative integer slack in (133), so at most B days are inefficient. Potential increases contribute additional slack and cause no exception. Equality forces all m useful inspections to be assigned to one cohort, called the day's *owner*: any nontrivial split makes the sum of the two positive-part bounds in (132) strictly less than d .

On an efficient day the owner loses potential d , while its untouched mate has constant potential. The owner's source potential is positive, so its survivor count r exceeds K . Its target potential is below C , so its next count, and therefore r , is below $h - K$. The interior expansion bound gives raw count loss at most $m - b = d$. Clipping cannot increase a positive loss; equality forces exactly m useful inspections, raw count loss d , and survivor surplus exactly b . Thus every efficient owner survivor has the form in Lemma 24.3.

An untouched cohort with potential strictly between zero and C cannot have constant potential: its actual count lies in the interior range and its neighborhood grows. The mate of an efficient owner therefore has potential zero or C . A consecutive run of efficient days has at most two owner blocks. Once a cohort has owned a day its potential is below C ; if ownership switches, that cohort can

remain unchanged only at zero, from which it cannot subsequently lose another d in the run. There are consequently at most $2(B + 1)$ owner blocks.

The matching-contracted graph is strongly connected for $n \geq 2$: neighboring pairs of columns allow movement in both rung directions, and rows are bidirectional. Every proper nonempty untouched support therefore grows by at least one room per day. Put $L = K + m + 1$. An untouched nonempty mate with zero potential cannot stay in that collar for L days, while an untouched mate with potential C has deficit at most K and becomes full within K days. Mark every inefficient day and the first L days of every owner block as *nonclean*. Their total number is at most

$$E = B + 2L(B + 1). \quad (134)$$

On each remaining *clean* day the mate is globally full or empty. A maximal clean block has one owner, a fixed frontier orientation, and both its owner beliefs and survivors have bounded-width descriptions: its first belief is already the neighborhood of the preceding efficient survivor. By (131), its average cut moves at the constant speed d/b toward clearance. Individual cuts need not move monotonically.

24.3 A physical interval protected from all exceptional behavior

At the start of every consecutive block of nonclean days, mark the preceding clean frontier's b cut columns, if there is one. Initially the board is full, so no cuts need be marked. Also mark every actual inspection column on every nonclean day, and the two board ends. There are at most

$$b(E + 1) + mE + 2 \leq F, \quad F = (b + m)(E + 1) + 2$$

marked columns. Enlarge each mark by radius $\rho = E + 2b + m + 2$. Outside these intervals, each cohort is locally full or empty throughout every nonclean block. At its start, all cuts are far away. During its at most E days there are no nearby inspections, and movement propagates information by at most one longitudinal column per day. An empty region thus stays empty in the protected interior. A full region stays full using the matching inside each column. This proves the assertion for all intermediate states as well as the endpoints, even if remote shots create arbitrary holes elsewhere.

For any prescribed W , the inequality

$$n > F(2\rho + 1) + (F + 1)(W + 2v + 4) \quad (135)$$

provides an unmarked interval of length at least $W + 2v + 4$. Trim $v + 2$ columns from each end and call the retained interval J ; it has length at least W .

A clean block has its endpoint frontiers outside the untrimmed interval: the preceding nonclean block ends uniformly there, and the cuts preceding the following nonclean block were marked. Its average frontier moves strictly toward clearance, with cut range at most v . Consequently it can change its owner's status on J only from full to empty, through a single complete passage. If its endpoint statuses agree, trimming removes any partial excursion onto J . It cannot start empty and end full, since that would move the average across the interval in the wrong direction. Nonclean blocks preserve the local status. Since each cohort starts full on J and ends empty there, each has exactly one clean crossing. These crossings occur at disjoint times; during either one the other cohort is globally full or empty.

24.4 Connecting physical cutoff frontiers

A direct height descent connects any compatible endpoints once the duration exceeds a short bound. This supplies the replacement segments needed for the exact period, without enumerating frontier states.

Put $A = 2b$ and $D = 2m - A = 2d$. For a left frontier, let a_v be the last occupied physical longitudinal coordinate in the fiber over $v \in Q$. The two cutoffs in transverse matched pair i are $c_i - 1$ and $c_i - 2$, in the order required by parity. Its actual next neighborhood advances both by one. A transverse edge uv consequently forces $a_u \leq a_v + 1$, and reversing the edge gives the opposite inequality. Their parities differ, so

$$|a_u - a_v| = 1 \quad \text{on every edge } uv \text{ of } Q. \quad (136)$$

This includes every extra transverse edge. Inside the protected crossing, all fiber cutoffs are far from both longitudinal ends. Conversely, any such interior prefix configuration satisfying (136) has neighborhood cutoffs exactly $a_v + 1$: longitudinal movement supplies them and transverse edges supply nothing larger. Right frontiers use reflected coordinates.

Lemma 24.4 (Height endpoints). *Let Q be connected and bipartite, with $A \geq 2$ vertices and diameter R . Fix $m > A/2$ and put $D = 2m - A$. Let a, b be integer height configurations satisfying (136). If*

$$t \geq \lceil AR/m \rceil, \quad b_v \equiv a_v + t \pmod{2} \quad \text{for every } v, \quad \sum_v a_v - \sum_v b_v = Dt,$$

then t global additions of one, each followed by m legal local maximum flips, take a to b . At completed days the mean height moves steadily between the endpoint means. All intermediate cutoffs, including partial inspection steps, lie in

$$[\mu(b) - R, \mu(a) + R + 1], \quad \mu(x) = A^{-1} \sum_v x_v.$$

Whenever this interval lies inside the cylinder's interior, these operations are actual movements and m useful inspections per day.

Proof. This is a corollary of the exact quota-word criterion in Lemma 12.2. Put $c = b - t\mathbf{1}$. Its mean difference from a is $-2mt/A$. Edgewise height differences bound the range of $c - a$ by $2R$, so for every vertex

$$c_v - a_v \leq -2mt/A + 2R \leq 0.$$

Thus $b \leq a + t\mathbf{1}$ and the number of required lowerings in that lemma is exactly mt . Its movement-first construction gives m useful lowerings after each of the t additions.

Every completed day lowers the mean by D/A . A configuration has range at most R , so its completed-day cutoffs lie between $\mu(b) - R$ and $\mu(a) + R$. Movement can raise the upper bound by one; subsequent lowerings descend coordinatewise to the next survivor. This proves the displayed movement-first margin, including partial inspection steps. Within a physical interior, all lowerings remove actual distinct rooms and the global additions are exact neighborhoods. $\square$

Set $g = \gcd(A, D)$, $\ell_0 = 2A/g$, and $\Delta = 2D/g$. These are even integers satisfying $A\Delta = D\ell_0$. Suppose a clean survivor path from a to b has length $t \geq \lceil AR/m \rceil + \ell_0$. After deleting Δ columns, a left crossing requires endpoints $a - \Delta\mathbf{1}, b$. Their sum difference is $D(t - \ell_0)$ and their phase difference

is $t - \ell_0$, since both shifts are even. Lemma 24.4 supplies a path of exactly that shorter length. For insertion use endpoints $a + \Delta \mathbf{1}, b$ and duration $t + \ell_0$. Reflection handles right frontiers. The mean and range bound controls each replacement inside the protected physical band, independently of the old path's details.

For the present even-order case,

$$\ell_0 = \frac{2b}{\gcd(b, d)}, \quad \Delta = \frac{2d}{\gcd(b, d)}, \quad K_0 = \left\lceil \frac{2b(2b-1)}{m} \right\rceil. \quad (137)$$

Here are conservative explicit margins that guarantee such a path well inside the protected interval. Set

$$U = K_0 + \ell_0 + 4, \\ W = 2(\Delta + v + m + 3) + \left\lceil \frac{d(U+4)}{b} \right\rceil + 2v + 2m + 10. \quad (138)$$

In a complete crossing, take the central segment after its average cut has entered J with margin $\Delta + v + m + 3$, and before it leaves at the opposite margin. Its average advances by $d/b \leq m$ per edge, so entry and exit overshoots lose at most $2m$ of spatial length. The remaining travel exceeds dU/b , providing more than U edges. Every cut is within v of the average, so the entire segment stays more than Δ from the ends of J . Its endpoint frontiers lie on opposite sides of a central interval of Δ columns. These margins also cover the entry and exit edges.

For definiteness, one may take

$$N = 4 + \left\lceil \frac{2K + m + d + 1}{b} \right\rceil + F(2\rho + 1) + (F + 1)(W + 2v + 4) + \Delta. \quad (139)$$

This ensures $C > d$, the protected-gap condition, and the same conditions after contracting by Δ . All constants depend only on b and m .

24.5 Deleting and inserting the common interval

Start with an optimal strategy on a cylinder long enough for the preceding construction. Delete Δ consecutive columns from the center of J . Outside the two central crossing segments, map rooms and inspections by ordinary column deletion. No nonclean inspection is lost. Every nonclean state is uniformly full or empty near the join, so neighborhood formation commutes with the deletion there. Other clean frontiers are entirely on one side of the join, and their transitions are either left fixed or translated as a whole by the even displacement Δ .

Within each central crossing, replace its t -edge path by the $(t - \ell_0)$ -edge path in Lemma 24.4. For a prefix frontier the new initial cutoffs are Δ lower and the final cutoffs are the old ones, matching the natural column map at both endpoints. For a suffix frontier use the reflected construction. The mean of the replacement decreases steadily between the mapped endpoint means, and its range is at most $2b - 1$. The chosen margins therefore keep every intermediate cutoff inside the shorter board. Every new daily edge is an actual movement followed by at most m inspections.

The mate is globally full or empty during this operation. Deleting an even number of days and columns preserves its phase and state. The two crossings are disjoint in time, so their edits do not interfere. An edge is one subsequent inspection day, hence exactly ℓ_0 days have been removed in each crossing. All remaining inspections respect the original budget. We have proved

$$T_m(Q \square P_{n-\Delta}) \leq T_m(Q \square P_n) - 2\ell_0. \quad (140)$$

For the reverse inequality, prepare an optimal strategy on the shorter, still sufficiently long cylinder and insert Δ columns in its protected interval. Extend the locally full or empty regions across the inserted band and translate the far-side rooms. Within each crossing, use Lemma 24.4 to replace its t days by $t + \ell_0$ days. A prefix frontier starts shifted right by Δ and ends at the old final cutoff; reflect this for a suffix. The same local transition and margin arguments apply. The full or empty mate returns to its phase, and the total added time is $2\ell_0$. Therefore

$$T_m(Q \square P_n) \leq T_m(Q \square P_{n-\Delta}) + 2\ell_0. \quad (141)$$

Combining (140) and (141), and using $2\ell_0 = 2b\Delta/d$, proves Theorem 24.2 for every $n \geq N$ after renaming the shorter length n . The even displacement preserves either longitudinal parity throughout.

24.6 Odd-order cross-sections: geometry in paired columns

When the cross-section has odd order, there is no transverse perfect matching. At even longitudinal lengths we instead match consecutive columns. This changes the smallest bulk expansion from a constant cost to two alternating costs. A single bit recording the previous expansion will recover constant progress.

Theorem 24.5. *Let Q be a finite bipartite graph with a Hamiltonian path and $A = 2b + 1 \geq 3$ vertices. Fix an integer $m \geq b + 1$ and put $D = 2m - A$. Put $g = \gcd(A, D)$ and $\Delta = 2D/g$. There is an effective positive integer N such that*

$$T_m(Q \square P_{n+\Delta}) = T_m(Q \square P_n) + 4A/g$$

for every even $n \geq N$.

Write $n = 2L$ and let X, Y be the bipartition of Q , with $|X| = b + 1$ and $|Y| = b$. A Hamiltonian path ensures these color sizes. Match columns $2j$ and $2j + 1$. Each cohort now identifies with an $A \times L$ array of vertices (v, j) , one for each $v \in V(Q)$ and paired-column index j . The contracted graph has loops, bidirectional Q -edges within every layer, and opposite one-way longitudinal flows on the two Q colors. Choose phase zero so that X fibers flow left and Y fibers flow right; phase one reverses both flows. There is no dependence on j . The ordinary neighborhood surplus is again the cardinality of the directed outside boundary. Throughout the paired-column argument put $h = AL$ and $K = A^2$.

Lemma 24.6 (Minimum surplus in paired columns). *Every survivor R with $K < |R| < h - K$ has surplus at least b . If its surplus equals b , then in phase zero all fibers are left prefixes. Across each Q -edge xy , $x \in X, y \in Y$, the X prefix has the same length as the Y prefix or is one position longer. In phase one the reflected statement holds, with right suffixes. Every fiber has at least two occupied and two unoccupied positions. Two consecutive middle survivors in an actual strategy cannot both have surplus b .*

Proof. Let B be the outside boundary, of size $s \leq b$. If one X fiber and one Y fiber both avoid B , all completely undeleted Q layers lie in one strongly connected component of the graph with B deleted. Each layer is strongly connected, and the two untouched, oppositely directed fibers connect these layers in both longitudinal directions. At most s layers contain a boundary point, so this component contains all but at most $As < K$ vertices, counting the deleted layers among the

exceptions. The outgoing-closed set R either contains or avoids the component, contradicting its middle size. There is at least one undeleted layer: $h > 2A^2$ implies $L > 2A > b$.

If $s < b$, both colors have an untouched fiber, a contradiction. If $s = b$, there is still an untouched X fiber, so every Y fiber must contain a boundary point. Thus B consists of exactly one point (y, c_y) on each Y fiber, and no points on X fibers.

In phase zero, the boundary-free X fibers flow left and are prefixes. Bidirectional Q -edges give, for every edge xy ,

$$R_y \subseteq R_x, \quad R_x \setminus R_y \subseteq \{c_y\}.$$

Adjacent cardinalities differ by at most one. Connectivity bounds the global range of fiber sizes by $A - 1$. Both the average size and average complement size exceed A , so each fiber has at least two occupied and two unoccupied positions. If $R_x = \{0, \dots, k - 1\}$, the inclusions make R_y either that prefix or that prefix with its boundary point removed. Removing a point strictly below $k - 1$ would leave $k - 1$ occupied; its right-going edge would then make k a second boundary point. Thus

$$R_y = \{0, \dots, c_y - 1\}, \quad k \in \{c_y, c_y + 1\}.$$

This proves the prefix description. Reflection proves the phase-one statement.

The next belief of a phase-zero minimum-surplus frontier leaves its X lengths unchanged and increases each Y length by one. It still omits position $L - 1$ in every fiber. A phase-one minimum-surplus middle survivor would be a nonempty right suffix in every fiber, hence could not be contained in that belief. Reflection handles the other direction. $\square$

Lemma 24.7 (The intervening frontier). *Suppose R is a phase-zero middle survivor of surplus b , and a phase-one middle survivor $S \subseteq N(R)$ has surplus $b + 1$. Then all fibers of S are left prefixes. Across each Q -edge, its Y prefix has the same length as its X prefix or is one position longer. The reflected assertion holds with both orientations reversed.*

Proof. For any support with boundary B , a Q -edge gives $S_u \setminus S_v \subseteq B_v$. Following a simple Q path shows that any two fiber cardinalities differ by at most $|B|$: its target fibers are distinct, so their boundary counts sum to at most $|B|$. At $|B| = b + 1$, an empty fiber would give $|S| \leq A(b + 1) < A^2$, contrary to the middle-size hypothesis. Every fiber is nonempty.

The containment $S \subseteq N(R)$ ensures that all fibers omit $L - 1$. In phase one, each X fiber flows right. If such a fiber had no boundary point, its nonempty set would be right-closed and would contain $L - 1$. Hence each of the $b + 1$ X fibers meets B . They exhaust the boundary, leaving no boundary point on any Y fiber. The latter flow left and are prefixes. Applying the same bidirectional-edge and endpoint argument as in Lemma 24.6, with X, Y exchanged, proves that the X fibers are prefixes as well, with the stated length relation. Reflection gives the other orientation. $\square$

Both frontier families have finitely many relative-length patterns. Their length range is at most $A - 1$, and choosing the zero-or-one length differences along a spanning tree gives at most 2^{A-1} patterns for fixed phase and orientation. Extra edges impose consistency conditions on these patterns.

24.7 A history bit and the paired-cylinder period

For each cohort, define $e_t = 1$ precisely when its preceding day's survivor was in the strict middle range and had surplus b ; otherwise set $e_t = 0$. Initially $e_0 = 0$. If its before-inspection count is

k_t , use the raw rank $2k_t + e_t$. Let p be its useful inspection count, so the current survivor has size $r = k_t - p$, surplus s , and next count $k_{t+1} = r + s$. The raw loss is

$$(2k_t + e_t) - (2k_{t+1} + e_{t+1}) = 2p - 2s + e_t - e_{t+1}. \quad (142)$$

For a middle survivor with $s = b$, Lemma 24.6 gives $(e_t, e_{t+1}) = (0, 1)$, so the loss is $2p - A$. If $s \geq b + 1$, the new bit is zero and the old bit at most one, again giving loss at most $2p - A$.

Set

$$L_0 = 2(K + m) + 1, \quad C = 2h - 4K - 2m - 1 > D, \quad \psi_t = \min\{C, (2k_t + e_t - L_0)_+\}.$$

A survivor of size at most K gives source rank at most L_0 , hence source potential zero. A survivor of size at least $h - K$ gives next rank at least $2h - 2K = L_0 + C$, hence target potential C , by the longitudinal matching. Clipping and (142) therefore give, in every case,

$$\psi_t - \psi_{t+1} \leq \max\{0, 2p - A\}. \quad (143)$$

For the two useful allocations $p + q \leq m$, the sum of these charges is at most D , with equality only if all m useful inspections belong to one cohort.

Both initial potentials equal C . The height upper bound $T \leq 2 \lceil A(n + 2)/D \rceil$ and its integer ceiling estimate imply

$$DT - 2C \leq B, \quad B = 8K + 8m + 2A.$$

Thus at most B days have total potential loss below D . On every efficient day the owner has positive source potential and target potential below C , which force its survivor into the strict middle range. Equality in the clipped loss then forces equality in (142). The only cases are

$$(s, e_t, e_{t+1}) = (b, 0, 1) \quad \text{or} \quad (b + 1, 1, 0). \quad (144)$$

The first survivor is a frontier by Lemma 24.6. In the second case the preceding survivor was a middle minimum-surplus frontier, so the actual containment in its neighborhood permits Lemma 24.7. Both efficient phases therefore have physical frontiers.

The paired contraction is strongly connected for $L \geq 2$. An untouched proper nonempty cohort grows by at least one room each day, while its history bit can drop by at most one. Its raw rank therefore increases by at least one, so its potential cannot remain constant strictly between zero and C . The owner-block argument in Section 24.2 applies unchanged. There are at most $2(B + 1)$ owner blocks. The zero-potential collar has at most $K + m$ possible rooms and the full-potential collar has deficit at most K . Marking the first $L_1 = K + m + 1$ days of every owner block, as well as every inefficient day, leaves at most

$$E = B + 2L_1(B + 1) \quad (145)$$

nonclean days. On each clean block the mate is globally full or empty, with bit zero, and the owner's surpluses alternate $b, b + 1$. The frontiers have a fixed orientation throughout this block: the bridge preserves orientation, and the next minimum-surplus phase has the same orientation as the preceding minimum-surplus phase.

There is a timing distinction when the graph vertices are survivors. At such a vertex write e for the *newly set* bit, referring to that survivor itself. If consecutive survivor sizes are r, r' and the first has surplus s , then $r' = r + s - m$. Their constant-drift rank is $2r - e$, rather than the before-inspection rank $2k + e$:

$$(2r - e) - (2r' - e') = D \quad (146)$$

in both alternating cases $(s, e, e') = (b, 1, 0)$ and $(b + 1, 0, 1)$. The ordinary average occupied length is nonincreasing and may be stationary for one round when $m = b + 1$. Its corrected version $(2r - e)/(2A)$ decreases by exactly $D/(2A)$ per edge and differs from it by at most $1/(2A)$. Use the enlarged margin $v = A + 2$ to cover this correction, the fiber range, and the direct height replacements.

These frontiers are also physical height configurations. If l_v is the occupied paired-column length and $\sigma_v \in \{0, 1\}$ its physical longitudinal parity, then

$$a_v = 2(l_v - 1) + \sigma_v.$$

In either phase, the boundary-free color has parity zero and its length equals its neighbor's length or exceeds it by one. Hence $|a_u - a_v| = 1$ across every Q -edge. The bridge lemma ensures this on the intervening phase as well. For a left frontier the sum of the parity indicators is $b + 1 - e$, so the sum of physical heights differs from $2r - e$ by a constant. The drift in (146) is therefore the same physical height drift used in Lemma 24.4.

Apply that lemma with $g = \gcd(A, D)$ and the conservative diameter bound $A - 1$. We may use

$$K_0 = \left\lceil \frac{A(A - 1)}{m} \right\rceil, \quad \ell_0 = 2A/g, \quad \delta = D/g, \quad \Delta = 2\delta. \quad (147)$$

A crossing of length at least $K_0 + \ell_0$ can be replaced by one shorter by ℓ_0 days, shifting the initial endpoint by δ paired columns, or by one longer by ℓ_0 with the opposite shift. The construction connects actual physical height endpoints, without an additional assumption about intermediate survivor shapes.

We spell out the changes to the protected-band argument to ensure that this is a physical reduction. Mark the paired-column index of every nonclean inspection and the initial frontier cuts of every nonclean block. Together with the ends there are at most

$$F = (A + m)(E + 1) + 2$$

marks. A physical move crosses at most one pair boundary, so radius $\rho = E + 2A + m + 2$ shields all nonclean blocks. Local full preservation uses the longitudinal matching *inside each pair*; it requires no perfect matching in the odd-order graph Q . Local empty preservation follows from finite propagation. Thus the argument of Section 24.3 applies in the paired coordinate. Choose a raw gap of length $W + 2v + 4$ and trim $v + 2$ at each end. The corrected mean, its nonzero speed, and the margin v give one complete clean crossing per cohort through the retained band. Their mates are globally full or empty and the crossings are disjoint in time.

Explicit constants are obtained by putting

$$\begin{aligned} U &= K_0 + \ell_0 + 4, \\ W &= 2(\delta + v + m + 3) + \left\lceil \frac{D(U + 4)}{2A} \right\rceil + 2v + 2m + 10, \\ L_* &= 4 + \left\lceil \frac{4K + 2m + 1 + D}{2A} \right\rceil + F(2\rho + 1) + (F + 1)(W + 2v + 4) + \delta, \end{aligned} \quad (148)$$

and taking $N = 2L_*$. The corrected mean moves by $D/(2A) \leq m$ per edge, so the entry and exit overshoots cost at most $2m$ paired positions. The stated W leaves at least U edges in a central path, with its endpoints on opposite sides of a δ -pair band and all intermediate cuts more than δ from the outer edges. The threshold ensures $C > D$, the raw gap, and the same conditions after contraction.

Delete that central band of δ pairs and replace each central crossing by the walk shorter by ℓ_0 supplied by Lemma 24.4. Start a left frontier shifted by $-\delta$ pairs, equivalently $-\Delta$ physical columns; its new endpoint is exactly the old endpoint. Reflect this for a right frontier. Every new edge is a valid physical-height transition, and the constant mean drift and bounded range keep its filled-tail configuration interior. Outside the crossings, neighborhood formation commutes with the natural column map because the protected neighborhood is uniformly full or empty, including every nonclean intermediate state. No nonclean inspection is deleted. The mate's full or empty state is preserved; the removed time is even, and paired-column translation preserves phase. As survivor edges count the following inspection, exactly ℓ_0 days are removed per cohort. This proves

$$T_m(Q \square P_{n-\Delta}) \leq T_m(Q \square P_n) - 2\ell_0.$$

Starting instead from a sufficiently long shorter board, insert the same number of pairs and use the longer endpoint connection to add ℓ_0 days to each crossing. The same filled-tail, endpoint, and mate arguments prove the reverse inequality, exactly as in Section 24.5. Since $2\ell_0 = 2A\Delta/D$, this proves Theorem 24.5 on every sufficiently long even cylinder, with the displayed effective threshold.

24.8 Combining the parity classes

Proof of Theorem 24.1. Every transverse box has a Hamiltonian path by the snake construction used in Theorem 23.8. If $A \geq 2$ is even, Theorem 24.2 already supplies the asserted recurrence for both longitudinal parities. If $A \geq 3$ is odd, every nontrivial transverse side is odd. Theorem 22.3 applies to odd longitudinal lengths, while Theorem 24.5 applies to even lengths. Both give exactly the same even period $\Delta = 2(2m - A)/\gcd(A, 2m - A)$ and the same increment $4A/\gcd(A, 2m - A)$. Taking the larger threshold proves the common recurrence on both parity classes.

If $A = 1$, the cylinder is a path. For $m = 1$ the formula $T_1(P_n) = 2n - 4$, $n \geq 3$, gives period two and increment four. For $m \geq 2$, take $\Delta = 2(2m - 1)$ and $N = m + 1$ in (1). This changes its ceiling by four and preserves both the exceptional congruence and the parity correction, giving the required increment $2\Delta/(2m - 1) = 4$.

For $A \geq 2$ and integer $m \leq A/2$, Theorem 23.8 gives impossibility whenever $n \geq 2 \lfloor A/2 \rfloor$. If $A = 1$, the only such budget is zero; no target on a nonempty path can be captured without an inspection.

For completeness the displayed sufficient time thresholds are polynomial. With $X = A + m + 1$, their definitions give $K, B, L_1 = O(X^2)$, $E = O(X^4)$, $F = O(X^5)$, and $\rho = O(X^4)$. Because $m > A/2$, we have $K_0 \leq 2(A - 1)$ and $\ell_0 \leq 2A$, so $U = O(A)$ and $W = O(X)$. The dominating threshold term $F\rho$ is $O(X^9)$. For the earlier odd-length box theorem, $S = \sum_i (n_i - 1) \leq A - 1$, hence its corner bound satisfies $Z = A(S + 1) \leq A^2$. Substitution into (111) also gives a polynomial bound, dominated by $O(X^9)$. The path threshold is smaller. Taking the larger parity threshold thus preserves a uniform polynomial sufficient bound. $\square$

At the eventual minimum budget $m = \lfloor A/2 \rfloor + 1$, the period is two. The time increase on adding two columns is $2A$ when A is even and $4A$ when A is odd. The polynomial onset bound is sufficient, not a claim of the earliest length at which this recurrence holds.

The theorem fixes the entire cross-section and the budget while the last side grows. It does not supply the smallest period, the finite exception values, or a simple formula for every arbitrary finite box. The proof covers every transverse box; for a general odd-order Hamiltonian graph the new paired-column theorem only asserts the even-length case. These eventual-period results have ordinary proofs and independent review, rather than complete Lean formalizations.

25 Exact finite interfaces for even-area cylinders

The efficient-frontier arguments yield more than an eventual formula. They identify an exact finite collection of local search problems. The symbolic height theorem then evaluates every connection through the interior directly, leaving a graph of bounded size. Arbitrary intermediate supports are retained inside the boundary problems; a daily pyramid normal form is not assumed.

Theorem 25.1 (An exact finite-interface representation). *Let Q be a fixed finite bipartite graph with a Hamiltonian path, with $A \geq 2$ vertices, and fix an integer $m > A/2$. For every positive n such that An is even, the following construction determines the exact value $T_m(Q \square P_n)$ and an optimal physical strategy.*

After an effective finite preprocessing depending only on Q, m , the value is obtained using $O_{Q,m}(1)$ integer arithmetic stages on $O_{Q,m}(\log(n+1))$ -bit integers. Its final graph has boundedly many vertices depending only on Q, m . Every vertex has at most one proper cohort, which is a physical pyramid; its mate is full or empty. Boundary edges are actual search blocks of bounded duration, retaining arbitrary nonpyramidal behavior in bounded longitudinal slabs. Interior edges have the exact symbolic costs in (55) and admit physical pyramid realizations. The sufficient size threshold and all bounds are explicit below.

The calculation gives a compressed description of an optimal strategy. Expanding it into individual inspections additionally costs its output length. No small bound on the parameter-dependent preprocessing is asserted.

For rectangles this covers every even-area board with both sides at least two, at every feasible budget, by choosing the shorter path as Q . The one-row case is already solved. Theorem 7.1 covers odd-area rectangles. These are uniform exact algorithms; the closed formulas in earlier sections remain useful explicit evaluations in their stated budget ranges.

This theorem is a geometric strengthening, rather than an asymptotic speed claim over an eventual-period formula with its entire finite prefix precomputed. It specifies the exact local transition problems and a physical strategy reconstruction. The finite preprocessing described here is a mathematical construction, not a claimed implemented generic compiler.

Put $\alpha = \lceil A/2 \rceil$ and define

$$\begin{aligned} K &= A^2, & B &= 8A^2 + 8m + 2A, \\ E &= B + 2(K + m + 1)(B + 1), & L &= E + 1, \\ q &= (m + \alpha)L, & R &= 100(q + A^2 + L + A + 1), \quad n_0 = 10R. \end{aligned} \quad (149)$$

These constants depend only on A, m . We prove the result for $n \geq n_0$. Every smaller admissible length is a finite exception, determined by the ordinary exact belief recurrence. The displayed threshold dominates the clipping guards in both parity cases of Section 24, and gives $An/2 > mL$.

25.1 Physical pyramids, including the end boundaries

Index Q by a Hamiltonian order $0, \dots, A-1$, so its bipartition is the parity of this index. A downward pyramid is a one-color set P such that

$$(v, y) \in P, \quad y > 0, \quad uv \in E(Q) \implies (u, y-1) \in P.$$

An upward pyramid is its longitudinal reflection. Empty and full color classes are allowed in either orientation. For current color p , put $s_v = (p - v) \bmod 2$. A downward pyramid has a spacing-two

prefix in each fiber. If that prefix has k_v rooms, its virtual last height is

$$a_v = s_v - 2 + 2k_v.$$

The predecessor condition is equivalent to

$$-2 \leq a_v \leq n-1, \quad a_v \equiv p-v \pmod{2}, \quad |a_u - a_v| = 1 \quad (uv \in E(Q)). \quad (150)$$

The values $-2, -1$ handle empty fibers. Two predecessor moves give spacing-two closure, and the higher cutoff forces the required adjacent predecessor; this proves both directions, including the bottom row. Along the Hamiltonian path the heights form a ± 1 walk, so their range is at most $A-1$ and there are at most $(n+2)2^{A-1}$ downward pyramids of each color. Extra edges of Q impose additional conditions on these walks.

Neighborhoods preserve this family at the physical ends. Indeed, let $(v, y) \in N(P)$, $y > 0$, and consider its predecessor $(u, y-1)$. A witnessing neighbor $(v, y-1) \in P$ is already adjacent to that predecessor. A witness $(v, y+1) \in P$ supplies $(u, y) \in P$. A transverse witness $(z, y) \in P$ supplies $(v, y-1) \in P$. In all cases the required predecessor belongs to $N(P)$.

Writing t'_v for the largest coordinate below n of parity $1 - s_v$, the exact neighborhood cutoff is

$$\eta(a)_v = \max(h_v, \{a_u : uv \in E(Q)\}), \quad h_v = \begin{cases} (1 - s_v) - 2, & k_v = 0, \\ \min(a_v + 1, t'_v), & k_v > 0. \end{cases}$$

In particular $\eta(a)_v \leq a_v + 1$. The simpler formula $a_v + 1$ alone need not be exact in an empty bottom fiber. A perfect matching of the cylinder is supplied by transverse Hamiltonian pairs when A is even, or by longitudinal column pairs when A is odd and n is even. Thus every support S satisfies $|N(S)| \geq |S|$. The two fiber parity sums differ by at most one, giving

$$0 \leq |N(P)| - |P| \leq \alpha, \quad |P| \leq |N^r(P)| \leq |P| + \alpha r. \quad (151)$$

Every $N^r(P)$ has the same pyramid orientation.

A pyramid with at most u rooms lies within $2u + A + 2$ columns of its filled end. Its complement in its color class is an oppositely oriented pyramid, so the same bound locates a deficit of at most u near the unfilled end. Finally, Lemma 11.2 bounds a pyramid with an empty physical fiber by $H_Q \leq A^2$ rooms. The height range also confines it to an end band of width at most A .

25.2 Canonical checkpoints of bounded separation

Call a full-board belief *canonical* if at most one cohort is nonempty and nonfull, that proper cohort is a pyramid, and its mate is full or empty. Initial full/full and final empty/empty beliefs are canonical.

The clean-day arguments already proved in Section 24 give the following interface. For even $A = 2b$, the potential loss is at most $m - b$ per day; every efficient owner survivor has surplus b and is a physical pyramid by Lemma 24.3 and (136). For odd $A = 2b + 1$ and even n , the history-bit potential has daily bound $2m - A$. Its two efficient cases are exactly (144). Surplus b gives a pyramid by Lemma 24.6; surplus $b + 1$ also gives a pyramid by the actual containment hypothesis in Lemma 24.7. Thus both efficient phases have the physical height condition (150) on every Q -edge.

The B in (149) dominates the slack bound in both cases. At most B days are inefficient, and there are at most $2(B + 1)$ efficient owner blocks. Mark their first $K + m + 1$ days and all inefficient

days. There are at most E marked days. The proved untouched-cohort growth argument makes the mate literally full or empty on every remaining clean day, rather than merely placing it in a flat potential collar. Its source is the neighborhood of a preceding efficient pyramid, so both its source and successor are canonical.

Partition an optimal physical strategy at the sources of its clean days, also including the initial and terminal states. Following the last successful inspection by the empty movement step does not add a day. Every resulting block has duration at most $E + 1 = L$: except for its possible first clean day, all intervening days are marked. If there are no clean days, the whole strategy has duration at most E .

Consequently some optimum is a path of physical blocks of length at most L between canonical states. Conversely every such concatenation is an actual strategy. No history bit is needed in its state: the bit establishes the existence of checkpoints in an optimum, whereas edge legality is checked on the physical supports themselves.

25.3 Localizing a physical transition

Lemma 25.2 (Exact endpoint localization). *Suppose an r -day block, $r \leq L$, takes a one-cohort pyramid P_0 to a pyramid P_1 , possibly of the opposite orientation. Set $P_2 = N^r(P_0)$ and $D = P_2 \setminus P_1$. Then $P_1 \subseteq P_2$ and $|D| \leq (m + \alpha)r$. Removing every inspection whose longitudinal distance from D exceeds r preserves the final support exactly. Moreover D lies in an interval of at most*

$$2(m + \alpha)r + 2A^2 + 4A + 4$$

longitudinal columns.

Proof. The matching gives $|P_1| \geq |P_0| - mr$, while (151) bounds $|P_2|$, proving the cardinality claim. Removing inspections enlarges the final support. Any newly surviving walk must finish in D , since its endpoint is freely reachable but not in P_1 . Every earlier point of its r -step walk is within longitudinal distance r of D . All original inspections that could hit it were retained, a contradiction. The endpoint is therefore unchanged.

For equal orientations, compare the two included height vectors of $P_1 \subseteq P_2$. Each fiber's cutoff difference is twice its contribution to $|D|$, and the individual height ranges are at most $A - 1$. This confines D to at most $2|D| + 2A + 2$ columns.

Suppose P_2 is downward and P_1 upward. If every fiber of P_1 is nonempty, it contains the topmost room of its color in each fiber, forcing P_2 to be full. Its difference D is then the downward complement of P_1 , of size at most $(m + \alpha)r$. Otherwise $|P_1| \leq A^2$, whence $|P_2| \leq A^2 + (m + \alpha)r$; this downward pyramid lies in the corresponding bounded initial band. Reflection handles the other case. Empty and full sets may be assigned either orientation. The stated bound covers all cases. $\square$

Apply the lemma separately to the two initial-color cohorts. Their current colors are disjoint on every day, so the retained inspections still respect the shared daily quota. The lemma concerns actual avoiding walks, with no condition on intermediate support shapes.

It also gives an effective finite test for an edge. If I is the difference band, all shots lie in its r -enlargement; only endpoints in its $2r$ -enlargement can be affected. Every r -step walk to such an endpoint remains in the $3r$ -enlargement. Enumerate the daily shot sets there, simulate from the restricted initial support in the largest band, and compare final membership on the middle

band. Outside that checked band the uninspected baseline already agrees with P_1 . First reject an endpoint unless $P_1 \subseteq N^r(P_0)$, an exact cutoff comparison. This proves an exact physical edge test, including the presence as well as absence of endpoint rooms.

25.4 Only one longitudinal coordinate remains

If the same cohort is proper at both ends with the same orientation, its cardinality changes between $-mr$ and αr . Height ranges are at most $A - 1$, and the parity-sum correction is at most one. The two cutoff positions consequently differ by at most $2(m + \alpha)r + 4A + 6 < R$. Away from the board ends the localized block depends only on relative height shapes, bounded displacement, color, and duration. Translation by two columns preserves its complete physical test.

Every other possibility is confined to end collars. A proper cohort becoming empty has initial size at most mr ; one becoming full has initial deficit at most αr . A full cohort becoming proper has final deficit at most mr , while an empty cohort cannot become proper. For an orientation change, the two cases in Lemma 25.2 show either that both endpoint deficits are at most $(m + \alpha)r$, or that both sizes are at most $A^2 + (m + \alpha)r$. All their cutoffs are therefore in bounded end collars. A full mate ending full needs no inspections by the same localization lemma; a full mate cannot become empty in $r \leq L$ days because $An/2 > mL$.

An interior state is encoded by its physical color, orientation, mate status, a relative height walk, and

$$z = (a_0 - p)/2.$$

For upward pyramids use the physical cutoff of their downward complement, so both orientations use the same longitudinal coordinate. Filter the Hamiltonian height walks by the additional Q -edge conditions. This gives $O_A(n)$ canonical states. The preceding bounds give finitely many interior edge types of bounded displacement and finitely many boundary edge types. Opposite-end gadgets are tested together with the shared daily quota; for $n \geq 10R$ their movement cones cannot communicate within L days. Their only dependence on length is the checkerboard parity at the far end.

Let $\mathcal{C}_n$ contain these states and all physical edges of duration $1, \dots, L$. It has $O_{Q,m}(n)$ edges. Its shortest-path value from full/full to empty/empty is exactly T_m : an optimum supplies a path by the checkpoint argument, and every graph path expands to an actual strategy. Nonpyramidal intermediate supports are kept inside the finite edge tests, rather than projected away.

25.5 Finite ports and exact symbolic interior costs

Use the common homogeneous counter interval

$$R \leq z \leq \lfloor (n - 2 - 2R)/2 \rfloor.$$

The first R and last R levels are called *ports*. Keep every vertex outside this interval, every port vertex, and all original boundary edges between retained vertices. There are only boundedly many of them depending on Q, m . The strict displacement bound $< R$ ensures that an edge crossing between the interior and a boundary family meets a port. Retain the original edges joining the two remote boundary families, with their shared-quota tests; their legality depends only on the longitudinal parity once $n \geq n_0$.

The proper cohort, its orientation, and its full or empty mate cannot change on an interior stretch: all such changes were confined to end collars above. Call these fixed data its *sector*.

Endpoint physical colors may differ, and give the required duration parity. For every ordered pair of ports in the same sector, add a direct edge of weight t_0 from (55), using the proper cohort's height vectors. Reflect longitudinal coordinates for an upward sector. The same construction includes connections returning to the same end.

We verify both directions of this replacement. Every original interior edge has duration $r \leq L$. The large collar in (149) implies the height margins in (53), and its initial proper support has more than $C_A + mr$ rooms. Indeed its cutoff position is at least $2R - A$ from its filled end, while $R > 100(mL + A^2 + L + A)$. Theorem 12.3 therefore bounds that edge against arbitrary intermediate supports by

$$z_v \leq a_v + r, \quad \sum_v a_v - \sum_v z_v + Ar \leq 2mr.$$

If its mate is full at both ends, omitting every mate inspection preserves that endpoint; an empty mate remains empty. Thus assigning the proper cohort the entire budget gives the relevant comparison. Summing these inequalities along any interior path makes the height sums telescope and gives exactly the criterion in (55). Its duration is at least t_0 . This lower bound applies the guarded theorem to each *bounded* edge, never to a long block with an unjustified $|P| - mt > C_A$ guard.

Conversely, port cutoffs are far enough from both physical ends that (54) holds. Proposition 12.4 realizes the direct edge in exactly t_0 days, allocating at most m inspections to the proper cohort and none to its full or empty mate. For $t_0 = 0$ the endpoints coincide. The realized block need not stay in the auxiliary counter interval: it stays within the actual physical board, which is the condition needed for an upper strategy.

Take an optimal path in $\mathcal{C}_n$ and split it at visits to the retained boundary families. Every omitted stretch has port endpoints in one sector. Replacing it by the corresponding direct edge cannot increase its cost. Conversely, every edge of the new graph is an actual physical block, so every new graph path is a valid strategy. The new finite graph therefore has exactly the same optimal value T_m . This argument allows arbitrary reverse travel, repeated visits to the same boundary, opposite-end gadgets, and owner changes within the retained boundary graph; it assumes neither two pure sweeps nor monotone excursions.

Port heights are affine functions of n , with the end parity fixed. Their edge weights require maxima, one integer ceiling, and one parity adjustment. One fixed-size shortest-path closure consequently evaluates all paths using $O_{Q,m}(1)$ integer arithmetic stages after finite preprocessing. The finitely many smaller lengths $n < n_0$ remain in that preprocessing, rather than being discarded by an asymptotic claim. Integer bit lengths are $O_{Q,m}(\log(n + 1))$.

Store a shortest path after removing cycles. It has boundedly many edges. Each boundary edge stores its finite local inspections. Each interior edge stores its endpoints, duration, the balanced quota rule in Proposition 12.4, and the legal height-descent rule. This gives a bounded-size description of an optimal physical strategy; expanding its individual inspections costs the output length. The construction does not claim that all parameter-dependent boundary tables have been implemented. Its strengthening is the exact symbolic elimination of the interior transition problems, preserving all physical boundary and shared-quota effects. This completes the proof of Theorem 25.1.

26 A spatial transfer matrix for a fixed deadline

The preceding reductions retain geometric information about the current belief. A different choice of state removes every parity restriction: fix the number of inspection days, and retain the history

of each column. The resulting matrix is independent of the longitudinal length. It gives an exact recurrence and an eventual formula for the inverse tradeoff.

Let Q be any nonempty finite simple graph, put $A = |V(Q)|$, and let $b_T(n)$ be the minimum daily budget guaranteeing capture within T days on $Q \square P_n$, where $n \geq 2$ and $T \geq 1$.

Theorem 26.1. *For fixed Q, T , there are effectively computable integers $n_0, p > 0$ and rational constants $c_0, \dots, c_{p-1}$ such that*

$$b_T(n) = \frac{An}{T} + c_{n \bmod p} \quad (n \geq n_0).$$

The period can be chosen divisible by $T/\gcd(A, T)$. The finitely many remaining values, and a winning strategy at any specified length, are also effectively computable. No bipartiteness assumption on Q is needed.

Here “effectively computable” asserts a terminating finite procedure, not a small uniform complexity bound. The state space below has $2^{2A(T-1)}$ states. In particular, this theorem does not supply a fixed matrix when the deadline itself grows with the box length.

26.1 Exact local certificates

The case $T = 1$ is $b_1(n) = An$. Assume $T \geq 2$, and use the finite alphabet $\mathcal{W} = (2^{V(Q)})^{T-1}$. A letter $b = (b_1, \dots, b_{T-1})$ records proposed survivor sets in one column. Put $b_T = \emptyset$, and let 0 denote the all-empty letter. For $a, b, c \in \mathcal{W}$ define

$$\begin{aligned} w_1(a, b, c) &= A - |b_1|, \\ w_t(a, b, c) &= |(N_Q(b_{t-1}) \cup a_{t-1} \cup c_{t-1}) \setminus b_t| \quad (2 \leq t \leq T). \end{aligned} \tag{152}$$

A word $b^1 \cdots b^n$, with boundary letters $b^0 = b^{n+1} = 0$, has day- t cost $\sum_{j=1}^n w_t(b^{j-1}, b^j, b^{j+1})$.

Lemma 26.2. *A successful T -day strategy with respective daily capacities $m_1, \dots, m_T$ exists if and only if such a word has day- t cost at most m_t for every t .*

Proof. Let R_t be the union of the word’s day- t column sets, with $R_T = \emptyset$. Define nominal beliefs $B_1 = V(Q \square P_n)$ and $B_t = N(R_{t-1})$ for $t \geq 2$, and inspect $S_t = B_t \setminus R_t$. The product-neighborhood identity shows that $|S_t|$ is exactly the local cost sum in (152).

The actual belief U_t satisfies $U_t \subseteq B_t$ inductively: its survivors are contained in R_t , so their neighborhood is contained in B_{t+1} . On day T , the empty survivor envelope makes $S_T = B_T$ cover every possibility. This argument intentionally does not require $R_t \subseteq B_t$; extra envelope vertices do not compromise soundness.

Conversely, take the actual survivors of a successful strategy, discard inspections outside the actual belief, and extend an earlier capture by empty inspections and survivors. Its column sections form a word whose costs are exactly the useful inspection counts. $\square$

Use pairs $(a, b) \in \mathcal{W}^2$ as matrix states. The transition $(a, b) \rightarrow (b, c)$ has monomial weight $\prod_{t=1}^T x_t^{w_t(a, b, c)}$. Let $M(\mathbf{x})$ be this matrix; let u indicate states $(0, b)$ and v indicate states $(b, 0)$. There is exactly one n -edge walk for each n -column word, so

$$F_n(\mathbf{x}) = u^\top M(\mathbf{x})^n v$$

has a positive coefficient at $\mathbf{x}^{\mathbf{c}}$ precisely when a certificate with exact cost vector $\mathbf{c}$ exists. Thus

$$\sum_{n \geq 0} F_n(\mathbf{x}) z^n = u^\top (I - zM(\mathbf{x}))^{-1} v$$

is a rational formal series encoding the exact daily-budget tradeoff. All coefficients are nonnegative; multiplicities of envelopes do not affect the existential criterion.

For a direct algorithm, retain the accumulated cost vector and discard it if any coordinate exceeds the proposed budget m . This uses at most $O(n|\mathcal{W}|^3(m+1)^T)$ transitions. Since an optimum never needs $m > An$, searching budgets remains polynomial in n for fixed Q, T . Backtracking recovers the word and the physical inspections in the lemma.

26.2 Why the optimum has an eventual formula

A subset of $\mathbb{N}^r$ is *linear* if it is a translate of a finitely generated additive monoid, and *semilinear* if it is a finite union of linear sets. Add a length coordinate one to each edge weight. The attainable vectors $(n, c_1, \dots, c_T)$ form an effectively semilinear set L . To see this directly, eliminate the finite automaton's states to obtain a regular expression. Its additive image is computed using union, Minkowski sum and additive closure, all of which preserve semilinearity. For the only less immediate operation, if $S = \bigcup_i (b_i + \text{monoid}(P_i))$, then

$$S^* = \bigcup_I \left(\sum_{i \in I} b_i + \text{monoid} \left(\{b_i : i \in I\} \cup \bigcup_{i \in I} P_i \right) \right).$$

Here I ranges over subsets of component indices. Reserve one base for each used component; extra bases account for repeated occurrences, and all periods can be assigned to the reserved occurrence. The empty I supplies zero. This proves both containments and effectivity.

By the effective equivalence of semilinear and Presburger-definable sets [8, Theorems 1.1 and 1.3], the relation

$$E(n, m) \iff n \geq 2 \text{ and } \exists \mathbf{c} ((n, \mathbf{c}) \in L \text{ and } c_t \leq m \text{ for all } t)$$

is effectively Presburger-definable. The same holds for $E(n, m) \wedge (m = 0 \text{ or } \neg E(n, m - 1))$, which is the graph of b_T . Convert that graph to a finite union of linear sets.

Every infinite linear component of a function's graph lies on a rational affine line. Indeed, a nonzero period cannot have zero first coordinate. For two periods $(a_i, d_i), (a_j, d_j)$ with positive first coordinates, using them a_j and a_i times gives the same input, so uniqueness of the output forces $a_j d_i = a_i d_j$. All periods therefore have one slope. The component's input projection is a translated numerical semigroup; it eventually contains exactly one residue class modulo the gcd of its positive generators. A threshold is computable by shortest paths among residues modulo any generator. Taking a common multiple of the finitely many periods and a maximum of their thresholds gives an affine formula on each eventual residue class. Finite components contribute only finitely many exceptions.

It remains to identify the slopes. The longitudinal matching between columns $(1, 2), (3, 4), \dots$ supplies $2A \lfloor n/2 \rfloor$ alternating target trajectories, pairwise disjoint on every inspection day. A probe intercepts at most one of these trajectories on that day. Consequently

$$Tb_T(n) \geq 2A \lfloor n/2 \rfloor \geq A(n - 1).$$

For the upper bound put $q = \lceil n/T \rceil$. When the possible rooms lie in a suffix beginning at column L , inspect its first $q + 1$ columns, clipped to the board. After movement the suffix begins at least at $L + q$. After $T - 1$ days at most $n - (T - 1)q \leq q$ columns remain; inspect all of them on the final day. Thus

$$b_T(n) \leq A(\lceil n/T \rceil + 1) < An/T + 2A.$$

Internal Q -edges never change the column index. The two bounds force every infinite affine component to have slope A/T . Enlarging the period by $T/\gcd(A, T)$ if necessary proves Theorem 26.1, including its effective exception table and integral additive increment.

The survivor-envelope equivalence and the exact product-fiber cost sum are verified in Lean by **SurvivorEnvelope**. The matrix and semilinearity arguments above are ordinary proofs. The reference transfer implementation was compared with the independently proved path and two-row formulas in 56 cases, with physical replay of its recovered inspection sets. Neither these finite comparisons nor the formal semantic lemma are being substituted for the general periodicity proof.

27 Formal verification, reproducibility, and open problems

The development comprises 83 source modules using Lean 4.33.1 and Std. There are no admitted proofs, project-specific mathematical axioms, or trusted external solver results. Finite certificates are reduced by the Lean kernel. Printed dependencies use only propositional extensionality, classical choice, and quotient soundness where needed.

27.1 Complete classifications and reusable formal theorems

Four classifications are complete for the physical game:

- **PathClassification** covers every positive path length and budget, including the one-room and one-probe cases.
- **LadderGame** covers every two-row length and budget, including the single-edge degeneration.
- **ThreeCubeTimeClassification** covers every budget on $3 \times 3 \times 3$, including the minimum-budget eighteen-day theorem.
- **FourCubeClassification** covers every budget on $4 \times 4 \times 4$. Its forty-day lower bound includes the concrete compression, shape-sensitive certificate, and interpretation for arbitrary room-valued inspection sequences.

Each includes feasibility, optimal time, and attaining schedules against all legal target walks. Independent semantic reviews check the physical graph, daily budget, and indexing: formal round zero is the first inspection, so last-round index t means $t + 1$ days.

The reusable developments formalize the following separate layers. **BeliefSemantics** equates possible positions with avoiding walks; **CaptureRecurrence** verifies the winning recurrence and losing traps. **StrategyCompression** compares whole strategies, including arbitrary daily budgets. Product lifting, finite stabilization, and the concrete square operators used for the four-cube theorem are also proved. The generic complement/inverse identity in **BipartiteProfileDuality** uses actual attained graph minima, with inverse arithmetic in **ProfileInverse**. **SurvivorEnvelope** connects local column certificates to actual walk capture

and global daily budgets. **MiddleIntervalTransfer** proves the exact bounded-counter interval theorem and constructs its legal daily allocations. **RootConvolution** defines the pronic and square roots internally and proves the exact two-candidate interval minimum. The more general **ConvexCapacityConvolution** constructs least inverses of monotone unbounded discrete-convex integer capacities and derives endpoint maximization from increasing increments. It proves an attained two-candidate interval minimum, with concrete instances for the square, pronic, and every punctured capacity C_h of (13). The triangular identity and the convexity of the maximum and truncated quadratic are proved inside Lean. **ProbeLocalization** proves on arbitrary directed graphs that retaining inspections only in the backward cone of unwanted endpoints preserves the final possible-position set exactly.

FastestAncestry proves concentration, low-total persistence, fastest-cohort ancestry after a pure reset, and the resulting midpoint obstruction for two deficit recurrences. Its inverse-concentration, proper-input, and bounded-secondary trace hypotheses are explicit. **ClippedPyramidErosion** uses actual nonnegative cylinder cells: it proves neighborhood inclusion after clipped erosion and the exact loss of occupied bottom roots, including empty fibers and both finite ends. Its room lists have unique entries and the asserted physical cardinalities.

27.2 Ordinary proofs and partial formalization

The other graph classifications and unbounded geometric theorems in this manuscript have ordinary proofs, not complete Lean proofs. This includes three through five rows; both general rectangle profile theorems; the uniform one-day determination and large-budget formulas; all-odd-box nesting and the other higher-dimensional reductions; the cylinder time formulas and eventual periods; and fixed-deadline semilinearity. The physical application of the formal middle-transfer lemma and the geometric bounded-interface application of probe localization are also ordinary proofs. The bounded odd frontier and its acceleration are not claimed to be Lean-verified by the root-convolution module.

Some arithmetic components of these results are formalized separately: the five-row scalar ranks, deficit consequences, four-probe transition and equality lemmas, three finite large-budget bounds, and the uniform minimum-budget scalar potential in **UniformOddRank**. These checks do not formalize their geometric hypotheses or the full unbounded classifications. Likewise, formal survivor-envelope semantics does not formalize the subsequent semilinearity argument. The uniform all-width minimum-budget formula also uses ordinary proofs of the rectangle inverse formulas, fresh-secondary clock, and physical midpoint reduction; **FastestAncestry** does not supply these instantiations. The punctured-quadrant and conditional half-strip neighborhood theorems, their every-size geometric attainment, and the corner propagation delays remain ordinary mathematics despite the formal inverse-capacity arithmetic. The erosion module does not prove the distance thresholds ensuring root occupancy, prescribed-size ideal enlargement, or maximal-erosion corrections at the top boundary. The detailed module-to-theorem map in `lean/README.md` records these distinctions. Successful compilation must not be described as formal verification of every theorem in the manuscript.

27.3 Reproduction and open scope

The companion at <https://angelraychev.com/princess/> provides the presentation, complete proof text, formal sources, and mathematical archive. The archive includes independent physical replays, finite certificates, counterexample searches, and dated research notes. The verification receipt records the compiler version, exact source hashes, dependency order, axiom reports, and exit

codes. Reproduce the complete check with `python3 src/check_lean.py -lean /path/to/lean` from the archive root. Compilation is sequential with one worker and requires no extra Lean packages.

The central open target is a simple explicit optimal-time formula for every rectangle and budget. Static neighborhood minimization is solved for all rectangles, but successive minimizing shapes can be incompatible. On odd rectangles the exact bounded frontier decides between the two times left by the solo recurrence; whether the two pure solo endpoints alone always suffice remains unresolved. The large-budget formulas now cover all side parities, with their stated quadratic thresholds. Below those thresholds, the exact algorithms and eventual theorems do not supply a single quotient-and-remainder formula. General higher-dimensional optimal time remains open as well.

27.4 Research and writing provenance

The path results originated in the author’s October 2019–March 2020 work. Reconstructing shorter proofs does not reassign that credit. The 2026 grid and box investigation, new arguments, experiments, formal developments, and manuscript were developed with substantial assistance from Astra 6 through the Codex harness. The research log distinguishes recovered results, established literature, new derivations with unverified priority, failed approaches, and open claims. The human author retains responsibility for the manuscript and any eventual submission.

Acknowledgments

The author thanks Dimitar Rusev for writing the preliminary section on monotonicity (Section 3) in the 2020 student-conference version of this work. By agreement, his contribution to that version is acknowledged here; both authors remain credited in its bibliographic entry. The author also thanks his mathematics teacher and mentor Dimitar Dimitrov, who encouraged him to pursue mathematical research and guided his early work.

References

- [1] Tatjana V. Abramovskaya, Fedor V. Fomin, Petr A. Golovach, and Michał Pilipczuk. How to hunt an invisible rabbit on a graph. *European Journal of Combinatorics*, 52:12–26, 2016. doi: 10.1016/j.ejc.2015.08.002. URL <https://fedorvf.github.io/articles/2016/2016g.pdf>.
- [2] Nikolay Beluhov and Emil Kolev. Search for a moving target in a graph. *Electronic Notes in Discrete Mathematics*, 57:39–46, 2017. doi: 10.1016/j.endm.2017.02.008. URL <https://doi.org/10.1016/j.endm.2017.02.008>.
- [3] Walid Ben-Ameur, Harmender Gahlawat, and Alessandro Maddaloni. Hunting a rabbit: Complexity, approximability and some characterizations. *Theoretical Computer Science*, 1075:115946, 2026. doi: 10.1016/j.tcs.2026.115946. URL <https://arxiv.org/abs/2502.15982>. Preprint first posted in 2025.
- [4] Sergei L. Bezrukov and Oriol Serra. A local–global principle for vertex-isoperimetric problems. *Discrete Mathematics*, 257(2–3):285–309, 2002. doi: 10.1016/S0012-365X(02)00431-4. URL [https://doi.org/10.1016/S0012-365X\(02\)00431-4](https://doi.org/10.1016/S0012-365X(02)00431-4).

- [5] Jessalyn Bolkema and Corbin Groothuis. Hunting rabbits on the hypercube. *Discrete Mathematics*, 342(2):360–372, 2019. doi: 10.1016/j.disc.2018.10.011. URL <https://arxiv.org/abs/1701.08726>.
- [6] John R. Britnell and Mark Wildon. Finding a princess in a palace: A pursuit-evasion problem. *The Electronic Journal of Combinatorics*, 20(1):P25, 2013. doi: 10.37236/2296. URL <https://arxiv.org/abs/1204.5490>.
- [7] Thomas Dissaux, Foivos Fioravantes, Harmender Gahlawat, and Nicolas Nisse. Further results on the Hunters and Rabbit game through monotonicity. *Information and Computation*, 305: 105302, 2025. doi: 10.1016/j.ic.2025.105302. URL <https://arxiv.org/abs/2309.16533>.
- [8] Seymour Ginsburg and Edwin H. Spanier. Semigroups, Presburger formulas, and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966. doi: 10.2140/pjm.1966.16.285. URL <https://msp.org/pjm/1966/16-2/pjm-v16-n2-p09-s.pdf>.
- [9] John Haslegrave. An evasion game on a graph. *Discrete Mathematics*, 314:1–5, 2014. doi: 10.1016/j.disc.2013.09.004. URL <https://doi.org/10.1016/j.disc.2013.09.004>.
- [10] Dmitry Kamenetsky. Entry A301426: Number of steps required in the worst case for three knights to find the princess in a castle with n rooms arranged in a line. The On-Line Encyclopedia of Integer Sequences, March 2018. URL <https://oeis.org/A301426>. Created 21 March 2018; conjectural formula added the same day. Accessed 11 September 2026.
- [11] Dmitry Kamenetsky. Entry A301337: Number of steps required in the worst case for two knights to find the princess in a castle with n rooms arranged in a line. The On-Line Encyclopedia of Integer Sequences, March 2018. URL <https://oeis.org/A301337>. Created 19 March 2018; conjectural formula recorded the same day. Accessed 11 September 2026.
- [12] János Körner and Victor K. Wei. Odd and even Hamming spheres also have minimum boundary. *Discrete Mathematics*, 51(2):147–165, 1984. doi: 10.1016/0012-365X(84)90068-2. URL [https://doi.org/10.1016/0012-365X\(84\)90068-2](https://doi.org/10.1016/0012-365X(84)90068-2).
- [13] Yota Otachi and Ryohei Suda. Bandwidth and pathwidth of three-dimensional grids. *Discrete Mathematics*, 311(10–11):881–887, 2011. doi: 10.1016/j.disc.2011.02.019. URL <https://arxiv.org/abs/1101.0964>.
- [14] Angel Raychev. Optimal strategies for finding a princess in a linear table. Bulgarian spring-conference manuscript, twentieth Student Section (US’20), 17 pages. Title translated into English, 2020.
- [15] Angel Raychev and Dimitar Rusev. Optimal strategies for finding a princess in a linear graph. Bulgarian student-conference manuscript, twentieth Student Conference (UK’20), 20 pages. Title translated into English, 2020.