Capture time on rectangular grids with a constant inspection budget

Angel Ivanov Raychev

September 2026
Research manuscript in preparation

Abstract

An invisible target starts in an unknown room of a rectangular grid. Each day at most k rooms are inspected simultaneously; after a miss the target must traverse exactly one edge. We study $T_k(a, b)$, the minimum worst-case capture time, including its infinite cases, and optimal inspections. The fixed research goal is an expression using an absolute $O(1)$ number of standard numerical operations, independent of a, b, k, with directly specified attaining inspections. This goal remains open. We give fixed numerical formulas at every budget for rectangles with at most five rows and for seven-row rectangles of odd length, above stated quadratic budget thresholds at every width and side parity, and on the even-width curve $w = 2r$, $k = r(r-1)$. Fixed-block evaluations provide further absolute-operation regions for odd solo clocks and even-width numerical lower and upper bounds; agreement of the latter gives additional exact subfamilies. Exact compatible profiles reduce odd rectangles to joint count evaluation. For $w = 2b + 1$, the solo midpoint criterion is now exact on every odd rectangle at feasible budgets $k \geq b + 1$ satisfying $k \leq 2b$ or $27k^3 \leq b^4$, in addition to the earlier large-budget and eventual regions. An ordinary ancestry argument and one complete finite complement establish these value and strategy results. Width-dependent solo clocks, including the minimum-budget width constant, still leave numerical evaluation gaps outside the fixed-operation domains. For even-area rectangles, joint boundary restrictions, corner and erosion frontiers, and persistent geometric information strengthen lower bounds without assuming that model transitions are physically attainable. They give $T_{13}(25, n) = 50n - 925$ for every even $n \geq 26$, and a one-day numerical bracket at width 24 and even length. Physical constructions and symbolic interior transitions give further bounds and exact algorithms, whose remaining boundary optimization is not uniformly evaluated. Inverse fixed-deadline consequences remain part of the same problem. Paths and two-row grids have complete physical-game Lean classifications; the new ordinary proofs and finite computational certificates are distinguished from formalized results. A parked companion preserves the higher-dimensional applications. The account reconstructs the author's October 2019–March 2020 path work and documents subsequent AI-assisted development.

Contents

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. Our active objective concerns two-dimensional rectangles only. 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 The fixed completion criterion

For every $a, b \geq 1$ and constant daily budget $k \geq 0$, the numerical research goal is a fixed-size closed-form expression for $T_k(a, b)$, including its infinite cases. Evaluation must use an absolute $O(1)$ number of standard numerical operations, independent of a, b, k. Arithmetic, powers, roots, floors, ceilings and a fixed finite case list are permitted. This is an operation-count goal, not constant bit complexity. Parameter-length sums or products, recurrences, iterative arrays, shortest paths and parameter-dependent preprocessing do not satisfy it. Renaming an algorithm as one operation does not change this. We do not claim that the goal has been attained or is necessarily attainable.

A separate requirement is a directly specified optimal sequence $S_t(a, b, k)$ when the answer is finite, with proofs of capture and of minimality against every competing inspection sequence. Printing the sequence need not take $O(1)$ time. Strategy existence, reconstruction by optimization, and direct inspection rules are distinguished below. The graph notation $T_m(G)$ and the local width/length letters in older proofs are retained, with $T_k(a, b) = T_k(P_a \square P_b)$ and $w = \min(a, b)$, $n = \max(a, b)$.

Relevant exact evaluators, lower and upper bounds, counterexamples, partial-state arguments and general proof tools remain in this paper. Their inclusion records progress; it does not lower the completion criterion. Independent higher-dimensional classifications and applications are preserved at `https://angelraychev.com/princess/extensions/` with a separate manuscript, sources, formalization map and parked restart notes. Scope selection is not mathematical unification, and this separation establishes no new cases.

1.2 Current rectangle coverage

Feasibility already has a bounded numerical criterion:

$$T_k(a, b) < \infty \quad \Longleftrightarrow \quad k \geq \lfloor \min(a, b)/2 \rfloor + 1.$$

For both sides at least two this is the rectangular-grid theorem of Abramovskaya et al. [1, Theorem 2]; paths and the one-room case are covered separately below. Thus the infinite-value cases do not remain an open part of the rectangle classification.

Fixed numerical formulas and attaining constructions are available for all budgets on rectangles with at most five rows, and above the stated quadratic budget thresholds at every width and side parity. Every seven-row odd-length budget is now covered (Theorem 30.2), as is the even-width budget curve $k = r(r-1)$ on width $2r$, $r \geq 3$, at every length (Theorem 31.4). The finite constants in the minimum-budget theorem also give fixed formulas at its listed widths through fifteen. These are proved ranges, not a complete arbitrary-width classification.

Every odd rectangle has exact compatible neighborhood profiles and an exact joint evaluator. Theorem 6.1 places its time in $\{2\tau - 1, 2\tau\}$; the bounded joint frontier in Theorem 7.1 decides which value occurs. The solo-residual rule suffices beyond the explicit threshold in Theorem 7.5 and at every length in the additional ranges of Theorem 30.5 and Corollary 30.6. In general these recurrences still require reduction to bounded numerical expressions. Even the all-width minimum-budget identity $T_{b+1}(2b+1, n) = 2(2b+1)n - C_{2b+1}$, for odd $n \geq 2b+1$, retains a width-only $O(b)$ clock in C_{2b+1} (Theorem 21.1). It is an exact general theorem, but not numerical completion when b varies.

For even-area rectangles, exact static neighborhood profiles and the physical interface theorem 27.1 retain the actual boundary choices. The symbolic transition theorem 13.3 evaluates the interior connections. Parameter-dependent boundary preprocessing and shortest-path optimization remain. The constants hidden in $O_{w,k}(1)$ are not an absolute $O(1)$ bound independent of all three parameters.

Uniform bounds also narrow a broad even-area region to one day: Theorem 29.19 covers $k \geq \lceil 4w/3 \rceil$, and Corollary 29.21 improves the threshold to $\lceil 13w/10 \rceil$ for even widths. The physical upper strategies are explicit, while choosing the optimal endpoint and evaluating every underlying clock in fixed size remain separate obligations.

The common corner-capacity and propagation arguments remain as tools for these gaps. The target is the optimum from the full rectangle. Universal optimality from every partial state, or classification of every legal boundary transition, is not an additional completion condition.

1.3 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 $n > m \geq 2$,

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

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

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

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

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

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

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

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

$$T_3(P_5 \square P_n) = 10n - 20. \tag{5}$$

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 $2\lceil (5n-9)/3 \rceil + 2\mathbf{1}_{\{3|n\}}$ 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 $m \geq b^2+1$.

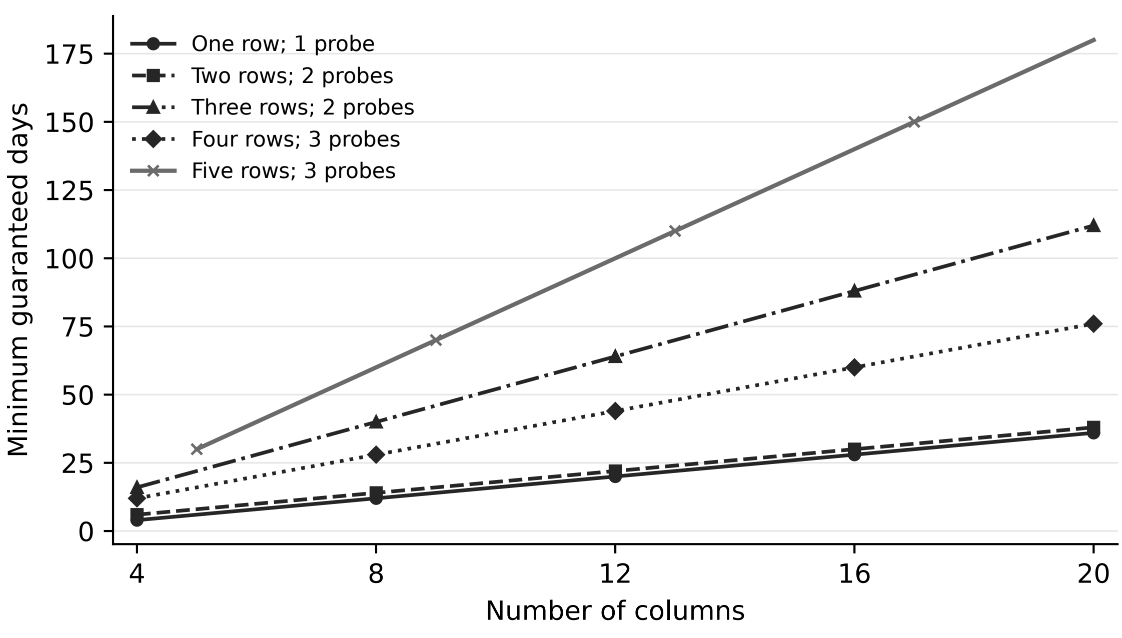

Figure 1: Optimal time at the smallest feasible daily budget for four narrow families and the five-row three-probe theorem. Each line is a proved formula, not a fit to search data.

1.4 Shared proof tools and proof status

The common difficulty is temporal. A smallest next neighborhood need not contain an equally favorable survivor for the following day. Whole- strategy compression, attained inverse-profile arithmetic, corner capacities, physical propagation and efficient-frontier arguments address different parts of that obstruction. General statements are retained where they are natural dependencies of rectangle proofs; their independent box applications are in the parked companion.

The exact fixed-deadline minimum *constant* budget is an inverse formulation of the same capture relation. Section 28 therefore retains its rectangle consequence and proof. It is not silently turned into a varying-budget objective.

Ordinary proofs, independent finite checks and physical-game Lean verification are different evidence layers. Section 32 states the precise boundary. A compiled collection of reusable components does not formally verify every unbounded theorem appearing here.

1.5 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 [13, 14]. 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. The parked companion separately records the classical closed-neighborhood isoperimetry used for its higher-dimensional applications. Those applications do not supply 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 32. 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 $t = 1, 2, \ldots$, 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. For positive rectangle side lengths define

$$T_k(a, b) := T_k(P_a \,\square\, P_b) = T_k(b, a).$$

The daily inspection set is $S_t(a, b, k)$, with $|S_t(a, b, k)| \leq k$. We use k for the global budget and retain the legacy letter m in individual proofs; locally $w = \min(a, b)$ and $n = \max(a, b)$. The separate no-inspection convention is $T_0(a, b) = \infty$, including the one-room board: inability to move is not capture. This convention is not derived from a vacuous target-walk quantifier. 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 $(S_t)_{t\geq1}$. 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

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

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 $v_1, \ldots, v_t = v$ 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. □

Since G has no isolated vertices, a nonempty set has a nonempty open neighborhood. Hence the capture condition is also equivalent to $B_{t+1} = \varnothing$. 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:

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

Then

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

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 $C : 2^V \to 2^V$ 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 $m_1, \ldots, m_T$ 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, $C(R_t) = \varnothing$, 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. □

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. □

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*

$$0, 2, \dots, 2q \qquad \textit{or} \qquad 1, 3, \dots, 2q-1.$$

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

For the rectangle objective take $H = P_w$ below. The prefix-fiber argument is retained in its natural generality as an exact evaluator; independent non-rectangle applications are parked in the companion.

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

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

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

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

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

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

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

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 $a \to F(b)$ whenever $b \le a$ and $d(a,b) \le m$. Then $T_m(a)$ is exactly the shortest-path distance to zero in this state graph, or infinity if zero is unreachable. Equivalently,

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

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

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

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

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

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. □

Remark 3.6 (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 $x+y = p, p+2, \ldots, p+2(r-1)$, and put $h = 2r+p$. For every integer $s \geq 0$, the largest possible support with neighborhood surplus at most s has cardinality*

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

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 $\left\lceil \sqrt{k} \right\rceil$ 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

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

List the occupied diagonals as $j_1 < \cdots < j_L$ 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

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

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

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

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

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

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

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 $a \ge b \ge 1$, with coordinates (x, y); thus x is the long coordinate. Write V_0, V_1 for its even and odd checkerboard classes, and

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

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

$$R(k) = \left\lceil \sqrt{k} \right\rceil, \qquad Q(k) = \min\{r \ge 1 : k \le r(r - 1)\}.$$

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,*

$$g_0(k) = k + \min\{R(k), b, R(E - k) - 1\}, \qquad 1 \le k \le E, \tag{15}$$
$$g_1(k) = k + \min\{Q(k), b + 1, Q(M - k)\}, \qquad 1 \le k \le M. \tag{16}$$

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.

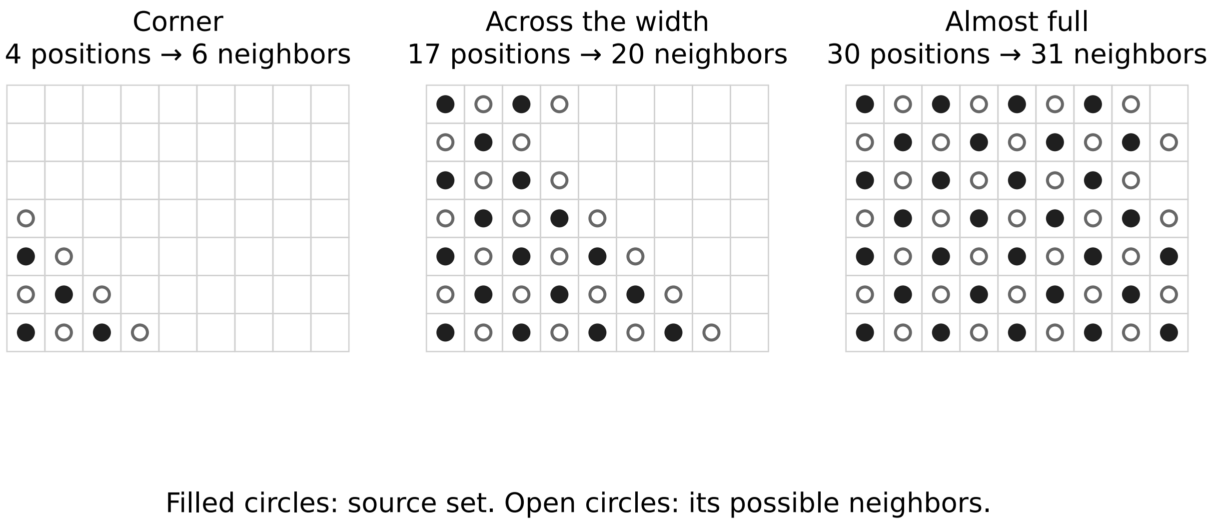

Figure 2: Minimizing even prefixes on a 9×7 rectangle. The three regimes are controlled by a corner, the short width, and the complement.

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*

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

If S has color 1, then

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

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,

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

For color 1 the identity is

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

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

$$\delta_{\text{quadrant}} = \delta + F + Z. \tag{19}$$

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

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

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 $\delta \geq 1$,

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

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

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

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

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

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

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

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

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

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:*

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

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. □

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. Theorems below also settle the all-length linear and superlinear budget ranges; only the shorter boards outside the proved ranges remain open.

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

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

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,*

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

Moreover $T_m = 2\tau - 1$ whenever

$$E + M - D_0(\tau-1) - D_1(\tau-1) \le m. \tag{24}$$

Necessity follows after the explicit onset in Theorem 7.5, as well as in the all-length minimum-budget and high-budget ranges and the additional ranges of Theorem 30.5 and Corollary 30.6. It remains open outside the proved 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

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

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; by itself it does not decide the scalar midpoint condition in the remaining unresolved parameter range.*

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

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

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*

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

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. □

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*

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

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. □

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. The later Theorem 30.5 and Corollary 30.6 also establish scalar necessity at every length in their budget ranges. The shorter boards outside the proved ranges remain 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:

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

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

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

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*

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

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 $c' \geq c$,

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

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

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

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 $H - 1 + 2m(\lfloor \Gamma/w \rfloor + 1) = K$. □

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

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

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

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

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

$$(j, k, \ell) \longmapsto (1-j, h_j(k+r), \ell + r), \qquad 0 \leq r \leq m. \tag{29}$$

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

$$(D_0', D_1') = (D_1 + m - b, D_0 + m - b - 1), \qquad D_p(t+2) = D_p(t) + d.$$

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. □

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 $M = (wn - 1)/2 \geq M_*$,

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

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 $0 \le s \le W+1$ the maximum is at most $J+\Gamma+m-1+ms \le M-J$, 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 $2K+m < J \le B$, 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 $m \ge b^2+1$ (Theorems 21.1 and 8.1) leave only intermediate budgets $b+2 \le m \le b^2$ 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

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

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 $v' = h_p(v+r)$, $0 \le r \le m$, 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*

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

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 $\mu_p(k) \le mt$; 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

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

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

Proof. Feasibility follows since $h_0(z+m) \ge z+1$ and $h_1(z+m) \ge z$. 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 $h_p(z) \ge k$. A last-step predecessor must be at least $v_0 = (z-m)_+$. For $v_0 \le v \le z$, monotone excess gives

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

For $v \ge z$ 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,

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

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

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

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. □

Remark 7.7 (Cyclic growth controls). The proof requires no special root identities. Let the phases form any nonempty finite cycle, let $m \ge 1$, and let each $h_p : \mathbb{N} \to \mathbb{N}$ 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

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

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 $i \le t$. 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 $z \ge H$, the lower root in I_p has reached its cap and the reflected root is nonincreasing. Hence $I_p(z) - z$ is nondecreasing, so

$$I_p(v) - I_p(u) \ge v - u \qquad (v \ge u \ge H).$$

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 $Q \ge \mu$ 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 $\Delta_s \ge 0$ 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 $m \ge b + 1$ 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. □

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

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

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 $k > 0$,

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

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

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

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

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*

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

Then

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

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 $k > 0$,

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

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

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

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 $0 \leq p \leq m$,*

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

Put $j(u) = \left\lfloor (u-w-2)_+ /D \right\rfloor$. 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

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

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. ☐

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. ☐

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

$$d' \le d + p. \tag{38}$$

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

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

and the monotone potential

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

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

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

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

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

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 $x = d+p \le B$, and the defining minimum gives $C - d \le p + F(\text{next rank})$. Thus the source is at most p plus each target branch. At capture, $k \le p \le m$ and $\Phi_z(k) \le F(2k+z) = k \le p$. 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. □

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

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

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

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

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

For $x \le B$, 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 $J(d_E(x)) \le x$ and $J(d_O(x)) \le x$ 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 $C \leq F(N)$. □

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,

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

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 $2J(r)+\delta(r) \le m$. Otherwise either full phase has rank at most $(q+1)D+w+1$, because $r+1 \le D$, and two $(q+1)$-day sweeps give the other upper bound. Early capture during a sweep only reduces its actual inspection costs.

For $m \ge N$ one inspection of the entire board suffices. For $h \le m < N$, 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 $m \ge B+b$, Lemmas 8.5 and 8.6 give the lower bound $\lceil 2F(N)/m \rceil$. It matches the construction except when

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

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

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

For the latter inequality, $m - r \ge 4$ for even r and $m - r \ge 3$ 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 $x \le B$,

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

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 $a \ge c-1$. Otherwise the bounds $k \ge (c-1)(c-2)+1$ and $x \le a^2$ 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 $a \ge 2$, $x \le a(a-1)$, and $k \ge (c-1)(c-2)+1$; 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 $p \le B$ 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 $x = d + p \le B$. Its output equals that of a full-cohort x-inspection departure. Its source quotient is either q or $q-1$, since $D \ge 2B-1$ and $r \ge 2$. If both source and target quotients equal q, the strict corner inequalities give $p + F(\text{target}) > C - d \ge \Phi(\text{source})$. If the source quotient is $q-1$, then

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

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/even, O/even, E/odd, O/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 $x + F(\text{target}) \ge C$, 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 $q \ge 1$ 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 $m \ge 3$ the same equality argument applies; a critical remainder has $m \ge 4$, $r = m - 3 \ge 1$, 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 $m \ge B + b$.

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

It remains to add $B + 1 \le m < B + b$. 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 $C \ge 5q + 3$, so the initial potential $2C \ge 10q + 6$ 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 $b \ge 3$. The minimum square size and $D \le 2B - 3$ give $q \ge 2$. 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

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

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 $p \ge m-1 \ge B$. 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 $n \ge w$ and $p \le m < B+b$. Here $r \le m-3 < D-3$ and $r \ge 2$, 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 $m \ge B+1$.

8.6 Removing the root functions

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

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

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 $2 \le w \le n$ 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 $\rho(d) = \min\{s \ge 0 : d \le s(s+1)\}$.

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

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

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

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

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 $0 \le x < e$, $0 \le y < o$. 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 $|B| = s \ge 0$, and R is outgoing-closed in $D_p - B$. If $s < b$, at most s columns contain B. Since $o \ge 2b \ge 2s + 2$, 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 $e/2 \ge b > s$. An outgoing-closed R either contains C or avoids it.

9.2 A sharp area bound outside the component

Suppose $U \cap B = \varnothing$, $N^+(U) \setminus U \subseteq B$, and every row avoiding B is empty in U. We prove $|U| \le s^2$. 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 $\alpha_i \le \alpha_{i+1} + t_{i+1}$ and hence $\alpha_i \le \sum_{i<j<r} t_j$. Horizontal matching gives $\beta_i \le \alpha_i + t_i$. 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 $q \le \alpha_i + t_i \le \sum_{i \le j < r} t_j \le s$, contrary to $q \ge b > s$. Thus in all cases each occupied row satisfies

$$|U_i| \le w_i + 2 \sum_{i<j<r} w_j. \tag{46}$$

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 $a \le W$. 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 $1, 3, \ldots, 2a - 1$, summing to a^2; every remaining vertex has coefficient at most $2a$. The total is therefore at most $a^2 + 2a(W - a) \le W^2$.

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 $|U| \le W^2 + (W')^2 \le s^2$. 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 $h - |R| = |U| + s \le s(s+1)$. 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 $s = \left\lceil \sqrt{k} \right\rceil < b$, 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 $0 \le x \le s^2$. 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

$$H_b = \{(x, y) : 0 \le x < 2b,\ y \ge 0\}, \qquad b \ge 1,$$

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

$$R(k) = \left\lceil \sqrt{k} \right\rceil, \qquad Q(k) = \min\{s \ge 1 : k \le s(s-1)\}, \qquad C_h(s) = \max\{s(s-1)/2, s(s-h)\}.$$

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	$\min\{R(k), b\}$
$c \notin S$	$\min\{Q(k), b+1\}$
$c' \in N(S)$	$\min\{Q(k), b\}$
$c' \notin N(S)$	$\begin{cases} R(k), & k < b^2, \\ b+1, & k \ge b^2. \end{cases}$

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

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

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

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

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 $a_x \le t-2+|x-v|$, 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 12.2 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 $k \ge b^2$, 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 $b \ge 2$. The joint condition forbids exactly the three source rooms

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

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 $C_3 \le C_2$ 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

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

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 $|S|+1 \le b^2-1$, or $|S| \le b^2-2$. 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 $b \ge 3$, convexity bounds the total by

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

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

$$\alpha_i \le \sum_{j<i} t_j, \qquad \beta_i \le \sum_{j>i} u_j,$$

and therefore

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

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 $b \ge 3$, 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 $c_i \le 2i$. 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 $b \ge 3$; 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 $s \le b$, 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 $b \ge 3$, 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*

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

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

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

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 $a \ge 1$ and has companions of total surplus $d \ge 2$, its total capacity is at most $a^2 + d(d-1) \le (a+d)(a+d-1)$. 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

$$(s-2)(s-3) + 2 \le s(s-2) \le C_2(s), \qquad s \ge 4;$$

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 $0, 2, \ldots, 2L - 2$: 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 $L \le s$. 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 $1, 3, \ldots, 2L - 1$, and the output origin supplies one more surplus unit, so $L + 1 \le s$.

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. □

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.

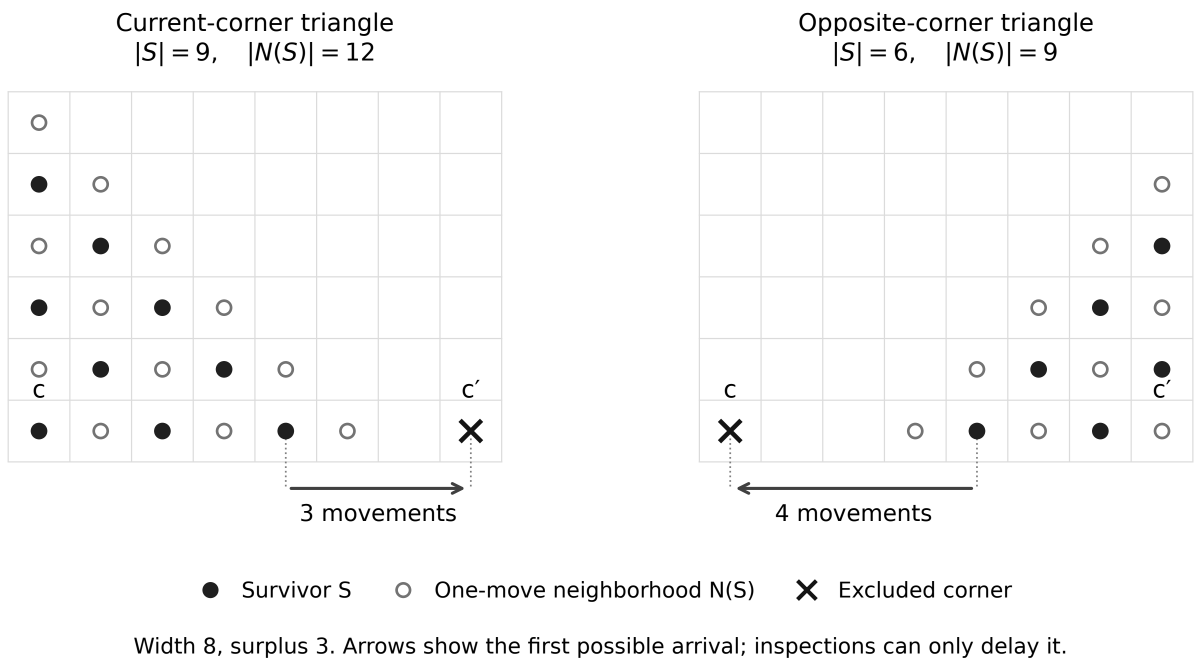

Figure 3: The sharp localization examples for $b = 4, s = 3$. Filled rooms are S; open rooms are $N(S)$. Crosses mark the excluded opposite corner. Its shortest arrival times are three and four movements.

11 Corner geometry and stronger lower bounds

The following argument retains the geometric information lost by a size-only profile, then proves which necessary models can be concentrated into one cohort. A model path need not be physically attainable. Its numerical evaluation therefore supplies a lower bound until a construction matches it.

11.1 A corner-count bound on every even-area rectangle

The preceding size profile can lose the corner information needed to chain two inspection steps. The following necessary bound retains both current and outgoing corner counts, without restricting the shape of the support. Let $G = P_w \,\square\, P_n$, where $2 \le w \le n$ and wn is even, and put $b = \lfloor w/2 \rfloor$ and $h = wn/2$. Each color contains exactly two board corners. For a one-color support S, write

$$a = |S|, \qquad s = |N(S)| - |S|, \qquad i = |S \cap \mathcal{K}_p|, \qquad j = |N(S) \cap \mathcal{K}_{1-p}|,$$

where $\mathcal{K}_p$ consists of the two corners of color p. Neighborhoods are always taken in the entire finite board.

For $0 \le i, j \le 2$, define

$$\Phi_{i,j}(s) = \begin{cases} -\infty, & s < i + 2j, \\ i + 2j - 2 + (s - i - 2j + 2)(s - i - 2j + 1), & s \ge i + 2j,\ j > 0, \\ i - 1 + (s - i + 1)^2, & s \ge i,\ j = 0 < i, \\ \max\{s(s-1)/2, s(s-2)\}, & s \ge 0,\ i = j = 0. \end{cases} \tag{49}$$

Put

$$\Psi_{i,j}(s) = \max_{0 \le d \le 2-i} \Phi_{2-j,d}(s - 2 + i + d). \tag{50}$$

This maximum has at most three candidates, with $d \in \{0, 1, 2\}$; both functions therefore use a fixed number of numerical operations.

Theorem 11.1 (Finite-board corner-count dichotomy)**.** *For every physical support S with $0 \le s < b$,*

$$a \le \Phi_{i,j}(s) \qquad \text{or} \qquad h - a - s \le \Psi_{i,j}(s). \tag{51}$$

No compression or compatibility assumption is required. The statement includes both parities of the shorter side and the empty and full supports.

Proof. Choose an even side e, and let o be the other side. Use the directed matching contraction in Section 9, with $e/2$ matched rows, o columns, contracted support R, and outside boundary B of size s. Since $e/2 \ge b > s$, some matched row avoids B. Since $o \ge 2b \ge 2s + 2$, two consecutive columns avoid B. Those two columns form a strongly connected ladder. Every undeleted matched row and every adjacent pair of undeleted columns therefore lie in one strongly connected component C of $D - B$. The outgoing-closed set R contains all of C or none of it.

First suppose R avoids C. Select an undeleted matched row and two consecutive undeleted columns. They contain neither R nor B. Physically they supply two consecutive empty transverse rows and two consecutive empty longitudinal columns. Cutting there splits S into four corner pieces with disjoint neighborhoods. Each neighborhood is unchanged when its piece is viewed in the whole quadrant based at its board corner: the empty separators keep it away from the two opposite physical board edges. Thus the four sizes and surpluses add.

At each source-colored corner present in S, the even quadrant piece has surplus at least one and capacity $C_0(u) = u^2$. At an omitted source-colored corner its capacity is $C_2(u)$. At each opposite-colored corner present in $N(S)$, the odd quadrant piece has surplus at least two and capacity $C_1(u) = u(u-1)$. If that outgoing corner is absent, both odd neighbors of its origin are forbidden, giving capacity $C_3(u)$. These are precisely the punctured-quadrant bounds of Theorem 4.1.

The i required even pieces and j required odd pieces consume surplus at least $i + 2j$. All four capacities are nondecreasing and discretely convex. Subject to their minimum surpluses, the sum is maximized by putting all remaining surplus into one piece: for two variable pieces, convexity moves the maximum to an endpoint of their allocation interval; repeat until one variable remains. Write $z = s - i - 2j$. Above its minimum, a required even piece gains $z^2 + 2z$, and a required odd piece gains $z^2 + 3z$. An optional piece gains at most $C_2(z) \le z^2$, with $C_3(z) \le C_2(z)$. Thus a required odd piece receives the excess when $j > 0$, a required even piece when $j = 0 < i$, and an omitted-even-origin piece when $i = j = 0$. This gives exactly (49) and proves the first branch.

Otherwise put $U = V_{1-p} \setminus N(S)$, so

$$|U| = h - a - s, \qquad N(U) \subseteq V_p \setminus S.$$

In the transposed matching contraction, U is outgoing-closed outside B and avoids C. The same four-quadrant argument applies. Its current corner count is $2 - j$. Write $d = |N(U) \cap \mathcal{K}_p|$ and $t = |N(U)| - |U|$; then $0 \le d \le 2 - i$ and $0 \le t \le s$. The complement slack has size

$$|(V_p \setminus S) \setminus N(U)| = s - t.$$

It contains the $2-i-d$ corners belonging to neither S nor $N(U)$. Consequently $t \le s-2+i+d$. The first-branch argument and monotonicity now give

$$h-a-s \le \Phi_{2-j,d}(t) \le \Phi_{2-j,d}(s-2+i+d) \le \Psi_{i,j}(s),$$

proving the second branch. The corner charge in the middle inequality is essential; merely using $t \le s$ would discard relevant information. $\square$

Corollary 11.2 (A uniform necessary inspection rule)**.** *Suppose a one-color belief A has size K and contains c of its physical-color corners. After p useful inspections, let the next belief have size K' and corner count j. Set $a = K-p$ and $s = K'-a$. If $s < b$, then* (51) *holds for some $i \in \{0,1,2\}$ satisfying $i \le c \le i+p$.*

Proof. The survivor $S \subseteq A$ has a rooms and i current-colored corners; removing p rooms can remove at most p corners. Apply the theorem. $\square$

For example, on 8×8 a 25-room belief with one current-colored corner cannot become a 23-room belief after five inspections. The size-only profile forces the survivor size to be 20 and surplus three. Its source-corner count is at most one. For $i=1$, the three dual capacities, indexed by $j=0,1,2$, are $2,4,6$, whereas $h-a-s=9$; for $i=0$ they are no larger. The small branch also fails. Thus the corner bound excludes this transition even though the static size profile permits it. For surplus $s \ge b$, retain the size-only profile; this section makes no claim that the corner relaxation is sharp there or that its allowed transitions are jointly attainable.

11.2 The critical corner bound on odd-width rectangles

When the shorter side is odd and the longer side is even, the preceding corner dichotomy extends to the critical surplus with one explicit exception. Let $w = 2b+1 \ge 3$, $n = 2\ell \ge w$, and $h = w\ell$. Use the corner counts i, j and capacities Φ, Ψ from Section 11.1.

Theorem 11.3 (Critical corner dichotomy)**.** *Every physical one-color support of size a and surplus b satisfies*

$$a \le \Phi_{i,j}(b) \quad \text{or} \quad h-a-b \le \Psi_{i,j}(b) \quad \text{or} \quad (i,j) = (2,0).$$

Together with Theorem 11.1, this is a uniform necessary corner bound at every surplus at most b. No assertion of attainability or dynamic compatibility is made.

Proof. Reflect to physical color zero and match longitudinal room pairs. The contracted graph has bidirectional transverse edges, downward edges on $U = \{0,2,\ldots,2b\}$, and upward edges on $V = \{1,3,\ldots,2b-1\}$. Let R be the contracted support and H its outside boundary, with $|H| = b$. Since $\ell > b$, a horizontal row avoids H.

If two adjacent columns also avoid H, the proof of Theorem 11.1 applies without change: the two separators lie in one strongly connected component, which R either contains or avoids. Accordingly R or its transposed complement splits into four physical corner quadrants. Their whole neighborhoods agree with the finite-board neighborhoods. Additive capacities and the same missing-corner slack give the first or second displayed alternative. That argument needs the two separators, not the strict surplus bound.

Otherwise the marked columns form a vertex cover of a path on $2b+1$ positions using at most b columns. They are exactly V, with one mark (v, c_v) each. Every U fiber is a prefix of length t_u, and horizontal closure gives, for adjacent uv,

$$R_v \subseteq R_u, \qquad R_u \setminus R_v \subseteq \{c_v\}.$$

Thus neighboring U lengths differ by at most one. For consecutive U columns u, u' surrounding v, one also has

$$|R_v| \leq \min\{t_u, t_{u'}\}, \qquad \ell - |R_v| \leq \min\{\ell - t_u, \ell - t_{u'}\} + 1.$$

These inequalities do not assume that R_v is a prefix.

An outgoing corner is present exactly when its endpoint U fiber is full, because the neighborhood adds no U vertices. A full U fiber forces every U fiber to be nonempty: its length drops by at most one across each of b intervening positions, and $\ell > b$. Consequently $j > 0$ implies $i = 2$. There are five remaining possibilities. If $(i, j) = (0, 0)$, both endpoint U fibers are empty, so $t_{2q} \leq \min\{q, b-q\}$; the first V bound sums to $a \leq b(b-1)/2 \leq \Phi_{0,0}(b)$. If $(i, j) = (1, 0)$, the one-sided distance bound instead gives $a \leq b^2 = \Phi_{1,0}(b)$. The pair $(2, 0)$ is precisely the stated exception. If $(i, j) = (2, 1)$, one endpoint U fiber is full; applying the second V bound to its deficiencies gives $h - a \leq b(b+1)$, hence $h - a - b \leq b^2 = \Psi_{2,1}(b)$. Finally, for $(i, j) = (2, 2)$ both endpoint deficiencies vanish, giving

$$h - a \leq b(b+1)/2, \qquad h - a - b \leq b(b-1)/2 \leq \Phi_{0,0}(b) = \Psi_{2,2}(b).$$

These alternatives exhaust the cases and prove the theorem. □

Corollary 11.4 (Rigidity at the pronic capacity)**.** *Suppose $b \geq 2$. A survivor S with*

$$|S| = b(b-1), \qquad |N(S)| - |S| = b,$$

and no corner of its own color is the full odd triangle based at one opposite-colored board corner: its relative diagonals are $1, 3, \ldots, 2b-3$. Consequently $N(S)$ contains exactly that one opposite-colored corner, and $N^2(S)$ contains neither original-colored corner. The latter omission persists after any intervening inspections.

Proof. Take radius $t = b$ in Theorem 11.12. Since the shorter side is odd, its outgoing allowance is $\kappa = 0$. □

Corollary 11.5 (Critical orientation down to the pronic threshold)**.** *Put $B = b(b+1)$ and $C = b(b-1)$. If a survivor S has surplus b and $C < |S| < h - B$, it contains at least one own-colored corner, while $N(S)$ contains neither opposite-colored corner. Two successive survivors in this interval cannot both have surplus b. Moreover, after such a survivor, every next survivor A of size $a \leq b^2$ satisfies*

$$|N(A)| \geq a + \min\{Q(a), b+1\}, \qquad Q(a) = \min\{q \geq 1 : a \leq q(q-1)\} \quad (a > 0), \quad Q(0) = 0.$$

Proof. Every $\Psi_{i,j}(b) \leq b^2$, so the upper size restriction excludes the dual branch of Theorem 11.3. If $i = 0$ or $j > 0$, then $\Phi_{i,j}(b) \leq C$, excluding its small branch. Its exceptional pair $(2, 0)$ already has the required corners. The next survivor omits both of its own corners, so the reflected conclusion proves incompatibility.

For its small-corner bound, suppose $a \leq b^2$ and its surplus is $s \leq b$. The critical exception is unavailable since its current corner count is zero. Also $h - a - s \geq h - b^2 - b > b^2 \geq \Psi_{0,j}(s)$, so the dual branch is excluded. The first branch gives $a \leq \Phi_{0,j}(s) \leq s(s-1)$, using the subcritical theorem when $s < b$. Therefore either $s \geq b+1$ or $s \geq Q(a)$, as required. □

Corollary 11.6 (The dual critical corner restriction)**.** *If $|N(S)| - |S| = b$ and $b^2 < |S| < h - b^2$, then S contains both corners of its own color and $N(S)$ contains at most one opposite-colored corner.*

Proof. The small branch of Theorem 11.3 is impossible since every $\Phi_{i,j}(b) \le b^2 < |S|$. Its exceptional pair $(2,0)$ already has the asserted corners. Otherwise the complementary size is $h - |S| - b > b(b-1)$. Every $\Psi_{i,j}(b)$ with $i < 2$ or $j = 2$ is at most $b(b-1)$, so the dual branch forces $i = 2$ and $j \le 1$. □

11.3 The critical corner bound on even-width rectangles

The critical surplus has two possible spanning orientations. The even one occurs on every even-width rectangle; the odd one occurs additionally on a rectangle whose sides differ by one. These are necessary exceptions, not a sufficiency assertion about the corresponding profiles.

Theorem 11.7 (Critical corner bound for even widths)**.** *Let $w = 2b \ge 2$, $n \ge 2b$, and $h = bn$. A physical one-color support of size a, surplus b, and corner counts i, j satisfies*

$$a \le \Phi_{i,j}(b) \quad \text{or} \quad h - a - b \le \Psi_{i,j}(b)$$

unless $(i,j) = (1,1)$, or $n = 2b+1$ and $(i,j) = (2,0)$.

Proof. Match adjacent rooms across the $2b$ side. The contraction has b rows, n columns, horizontal bidirectional edges, and alternating downward and upward rungs. Reflect to make even columns downward. Write R for the contracted support and B for its outside boundary, with $|B| = b$. Then R is disjoint from B and outgoing-closed in $D - B$.

An unmarked adjacent pair. Two adjacent columns avoiding B form a strongly connected ladder; no unmarked matched row is needed for this assertion. If R avoids this ladder, the two empty physical columns split its support into two pieces whose neighborhoods are disjoint and unchanged in the whole half-strips of width $2b$ based at the board's two ends. If both pieces are nonempty, each surplus is less than b. Each piece then splits across an empty matched row into two quadrants, and the capacity allocation in Theorem 11.1 gives Φ. If only one is nonempty, Theorem 10.1 bounds its critical capacity by $C_2(b)$, $b(b-1)$, or $b^2 - 1$ for corner bits $(0,0)$, $(0,1)$, or $(1,0)$, respectively. These fit Φ; the remaining pair $(1,1)$ is the stated exception.

If R contains the ladder, use $U = V_{1-p} \setminus N(S)$ in the transposed contraction. Its surplus t is at most b. The same first-branch argument applies to U, except possibly at its critical pair $(1,1)$. Its outgoing count d satisfies $t \le b - 2 + i + d$ by the missing-corner slack, giving Ψ as before. If $t = b$ and U has pair $(1,1)$, zero slack forces $i = j = 1$ for S. This proves the theorem whenever an adjacent pair is unmarked, and hence for $n \ge 2b+2$.

The near-square $n = 2b+1$. If no adjacent pair is unmarked, the b marked columns are exactly the odd columns, each with one mark. Clean even fibers are prefixes of lengths $t_0, \dots, t_b \in \{0, \dots, b\}$. For a marked fiber R_v between clean fibers $q, q+1$, horizontal closure gives

$$|t_{q+1} - t_q| \le 1, \qquad |R_v| \le \min(t_q, t_{q+1}), \qquad b - |R_v| \le \min(b - t_q, b - t_{q+1}) + 1. \tag{52}$$

Current corners correspond to nonempty endpoint clean fibers, and outgoing corners to full ones. The distance sums in the proof of Theorem 11.3 give the same bounds here: $i = j = 0$ gives $a \le b(b-1)/2 \le \Phi_{0,0}(b)$; $(i,j) = (1,0)$ gives $a \le b^2 = \Phi_{1,0}(b)$; $(i,j) = (2,1)$ gives $h - a - b \le b^2 = \Psi_{2,1}(b)$; and $i = j = 2$ gives $h - a - b \le b(b-1)/2 \le \Psi_{2,2}(b)$. The pair $(2,0)$ is allowed. The only extra possibility at matched height b is one full endpoint and one empty endpoint, giving $(1, 1)$. Thus this case also follows.

The square $n = 2b$. Again assume that no adjacent pair is unmarked. The b unmarked columns form a maximum independent set of the $2b$-vertex path. They are

$$0, 2, \dots, 2\ell - 2; \quad 2\ell + 1, 2\ell + 3, \dots, 2b - 1 \qquad (0 \le \ell \le b). \tag{53}$$

Indeed b nonadjacent positions need span at least $2b - 2$, leaving one unit of slack, at an end or at one interior gap. Every marked column has exactly one boundary vertex. Clean even fibers are prefixes and clean odd fibers suffixes; the bounds (52) apply across each single marked column.

For the pure pattern $\ell = b$, write $t_0, \dots, t_{b-1}$ for its prefix lengths. The final odd column is marked and has just one clean neighbor. Its fiber is contained in that neighbor and differs by at most one boundary vertex. In particular $t_{b-1} > 0$: otherwise the final column has no source neighbor and cannot contain an actual boundary mark. Its bottom outgoing corner is consequently always present. The left current corner is present exactly when $t_0 > 0$, the left outgoing corner exactly when $t_0 = b$. The right current corner requires $t_{b-1} = b$, which forces $t_0 \ge 1$.

If $i = 0$, then $t_0 = 0$ and $j = 1$. Summing $t_q \le q$ and the marked bounds gives $a \le b(b-1) = \Phi_{0,1}(b)$. The pair $(1, 1)$ is allowed. For $(1, 2)$ or $(2, 1)$, a clean endpoint is full. The deficiencies are bounded by their distances from that endpoint; the marked deficiencies have one extra unit each. Their sum is at most b^2, so $h - a - b \le b(b-1) = \Psi_{1,2}(b) \le \Psi_{2,1}(b)$. For $i = j = 2$, both endpoints are full and $d_q = b - t_q \le \min(q, b - 1 - q)$. Hence

$$h - a - b \le \sum_q d_q + \sum_{q=0}^{b-2} \min(d_q, d_{q+1}) \le (b-1)(b-2)/2 \le C_2(b) = \Psi_{2,2}(b).$$

This includes $b = 1$; reflection handles $\ell = 0$.

It remains to handle $1 \le \ell \le b - 1$ in (53). Put $d = b - \ell$. There are ℓ clean prefixes on the left and d clean suffixes on the right. Let the two central clean fibers be $P = [0, L - 1]$ and $Q = [b - R, b - 1]$. Between them are two marked fibers A, Z, with respective marks α, β. Horizontal closure gives

$$A \subseteq P, \quad P \setminus A \subseteq \{\alpha\}, \quad Z \subseteq Q, \quad Q \setminus Z \subseteq \{\beta\}, \quad A \setminus Z \subseteq \{\beta\}, \quad Z \setminus A \subseteq \{\alpha\}.$$

Thus $P \triangle Q \subseteq \{\alpha, \beta\}$, and

$$L + R \le 2 \quad \text{or} \quad L + R \ge 2b - 2. \tag{54}$$

The unit slopes from the center and the marked-size bounds give

$$a \le 2\ell L + 2dR + (\ell - 1)^2 + (d - 1)^2. \tag{55}$$

For example, the clean left fibers contribute at most $\ell L + \ell(\ell - 1)/2$, their internal marked fibers at most $(\ell - 1)L + (\ell - 1)(\ell - 2)/2$, and the central marked fiber at most L. The right calculation is identical.

First suppose $L + R \le 2$. Equation (55) is at most b^2. If an outgoing corner is present, an endpoint clean fiber is full. On the left this requires $b \le L + \ell - 1$ and hence $d = 1, L = 2, R = 0$; the other endpoint is empty. This is the allowed pair $(1, 1)$, and reflection handles the right side. Therefore assume $j = 0$. For $i = 0$, bound each half from its empty outer endpoint instead:

$$a \le \ell(\ell - 1) + d(d - 1) \le (b-1)(b-2) \le C_2(b).$$

For $i = 1$, the bound $a \le b^2$ suffices. For $i = 2$ the required bound is $(b-1)^2 + 1$. If $L + R \le 1$, it follows from (55). If $(L, R) = (2, 0)$, the nonempty right endpoint requires $d \ge 2$, and the same equation is at most $b^2 - 2(d-1)(\ell+1) \le (b-1)^2 + 1$; reflect for $(0, 2)$. If $L = R = 1$ and $b \ge 3$, both central marked fibers are empty. Indeed $P = \{0\}$ and $Q = \{b-1\}$; the horizontal relations permit either both marked fibers empty or both singletons. The latter forces $\alpha = b-1, \beta = 0$, but the upward rung out of $A = \{0\}$ then requires $1 = \alpha$, contradicting $b \ge 3$. Removing the two central estimates gives $a \le \ell^2 + d^2 \le (b-1)^2 + 1$. For the sole endpoint $b = 2$, $\ell = d = L = R = 1$, empty central fibers give $a = 2 = \Phi_{2,0}(2)$; two singletons give $a = 4$ and $h - a - b = 2 = \Psi_{2,0}(2)$.

Finally consider the other branch of (54), after disposing of the first: $L + R > 2$ and $L + R \ge 2b-2$. For $U = V_{1-p} \backslash N(S)$ in the transpose, the clean central lengths after reflection are $b-L, b-R$, totaling at most two. It is closed outside B, so the area estimate alone gives $|U| \le b^2$. If its true surplus $t < b$, its complementary branch in Theorem 11.1 is impossible because

$$h - |U| - t \ge b^2 - b + 1 > (b-1)^2 \ge \Psi_{i',j'}(t).$$

Its small branch and the missing-corner slack give the desired Ψ bound for S. If $t = b$, its true boundary equals B, and the preceding small-center proof applies to U. A small branch gives Ψ for S; its exception $(1, 1)$ transfers to the same pair by zero slack. The only small-center case that instead used a complementary branch had $b = 2$ and central total two. It cannot occur here, since $b = 2, L + R > 2$ makes the complement's central total strictly less than two. There is therefore no circular appeal to the square theorem. This finishes every pattern. ☐

Corollary 11.8 (A symmetric critical corner rule). *On every nondegenerate even-area rectangle put $b = \lfloor \min(w, n)/2 \rfloor$. The Φ/Ψ dichotomy holds at surplus b except possibly for $(i, j) = (1, 1)$ when a side equals $2b$, or $(i, j) = (2, 0)$ when a side equals $2b + 1$.*

Proof. Combine Theorems 11.3 and 11.7. The two exceptional conditions may both apply only to the near-square $2b \times (2b + 1)$. ☐

11.4 Retaining the corners of an erosion

Let an even-area rectangle have $4 \le w \le n$, and put $b = \lfloor w/2 \rfloor$, $h = wn/2$. For a one-color support A, let $c(A)$ count its own-colored board corners and let $e(A)$ count the opposite-colored corners whose two neighbors both belong to A. Thus $e(A)$ is the corner count of $E(A) = \{v : N(v) \subseteq A\}$ in the opposite color. Keeping this second count distinguishes reaching a corner from retaining both of its neighbors. It yields necessary rules for arbitrary physical beliefs.

Lemma 11.9 (One forbidden neighbor of a quadrant origin). *For odd-colored supports in the whole nonnegative quadrant with one specified neighbor of the origin forbidden, the maximum cardinality at neighborhood surplus at most s is*

$$H(0) = 0, \qquad H(s) = \max\{s(s-1)/2,\ s(s-3) + 1\} \quad (s \ge 1).$$

This asserts attainment of the maximum, not every intervening cardinality.

Proof. Compress each odd diagonal toward the permitted neighbor. This preserves the prohibition and cardinality and cannot increase the neighborhood. Use the occupied-layer notation of the punctured-quadrant proof, with surplus δ, L occupied diagonals, and f_t full diagonals among the first t. The same layer bound is $a_{j_t} \le \delta - L + t - 1 + f_t$. If the first diagonal is empty, $C_3(\delta) \le H(\delta)$

applies. If it has one room and no diagonal is full, the no-full-layer bound is $\delta(\delta-1)/2$. Otherwise the first full diagonal has occupied rank $u \ge 2$ and length at least $2u$, while the layer bound gives length at most $\delta - L + u$. Hence $L \le \delta - 2$. For $t \ge 2$ use $f_t \le t-1$ and retain the first layer's exact size one:

$$|A| \le 1 + \sum_{t=2}^{L}(\delta - L + 2t - 2) = 1 + \delta(L-1) \le 1 + \delta(\delta - 3).$$

Monotonicity of H permits $\delta \le s$. Proper initial diagonal counts $1, 2, \ldots, s-1$ attain the first branch. For $s \ge 5$, a full odd triangle of surplus $s-1$, with the forbidden origin neighbor deleted, retains its neighborhood and has size $(s-1)(s-2)-1 = s(s-3)+1$. The empty endpoints are immediate. □

For $0 \le i \le 2$ and $0 \le u \le j \le 2$, put $z = s - i - 2j$ and define

$$\mathcal{P}_{i,u,j}(s) = \begin{cases} -\infty, & z < 0, \\ i + j + u + z^2 + 3z, & z \ge 0,\ u > 0, \\ i + j + z^2 + 2z, & z \ge 0,\ u = 0 < i, \\ j - 1 + H(z+2), & z \ge 0,\ i = u = 0 < j, \\ C_2(s), & s \ge 0,\ i = u = j = 0. \end{cases} \tag{56}$$

It uses a fixed number of numerical operations and satisfies $\mathcal{P}_{i,u,j}(s) \le \Phi_{i,j}(s) \le s^2$ whenever finite.

Theorem 11.10 (Erosion-corner dichotomy)**.** *Let S be one-colored, $B = N(S)$, and set*

$$q = |S|, \quad i = c(S), \quad u = e(S), \qquad t = |B|, \quad j = c(B), \quad v = e(B), \quad s = t - q.$$

Then $u \le j$, $i \le v$, and, for $0 \le s < b$,

$$q \le \mathcal{P}_{i,u,j}(s) \quad \text{or} \quad h - t \le \Phi_{2-j,2-v}(s + i - v). \tag{57}$$

For odd $w = 2b+1$ and even $n \ge w$, this also holds at $s = b$ unless $(i,j) = (2,0)$; that exception necessarily has $u = 0, v = 2$. The refined critical assertion here does not include even w.

Proof. The two feature inequalities follow directly from the definition of a neighborhood. In the small side of the matching separation proof of Theorem 11.1, the i required even-origin pieces consume surplus one and area one each. Of the j touched odd origins, u have both neighbors present and $j-u$ have exactly one. They consume surplus two each, with respective initial areas two and one. The remaining surplus is z. Discrete convexity concentrates it in one piece. The gains for a required two-neighbor odd piece, a required even piece, and a required one-neighbor odd piece are respectively

$$z^2 + 3z, \qquad z^2 + 2z, \qquad H(z+2) - 1 = \max\{z(z+3)/2, z^2 + z - 2\}.$$

The third gain lies between $C_2(z)$ and $z^2 + 2z$; an optional piece gains at most $C_2(z)$. This proves the four branches of (56).

On the other side, put $U = V_{1-p} \setminus B$. The identity $E(B) = V_p \setminus N(U)$ gives the exact counts $c(U) = 2 - j$, $c(N(U)) = 2 - v$. If $\sigma = |N(U)| - |U|$, the slack $(V_p \setminus S) \setminus N(U)$ has size $s - \sigma$ and contains $v - i$ corners. Thus $\sigma \le s + i - v$. The same common matching component and separators put U on the small side. Applying its Φ capacity proves the second alternative.

At critical odd width, the separator case is unchanged. In the alternating-column case of the critical corner proof, the possible corner pairs are $(0,0),(1,0),(2,0),(2,1),(2,2)$. In the first two, $u=0$ and the refined small capacity equals the old one. The middle pair is the stated exception. In the last two, $v=2$ and the fixed dual capacity $\Phi_{2-j,0}(b)$ equals the previous $\Psi_{2,j}(b)$, so the same deficiency bound applies. □

Proposition 11.11 (Corner closure and retention)**.** *In the preceding notation, put $\delta = v-i$ and $\sigma = s-\delta$. Then $\sigma \geq 0$. Whenever σ is in the theorem's proved range, its stronger closed form is*

$$q+\delta \leq \mathcal{P}_{v,u,j}(\sigma) \quad \text{or} \quad h-t \leq \Phi_{2-j,2-v}(\sigma),$$

with the same critical exception $\sigma = b, (v,j) = (2,0)$ on odd width. This applies even when the original surplus s exceeds b.

Every physical transition from a belief with features (a,c,e) can be realized, at no larger inspection cost p, with survivor features (q,i,u) satisfying

$$q = a-p, \quad i \leq c, \quad u \leq e, \quad (c-i)+(e-u) \leq p, \quad i \geq \max\{0, c+v-2\}.$$

Every physical feature (a,c,e) satisfies $c+2e \leq a \leq h-4+c+e$.

Proof. Adjoin to S its $v-i$ missing own-colored corners in $E(B)$. Their neighbors are already in B, so the neighborhood stays B. They are disjoint from the opposite-corner neighbor pairs, so u is unchanged. The augmented support has features $(q+\delta, v, u)$ and surplus $\sigma \geq 0$ by the board's perfect matching. Apply the theorem. These adjoined corners are used only in the inequality and need not belong to the previous belief.

For retention, replace an actual survivor by $A \cap E(B)$. It contains the original survivor, has neighborhood exactly B, and costs no more inspections. Its own corners are the intersection of two subsets of a two-element set of sizes c, v, giving the lower bound on i. The two own corners and the two opposite-corner neighbor pairs are pairwise disjoint when both sides are at least four. Removing an own corner and destroying an eroded corner therefore require distinct inspected rooms, proving the combined deletion charge. The same disjointness proves the feature lower bound and forces at least $(2-c)+(2-e)$ missing rooms, proving the upper bound. Since retention may reduce cost, all costs $p \leq k$ must be allowed in a necessary model. □

11.5 Equality at the pronic corner capacity

The corner bounds give a uniform history restriction at every radius where the odd-corner capacity is attained. The critical even-square endpoint must retain the corners reached across both sides.

Theorem 11.12 (Pronic rigidity on every even-area rectangle)**.** *Let $4 \leq w \leq n$, wn be even, $b = \lfloor w/2 \rfloor$, and $h = wn/2$. If $2 \leq r \leq b$ and a one-color support S satisfies*

$$|S| = r(r-1), \qquad |N(S)| - |S| = r, \qquad c(S) = 0,$$

then it is exactly a full odd triangle based at an opposite-colored board corner, with relative diagonals $1, 3, \ldots, 2r-3$. In particular, $N(S)$ has features $(r^2, 1, 0)$ in the notation of Section 11.4. Its second neighborhood contains exactly κ corners of the original color, where

$$\kappa = \begin{cases} 0, & r < b \text{ or } w \text{ is odd}, \\ 1, & r = b,\ w \text{ is even},\ n > w, \\ 2, & r = b,\ w = n \text{ is even}. \end{cases} \tag{58}$$

Consequently the next belief after any inspections of $N(S)$ has at most κ such corners.

Proof. Apply the subcritical or critical corner dichotomy. Its critical exceptions have current-corner count one or two, so neither applies. The complementary side has size $h - r^2 \geq r^2$, whereas every finite $\Psi_{0,j}(r)$ is strictly less than r^2. Thus the small branch holds. Its capacities for $j = 0, 2$ are strictly below $r(r-1)$, so $j = 1$ and equality holds in $\Phi_{0,1}(r) = r(r-1)$.

When the proof uses quadrant separation, equality in its convex surplus allocation puts all surplus into the required odd-origin component. Equality in the diagonal-layer capacity bound makes every initial odd diagonal full. Compression preserved every diagonal's cardinality, so the original support itself is that full triangle.

At a critical endpoint, the proofs of Theorems 11.3 and 11.7 supply the other possible decompositions. The odd-width and near-square alternating cases with no current corner have no outgoing corner and therefore cannot give $j = 1$. In the even square's pure alternating case, equality forces successive clean-fiber lengths $0, 1, \ldots, b-1$ and equality in every intermediate and final fiber bound; these are exactly the same triangle. The mixed square case with no current corner has no outgoing corner and is again excluded.

In a single critical whole-half-strip piece, a boundary-free matched row reduces to the quadrant case. If no such row exists, each matched row has one outside-boundary vertex and its occupied set is an initial interval of length c_i. The omitted current corner gives $c_0 = 0$; alternating rungs give $c_{i+1} \leq c_i + 2$. Equality in

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

forces $c_i = 2i$ for every i. In physical coordinates $x = 2i + (y \bmod 2)$ these intervals are the odd triangle based at $(2b-1, 0)$. This verifies the half-strip equality case directly.

The first neighborhood has relative even diagonals through $2r-2$, so its size is r^2 and its own corner count is one. At any other board corner reached by the second neighborhood, exactly one of its neighbors lies in that first neighborhood; therefore $e(N(S)) = 0$. The second neighborhood has radius $2r-1$. Among original-colored corners it can reach precisely those across an even board side of length at most $2r$. This gives (58). Inspections only decrease the following neighborhood. $\square$

Corollary 11.13 (Dual pronic restriction)**.** *Under the same board and radius assumptions, suppose*

$$|S| = h - r^2, \qquad |N(S)| - |S| = r, \qquad c(N(S)) = 2.$$

Then $c(S) = e(N(S)) = 1$ and $e(S) = 2 - \kappa$.

Proof. Let $U = V_{1-p} \setminus N(S)$. It has size $r(r-1)$, no own corner, and surplus at most r. The unconditional size profile forces surplus at least r: its square term is r, the width cap is at least r, and the complementary pronic term is at least r since $h - r(r-1) \geq r(r+1)$. Hence its surplus is exactly r and $N(U) = V_p \setminus S$. The theorem makes U a full odd triangle. Its neighborhood contains one corner, so S contains the other one. An opposite-colored corner has both neighbors in S exactly when it is absent from $N^2(U)$, giving $e(S) = 2 - \kappa$. Finally $E(N(S)) = V_p \setminus N(U) = S$, so $e(N(S)) = 1$. $\square$

These are physical restrictions on every strategy. They may be imposed as a one-step memory and a dual transition condition in a necessary model. Their validity does not assert that the remaining model paths are physically attainable. In particular, the zero-corner continuation valid on odd widths cannot be imposed at the critical even-square radius.

11.6 Localization throughout the pronic band

Exact equality is not needed to recover a support's corner location. Throughout this section let $w = 2b + 1 \geq 5$, let $n \geq w$ be even, and put

$$h = wn/2, \qquad C = b(b-1), \qquad D = C_2(b) = \max\{b(b-1)/2, b(b-2)\}.$$

Write $c(A)$ for the number of own-colored board corners in A. Write $e(A)$ for the number of opposite-colored corners whose two neighbors both belong to A.

Theorem 11.14 (Localization above the omitted-corner capacity)**.** *Let U be a one-color support with $c(U) = 0$ and*

$$D < |U| \leq C, \qquad |N(U)| - |U| \leq b.$$

Its surplus is exactly b, and it is connected under the relation of sharing a neighbor. There is exactly one board corner $o \in N(U)$, and every $u \in U$ has Manhattan distance at most $2b - 3$ from o.

Proof. Consider any nonempty sharing-neighbor component U_i and its surplus $s_i \leq b$. Theorems 11.1 and 11.3 apply with current-corner count zero. Their critical exception is unavailable, and the large branch is excluded by

$$h - |U_i| - s_i \geq h - b^2 > b^2 \geq \Psi_{0,j}(s_i).$$

Hence $|U_i| \leq s_i(s_i - 1)$ and $s_i \geq 2$. Components have disjoint neighborhoods, so their sizes and surpluses add. For $b \geq 4$, two or more components have total size at most

$$(b-2)(b-3) + 2 \leq D,$$

by convex allocation of their surplus; for $b < 4$, two components already cost too much surplus. Thus U is connected. Surplus at most $b - 1$ would give $|U| \leq (b-1)(b-2) \leq D$, so its surplus is b. The bounds $\Phi_{0,0}(b) = D$ and $\Phi_{0,2}(b) \leq D$ then force exactly one outgoing corner o.

The critical separator proof supplies an actual corner quadrant here. Indeed the longitudinal matching has more than b rows, so a row avoids the boundary. If no adjacent columns avoid it, the alternating case with current-corner count zero has outgoing-corner count zero, contrary to the preceding conclusion. Otherwise the two separators split the source or its transposed complement into quadrants; the complement case was excluded by the large-branch inequality above. Their neighborhoods agree with the whole-quadrant neighborhoods. Connectedness therefore puts all of U in the quadrant based at o.

No compression is needed to bound its radius. Touching o makes diagonal 1 occupied. A sharing-neighbor step changes the diagonal index by zero or two, so its occupied diagonals are $1, 3, \ldots, 2L - 1$ without gaps. The upper shadow of each occupied diagonal has at least one more room than that diagonal's support: for a nonempty finite integer set X, $|X \cup (X + 1)| \geq |X| + 1$. These upper shadows are disjoint, and the output origin supplies one additional room. Consequently

$$|N(U)| \geq |U| + L + 1, \qquad L \leq b - 1,$$

which gives the stated radius. □

Corollary 11.15 (Dual feature rule throughout the pronic band)**.** *Let S be a one-color survivor, put $B = N(S)$, and suppose*

$$c(B) = 2, \qquad D < h - |B| \leq C, \qquad |B| - |S| \leq b.$$

Then

$$|B| - |S| = b, \qquad c(S) = 1, \qquad e(S) = 2, \qquad e(B) = 1.$$

Proof. Put $U = V_{\mathrm{opp}} \setminus B$. Its own corners are absent, and $N(U) \subseteq V_{\mathrm{src}} \setminus S$, so its surplus is at most $|B|-|S| \leq b$. The localization theorem makes that surplus exactly b. Equality of sizes follows throughout, giving $N(U) = V_{\mathrm{src}} \setminus S$ and $|B| - |S| = b$. The unique corner of $N(U)$ gives $c(S) = 1$. Also $N^2(U)$ lies within radius $2b-1$ of its localization corner, whereas every opposite-colored corner has distance at least $n - 1 \geq 2b + 1$. Thus both such corners have all their neighbors in S, giving $e(S) = 2$. Finally $E(B) = V_{\mathrm{src}} \setminus N(U) = S$, so $e(B) = c(S) = 1$. $\square$

Corollary 11.16 (No consecutive transitions into the pronic band)**.** *For a physical inspection transition $A \longmapsto A' = N(A \setminus I)$, put $p = |I \cap A|$ and $s = |A'| - |A| + p$. Consider the condition*

$$c(A') = 2, \qquad D < h - |A'| \leq C, \qquad s \leq b. \tag{59}$$

Two consecutive transitions cannot both satisfy (59)*, regardless of their inspection budgets.*

Proof. The first transition has $e(A') = 1$ by Corollary 11.15. Every next survivor $S' \subseteq A'$ therefore has $e(S') \leq 1$, whereas another transition satisfying (59) would require $e(S') = 2$. $\square$

The existing erosion-corner feature thus records this restriction without an additional history bit. No assertion that the features characterize arbitrary reachable beliefs is needed.

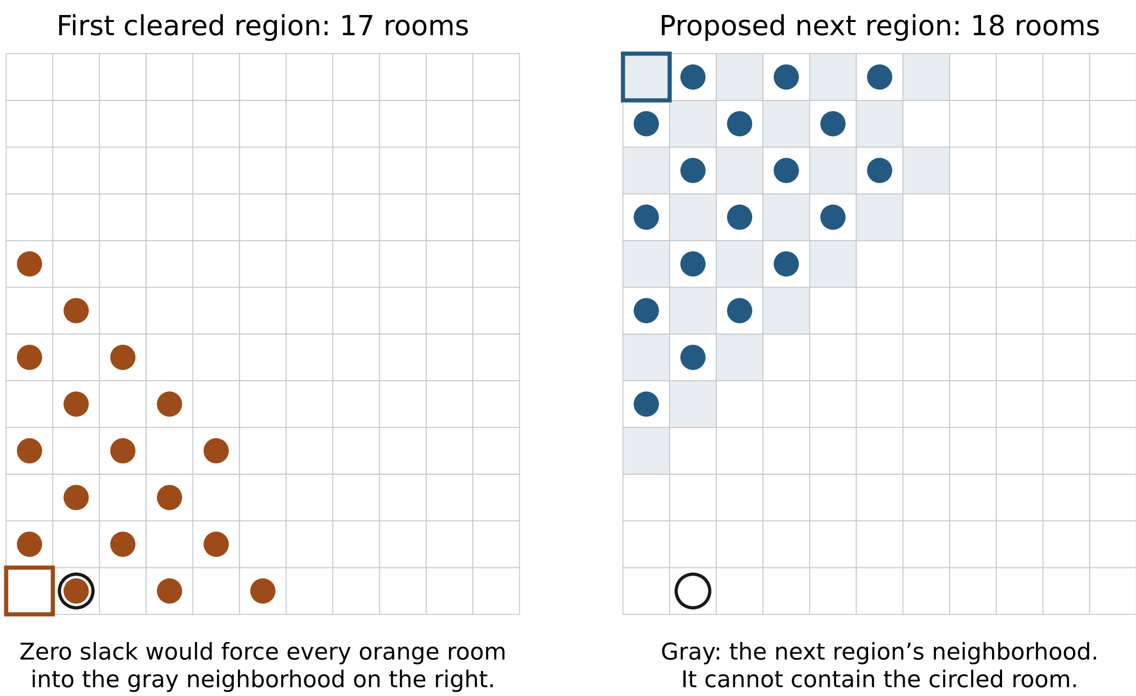

Figure 4: Two individually possible cleared regions on an 11×12 board, of sizes 17 and 18, each with surplus five. The gray region is the neighborhood of the proposed second region. At zero complement slack, every first-region room would have to lie there; the circled room prevents this. The localization proof applies to arbitrary supports in the stated band, not only the prefixes drawn here.

11.7 A near-square rigidity and a boundary-history obstruction

The corner capacities contain geometric information beyond their numerical values. Close to the square capacity, a support cannot spread across distant diagonals.

Theorem 11.17 (Near-square localization)**.** *Let U be a finite even-color support in the nonnegative quadrant, with $|N(U)| - |U| \le s$, where $s \ge 1$. If*

$$|U| > s(s-1),$$

then

$$U \subseteq \{(x,y) : x + y \in \{0, 2, \ldots, 2s-2\}\}.$$

In particular, U lies inside the s^2-room triangle based at the origin.

Proof. Apply the diagonal compression used in the proof of Theorem 4.1, which preserves the number of rooms on each diagonal. Use its notation a_j, δ_j, L, with compressed surplus $\delta \le s$. The branch with no full diagonal has at most $s(s-1)/2$ rooms and is excluded. In the other branch, that proof gives $L \le \delta$ and $|U| \le L\delta$. Hence $|U| > s(s-1)$ forces $L = \delta = s$. Each occupied diagonal contributes at least one to the surplus, so $\delta_j = 1$ there and every other contribution, including δ_{-1}, is zero.

There can be no initial or internal gap before an occupied diagonal $j \ge 1$. Such a gap would give $0 = \delta_{j-1} \ge a_j - \epsilon_j$, forcing the nonempty diagonal to be full of size one, whereas its length is $2j+1 > 1$. Thus the occupied diagonals are exactly $0, \ldots, s-1$. Their cardinalities were unchanged by compression, so the same localization holds for the original support. □

If a support is in the small branch of the finite-board quadrant decomposition, the same threshold forces it into one genuine corner quadrant. Indeed, if at least two nonempty pieces have positive integer surpluses summing to at most s, their capacities are at most

$$(s-1)^2 + 1 \le s(s-1) \qquad (s \ge 2).$$

The case $s = 1$ cannot have two nonempty pieces. A surviving single quadrant must be based at an occupied corner of the support's color: the other corner and omission profiles have capacity at most $s(s-1)$. After a reflection, Theorem 11.17 applies.

Proposition 11.18 (A physical boundary-history barrier)**.** *On the* 24×24 *board with daily budget* 13*, a solo history starting from an entire color class has at least* 177 *possible rooms after* 28 *unsuccessful inspection-and-move steps. Consequently its guaranteed solo capture time is at least* 103 *days.*

Proof. The finite necessary corner model admits no state smaller than 176 at that time, and its only 176-room feature has one occupied corner. Suppose a physical history reached such a belief A, and set $U = V_p \setminus A$. Thus $|U| = 112$ and U has one corner of its color. Reversing the first 28 inspections gives a 28-day capture strategy from $N(U)$: an avoiding path there, preceded by its edge from U, would reverse to an original avoiding path ending in U.

A checked lower potential on the same necessary model assigns at least 29 days to every feature of size at least 124. Hence $|N(U)| \le 123$. The unconditional profile gives the reverse inequality, so $|N(U)| = 123$ and the surplus is 11. In the finite corner dichotomy, the large branch is impossible: its complementary size is 165, exceeding each relevant dual capacity. The small branch also forces zero outgoing corners. Since $112 > 11 \cdot 10$, the preceding localization places U in a single genuine corner triangle of size 121.

Reflect this corner to the origin and apply the simultaneous square diagonal strategy compression C. The triangle is fixed, so $C(U)$ still lies in it and has size 112. Moreover $N(C(U)) \subseteq C(N(U))$. The unconditional profile and cardinality preservation give both sets size 123, hence equality. There are exactly eight fixed pyramid supports obtained by deleting nine rooms from the 121-room triangle. Their neighborhoods all have size 123. An independent finite certificate excludes capture in 28 days from each such neighborhood.

For completeness, the certificate uses 146 failed states for the anchored prefix and one additional failed initial state for each of the other seven neighborhoods. An independent verifier enumerates every full-quota survivor by local-maximum deletion, computes literal board neighborhoods, and checks either a decreasing-time certificate edge or a separately checked lower-potential inequality. Thus it covers all pyramid strategies. The fixed-state clause of Theorem 3.2 makes this an unrestricted lower bound from each of these initial pyramids. Strategy comparison gives

$$T_{13}(N(U)) \geq T_{13}(C(N(U))) = T_{13}(N(C(U))) \geq 29,$$

contradicting the reversed 28-day strategy.

Finally, recomputing the necessary solo model with this physical time-28 barrier removes every 102-day path; its minimum becomes 103. This last finite calculation retains every allowed target, not merely a path attaining the successive minimum cardinalities. $\square$

The physical restriction in this proposition cannot automatically be imposed on the synthetic merger of two cohorts. The following finite equality check supplies the required additional bridge.

Corollary 11.19 (The remaining 24×24 bracket)**.** *For the full board,*

$$206 \leq T_{13}(24, 24) \leq 207.$$

Proof. The established subcritical-model merger maps a pair of cohort features $(a, c), (b, d)$ to $(a + b - 288, \max(0, c + d - 2))$ while the merged size exceeds the budget. Its shortest solo path from $(288, 2)$ to $(176, 1)$ has length 28. Consequently, any joint history attaining merged feature $(176, 1)$ at that time maps to a shortest path: at stage t, its merged feature has forward distance t and backward distance $28 - t$.

A finite equality certificate enumerates every component-feature pair over these shortest-path layers and every allowed combined-quota transition. Its unique terminal pair is

$$((176, 1), (288, 2)).$$

Thus equality forces one cohort still to be full and the other to have the physical belief excluded by Proposition 11.18. This conclusion uses the final pair, and does not assume that the full cohort was never inspected earlier. Therefore every physical merged history also obeys the time-28 restriction.

An independent verifier rebuilds the weighted relations from separate budget-indexed bitset relations and expands component transitions. It checks 380 358 weighted edges and 66 295 equality-graph transitions. The resulting time-filtered subcritical solo model has clock 103 and minimum remaining size 9 after 102 steps. The guarded merger therefore shows that every full-board history of 102 unsuccessful steps leaves at least $288 + 9 = 297$ possible rooms. For a purported 205-day strategy, apply this to its first 102 inspections and to its final 102 inspections in reverse order. Their central possible sets intersect in at least $2 \cdot 297 - 576 = 18$ rooms, exceeding the central budget 13. This proves the lower bound 206.

For the upper bound, the independently replayed 104-day solo construction ends with a two-room final inspection. Use its first 103 inspections, then the union of that final inspection with its horizontal reflection, then the reversed horizontal reflections of the first 103 inspections. The central day uses four inspections, and every other day uses at most 13. Literal full-board replay verifies capture by day 207. $\square$

11.8 Concentration in the necessary corner model

The corner-count inequality also admits a concentration principle. This reduces a lower-bound calculation to one cohort; it does not assert that the model's paths are physically attainable. We first prove the principle for the base model below, then prove its closure under critical-surplus filters and one geometric memory rule. The additional restrictions are not inferred automatically from the base merger theorem.

A feature (a, c) satisfies $c \in \{0, 1, 2\}$ and $c \le a \le h - 2 + c$. A cost-p transition, $0 \le p \le k$, chooses a valid survivor feature (q, i) with

$$q = a - p, \qquad \max\{0, c - p\} \le i \le c,$$

and a valid target (t, j). For $q = 0$ require $(t, j) = (0, 0)$. Otherwise put $s = t - q \ge 0$, require the size profile of Theorem 9.1, and, when $s < b$, require

$$q \le \Phi_{i,j}(s) \quad \text{or} \quad h - q - s \le \Psi_{i,j}(s).$$

At $s \ge b$ impose no further corner condition. Every physical transition is included by Corollary 11.2. Joint transitions allocate a total of at most k inspections to two features.

Theorem 11.20 (Merger before a solo clearing day)**.** *Let $b \ge 1$, $h \ge 2b^2$, and $1 \le k < h$. Define*

$$\mathcal{M}((a, c), (a', c')) = (a + a' - h, (c + c' - 2)_+).$$

If the merged source has size greater than k, every joint model transition induces a solo model transition from the merged source to the exact merged target, using the same total inspection cost.

Proof. We first record the needed capacity algebra. Write $A = 2 - i$, $B = 2 - j$. Expanding the at most three candidates in (50) gives $\Psi_{i,j}(s) = -\infty$ for $s < A + B$. For $z = s - A - B \ge 0$ it gives

$$\Psi_{i,j}(s) = \begin{cases} z^2 + 2z + B, & B > 0, \\ z^2 + z, & B = 0 < A, \\ C_2(z), & A = B = 0. \end{cases} \tag{60}$$

For $j < 2$, the $d = 0$ candidate dominates: writing $v = s + i + j - 3$, its differences from the $d = 1, 2$ candidates are $v - 1, 3v - 5$, and feasibility of the latter gives $v \ge d + 1$. For $j = 2, i < 2$, the $d = 1$ candidate dominates whenever feasible; at the initial argument the $d = 0$ value is zero and agrees with (60). The case $i = j = 2$ has only $d = 0$. In particular every finite $\Psi_{i,j}(s)$ is at most s^2.

The dual capacities are superadditive under the merger of corner counts:

$$\Psi_{i_1,j_1}(s_1) + \Psi_{i_2,j_2}(s_2) \le \Psi_{(i_1+i_2-2)_+,(j_1+j_2-2)_+}(s_1 + s_2) \tag{61}$$

whenever both left terms are finite. Indeed the merged missing counts are $A = \min\{2, A_1 + A_2\}$ and $B = \min\{2, B_1 + B_2\}$. Put $\delta = A_1 + A_2 - A + B_1 + B_2 - B \geq 0$. The new variable is $z = z_1 + z_2 + \delta$. If $B > 0$, bound each source term by $z_l^2 + 2z_l + B_l$. The difference from the target is at least

$$2z_1 z_2 + 2\delta(z_1 + z_2) + \delta^2 + 2\delta - (B_1 + B_2 - B) \geq 0.$$

If $B = 0 < A$, use the source bounds $z_l^2 + z_l$ instead. If $A = B = 0$, all four missing counts vanish and convexity of C_2, with $C_2(0) = 0$, gives its superadditivity. This proves (61).

Two deletion inequalities handle the other combination:

$$\Phi_{(i-1)_+,j}(s+1) \geq \Phi_{i,j}(s) - 1, \qquad \Phi_{i,(j-1)_+}(s+1) \geq \Phi_{i,j}(s). \tag{62}$$

For positive j, deleting a current corner increases the pronic quadratic. At $j = 0, i = 2$ compare $(s+1)^2$ with $1 + (s-1)^2$; at $j = 0, i = 1$ use $C_2(s+1) \geq s^2 - 1$. For the second inequality, $j = 2$ gives a larger pronic quadratic; $j = 1, i > 0$ gives a larger square; and $j = 1, i = 0$ uses $C_2(s+1) \geq s^2 - 1 \geq s(s-1)$. A zero index uses monotonicity. These comparisons also preserve the domains of the capacities.

Let the chosen survivors be (q_l, i_l), the targets (t_l, j_l), and the costs p_l. Put

$$Q = q_1 + q_2 - h > 0, \quad T = t_1 + t_2 - h, \quad s_l = t_l - q_l, \quad s = s_1 + s_2,$$

$$I = (i_1 + i_2 - 2)_+, \qquad J = (j_1 + j_2 - 2)_+.$$

If $s < b$, both component profiles are subject to the corner inequality. They cannot both use its small branch: that would give $q_1 + q_2 \leq s_1^2 + s_2^2 \leq s^2 < b^2 < h$. If both use its complementary branch, their complement sizes add and (61) gives the complementary merged branch. If one is small, call its counts i, j and surplus u. Let the other omit A current corners and B outgoing corners, have surplus v, and leave complement size $d = h - t_{\text{high}} \geq 0$. Equation (60) gives $v \geq A + B$. Applying (62) A and B times gives

$$\Phi_{(i-A)_+,(j-B)_+}(u + v) \geq \Phi_{i,j}(u) - A.$$

But $Q = q_{\text{small}} - (v + d)$ and $v + d \geq A$. The small merged branch follows. Either branch implies the unconditional size profile. For $s \geq b$, that profile is automatic.

It remains to check feature validity. If the merged source is (m, c), then $m > k \geq 1$ gives $m \geq 2 \geq c$, and adding the source upper bounds gives $m \leq h - 2 + c$. The only invalid raw survivor possibility is $Q = 1, I = 2$; choose $I^* = \min\{Q, I\}$. The inspection inequalities survive because both I and $Q = m - p_1 - p_2$ are at least $c - p_1 - p_2$, and $I \leq c$. All survivor upper bounds also hold. In the subcritical complementary/ complementary case,

$$Q \geq h - s - s^2 \geq b^2 + b \geq 2,$$

so no clipping occurs. In the mixed case $\Phi_{1,J}(s) \geq \Phi_{2,J}(s)$ preserves the proved small bound: for $J = 0$ compare s^2 with $1 + (s-1)^2$, and for $J > 0$ use the pronic formula. At $s \geq b$ the profile does not depend on I.

The target upper bound follows by adding the two original upper bounds. Its lower bound could fail only at $T = 1, J = 2$. Then $Q = 1, s = 0$; both nonempty original survivors have zero surplus, so the size profile forces $q_1 = q_2 = h$, contradicting $Q = 1$. Thus (T, J) is valid. The chosen survivor (Q, I^*) and target (T, J) constitute the required solo transition of cost $p_1 + p_2$. □

Corollary 11.21 (Closure under critical filters)**.** *The merger theorem remains valid if an additional condition is imposed at surplus b, provided that it retains every profile satisfying the Φ/Ψ alternative. Keep the base rule at surplus greater than b.*

Proof. Only total surplus b is new. If both component surpluses are below b, the same capacity algebra applies. The two small branches are impossible since their survivor sizes sum to at most $b^2 < h$. In the two complementary branches, $Q \geq h - b - b^2 \geq b^2 - b \geq 2$ when $b \geq 2$, so clipping is still unnecessary. When $b = 1$, two subcritical surpluses cannot sum to b. Mixed-case clipping preserves Φ as in the theorem. The merged profile is therefore retained by the assumed filter.

If one component has surplus b, the other has surplus zero. Its survivor is nonempty and the size profile forces it to be full. Its source and target are consequently full and its cost is zero. The merger is the identity on the critical component, which already obeys the filter. □

Corollary 11.22 (Closure under the odd-triangle memory rule)**.** *On $w = 2b + 1$, with even $n \geq w$ and $b \geq 2$, equip the necessary model with the following flag, initially zero. A transition sets its new flag to one exactly when its chosen survivor and target satisfy*

$$q = b(b-1), \quad i = 0, \quad s = b, \quad (t, j) = (b^2, 1).$$

Every other transition resets the flag to zero. A flagged source permits only targets with zero corners. Every physical history satisfies this rule, and the precapture merger remains valid for the augmented model.

Proof. Physical necessity is Corollary 11.4. Define the merged history's flag from its own preceding action. We show that a newly flagged merged state comes from a flagged component with the other component full. Its total surplus is b. If both component surpluses are below b, the mixed case would give

$$Q \leq q_{\text{small}} \leq (b-1)^2 < b(b-1),$$

whereas the two complementary branches would give

$$h - b^2 \leq \Psi_{0,1}(b) \leq (b-2)^2 < b^2 \leq h - b^2.$$

Both are impossible; the two small branches were already excluded. Thus one component has surplus b and the other is a zero-cost full-to-full action. The merger is the identity on the first, so it receives the same flag.

If the current merged flag is one, its flagged component has next corner count zero. The merged next count is therefore $(0 + j_{\text{other}} - 2)_+ = 0$, satisfying the flag restriction. When the merged action does not trigger, its own flag resets to zero; no assertion that it equals the disjunction of the component flags is needed. Induction proves the claimed history merger. □

Corollary 11.23 (A one-cohort lower clock)**.** *Let $\widehat{\tau}$ be the finite minimum solo capture time in this necessary model (including either preceding augmentation), starting from $(h, 2)$, and let f_t be its minimum reachable belief size after t inspections and moves. For every $t < \widehat{\tau}$, the minimum joint model total size is exactly $h + f_t$. Consequently every physical strategy satisfies*

$$T_k(w, n) \geq 2\widehat{\tau} - 1, \qquad T_k(w, n) \geq 2\widehat{\tau} \quad \textit{if } 2f_{\widehat{\tau}-1} > k.$$

Proof. Merge the full joint start into $(h, 2)$. Inductively, if a merged source had size at most k before the last required transition, the solo path already constructed could clear on the next day, before $\widehat{\tau}$. Hence Theorem 11.20 applies at every step up to time $t < \widehat{\tau}$. It gives joint size at least $h + f_t$. Equality is attained in the model by following a minimizing solo path while leaving the other cohort full.

Put $L = \widehat{\tau} - 1$. A supposed $2L$-day physical capture has a forward belief after L inspections and moves of size greater than h. Read its final L inspections backward: after the first $L-1$ moves the reverse belief has at least $h + f_{L-1} > h + k$ rooms, and its last inspection leaves more than h. The two sets occupy the same cut and intersect, giving an avoiding walk. Thus at least $2\widehat{\tau} - 1$ days are needed. For a $(2L+1)$-day search, the two beliefs at the central day each have at least $h + f_L$ rooms. Their intersection has at least $2f_L$ rooms, which must all be inspected centrally. This proves the second bound. □

The clock is an exact evaluator inside the necessary model. Neither the merger theorem nor the lower-clock corollary asserts physical attainment of its solo paths or a fixed-size formula for its numerical value.

11.9 Concentration with eroded-corner information

The enriched inequalities still admit a one-cohort lower calculation. One deliberate relaxation makes the merger closed: for positive states we retain the upper feature bound but omit the lower cardinality bound. This only adds states to a necessary model. The theorem concerns odd-width, even-length rectangles, where the enriched critical profile has been proved; it does not assert physical attainment of model paths.

Let $b \geq 2$, $h \geq 2b^2$, and $1 \leq k < h$. A feature (a, c, e) satisfies

$$c, e \in \{0, 1, 2\}, \qquad 0 \leq a \leq h - 4 + c + e, \qquad a = 0 \Longrightarrow c = e = 0.$$

We omit $c + 2e \leq a$ when $a > 0$, and omit its survivor counterpart. The full feature is still necessarily $(h, 2, 2)$. A cost-p action, $0 \leq p \leq \min\{k, a\}$, chooses $q = a - p$. If $q = 0$, its target is empty. Otherwise choose a survivor (q, i, u) obeying the retained feature bound and

$$i \leq c, \quad u \leq e, \quad (c - i) + (e - u) \leq p,$$

and a target (t, j, v) obeying the feature bound, with $t \geq q$, $u \leq j$, $i \leq v$. Require the size profile and the maximal-retention condition $i \geq (c + v - 2)_+$. Writing

$$s = t - q, \quad \delta = v - i, \quad x = q + \delta, \quad \sigma = s - \delta,$$

require $\sigma \geq 0$. Whenever $\sigma \leq b$, require

$$x \leq \mathcal{P}_{v,u,j}(\sigma) \quad \text{or} \quad h - t \leq \Phi_{2-j,2-v}(\sigma) \quad \text{or} \quad \sigma = b, (v, j) = (2, 0). \tag{63}$$

Here $\mathcal{P}$ is (56). The reduced surplus controls this filter even when $s > b$. Theorem 11.10 and Proposition 11.11 include every physical trajectory on $w = 2b + 1$, even $n \geq w$, $h = wn/2$, after the harmless maximal-retention normalization.

Lemma 11.24 (Capacity comparisons for erosion merging)**.** *For finite inputs one has*

$$\Phi_{\min(2,a+a'),\min(2,d+d')}(s + t) \geq \Phi_{a,d}(s) + \Phi_{a',d'}(t), \tag{64}$$

$$\mathcal{P}_{(i-1)_+,u,j}(s + 1) \geq \mathcal{P}_{i,u,j}(s) - 1, \tag{65}$$

$$\mathcal{P}_{i,(u-1)_+,j}(s + 1) \geq \mathcal{P}_{i,u,j}(s) - 1. \tag{66}$$

For $j > 0$ and $i > 0$, respectively, the stronger comparisons are

$$\mathcal{P}_{i,0,j-1}(s + 1) \geq \Phi_{i,j}(s), \tag{67}$$

$$\mathcal{P}_{i-1,0,j}(s + 2) \geq \Phi_{i,j}(s). \tag{68}$$

Consequently, if $A, B, C \in \{0, 1, 2\}$, $C \geq A$, and $\Phi_{A,B}(t) \geq \max\{0, B + C - t\}$, then

$$\mathcal{P}_{(v-B)_+,(u-C)_+,(j-A)_+}(s + t) \geq \mathcal{P}_{v,u,j}(s) - \max\{t, B + C\} \tag{69}$$

whenever its right side is positive.

Proof. For (64), combine the two quadrant allocations in the proof of Theorem 11.1, temporarily allowing more than two required pieces of each type. Convexity combines the optional pieces. Truncating a required count greater than two frees one or two surplus units to another required piece of the same type; its convex gain covers the removed area. The resulting capacity is the left side.

The remaining comparisons follow from the branches of $\mathcal{P}$ and the one-neighbor capacity H of Lemma 11.9. For (65), the only possible loss is $i = 1, j = u = 0$, where $C_2(s+1) \geq s^2 - 1$. The other branches increase; when $i = 1, u = 0, j > 0$, put $z = s - 1 - 2j$ and use $H(z+4) \geq (z+4)(z+1)+1$. For (66), the only possible loss is $i = 0, u = 1$, where $H(z + 3) \geq (z + 3)z + 1$ loses at most one area unit; the other branches increase. Zero indices use monotonicity. For (67), $j = 1, i = 0$ uses $C_2(s + 1) \geq s^2 - 1 \geq s(s - 1)$, and $j = 2, i = 0$ uses $H(s + 1) \geq s^2 - s - 1 \geq s^2 - 5s + 8$ for $s \geq 4$. When $i > 0$, the new even-piece quadratic dominates the old Φ. For (68), $i = 1, j = 0$ uses $C_2(s + 2) \geq s(s + 2)$; at $i = 1, j > 0$ use $H(z + 5) \geq (z + 5)(z + 2) + 1$; at $i = 2$ the new even-piece quadratic again dominates. These substitutions also preserve the capacity domains.

To prove (69), first suppose $A > 0$. Apply (67) A times, resetting u to zero, and (65) B times. A zero outgoing index uses monotonicity. This needs $A + B$ surplus and loses at most B; feasibility gives $t \geq A + 2B$. Monotonicity in u permits the desired possibly positive final count instead of zero. If $A = 0$ and $t \geq B + C$, use (65) B times and (66) C times. If $A = 0$ and $t < B + C$, feasibility and $\Phi_{0,0}(0) = \Phi_{0,0}(1) = 0$ leave only $B = 1, C = 2, t = 2$. For $v = 0$, two pair deletions lose at most two; for $v > 0$, (68) applies. Both suffice for the permitted loss three. □

Theorem 11.25 (Enriched precapture merger)**.** *Define*

$$\mathcal{M}((a, c, e), (a', c', e')) = (a + a' - h, (c + c' - 2)_+, (e + e' - 2)_+).$$

When the merged source has size greater than k, every joint transition of total cost at most k in the stated model induces a solo transition to its exact merged target, at the exact same cost.

Proof. For component survivors (q_l, i_l, u_l) and targets (t_l, j_l, v_l), put

$$Q = q_1 + q_2 - h > 0, \quad T = t_1 + t_2 - h, \quad I = (i_1 + i_2 - 2)_+, \quad U = (u_1 + u_2 - 2)_+,$$

$$J = (j_1 + j_2 - 2)_+, \quad V = (v_1 + v_2 - 2)_+.$$

Write $\delta_l = v_l - i_l$, $x_l = q_l + \delta_l$, $\sigma_l = t_l - q_l - \delta_l$. The merged closure has $\delta = V - I \leq \delta_1 + \delta_2$. With $\eta = \delta_1 + \delta_2 - \delta \geq 0$, its closed size and surplus are

$$Q + \delta = x_1 + x_2 - h - \eta, \qquad \sigma = \sigma_1 + \sigma_2 + \eta. \tag{70}$$

If $\sigma > b$, the profile is automatic. Otherwise both component filters apply. If one uses the critical exception, say $\sigma_1 = b$, then $\sigma_2 = \eta = 0$. At zero reduced surplus the positive closed support cannot use the small branch; its high branch forces $(t_2, j_2, v_2) = (h, 2, 2)$ and $x_2 = h$. Thus the merged closed profile is exactly the critical one, with $V = 2, J = 0$. No claim that this component's original source is full is needed here.

Otherwise both use capacity branches. They cannot both be small because $q_1+q_2 \le x_1+x_2 \le \sigma_1^2+\sigma_2^2 \le b^2 < h$. If both are high, their target-hole sizes add, and (64) with surplus monotonicity in η gives the high merged branch. If one is small, put $A = 2-j_H$, $B = 2-v_H$, $C = 2-u_H$, $t = \sigma_H$, and $d = h - t_H$ for the high component. Then $C \ge A$, $d \le \Phi_{A,B}(t)$, and its survivor upper bound gives $t+d = h - x_H \ge B+C$. The merged closed size is $x_L - (t+d) - \eta$. Apply (69) and then surplus monotonicity in η to get the low merged branch. Either branch implies the old size profile since both capacities are at most the square of their surplus.

It remains to check the features and inspection charge. Adding the retained upper bounds proves each merged upper bound, because $\min(c_1+c_2,2)+\min(e_1+e_2,2) \le 4$. The merged source, survivor, and target are positive, so no empty-state or lower cardinality issue arises. Monotonicity gives $U \le J$, $I \le V$. The positive-part map is one-Lipschitz, so the losses of the merged current and eroded counts sum to at most the four component losses, hence at most p_1+p_2. Finally component maximality gives

$$I \ge (c_1+c_2+v_1+v_2-6)_+ \ge \Big((c_1+c_2-2)_+ + V - 2\Big)_+ .$$

This is exactly merged maximality. The raw survivor (Q, I, U) therefore supplies the required transition at the exact total cost. $\square$

Corollary 11.26 (Uniform pronic restrictions preserve merging)**.** *On the stated odd-width, even-length boards, the merger remains valid with all the following restrictions, for every $2 \le r \le b$:*

- *A survivor $q = r(r-1)$, $i = 0$, $s = r$ has target $(r^2, 1, 0)$. Flag that target; on the next step the current-corner count must be zero, then reset the flag unless a new trigger occurs.*
- *If $q = h - r^2$, $s = r$, and $j = 2$, require $(i, u, v) = (1, 2, 1)$.*

These are necessary physical rules by Theorem 11.12.

Proof. A merged primal trigger has original total surplus r. If both component surpluses were positive, they would each be below r. A small component would give $Q \le x_L \le (r-1)^2 < r(r-1)$. With two high components, the merged hole size $h - r^2 \ge r^2$ is bounded by $\Phi_{2-J,2-V}(r-V)$, since $I = 0$. For $V > 0$ this is at most $(r-1)^2$, and for $V = 0$ its second index is two, again giving a value strictly less than r^2. Both are impossible. Consequently one component has original surplus zero. Its survivor, source, and target are all full, with cost zero. Merging is the identity on the other component, which has exactly the same trigger and flag.

For a merged dual event, $J = 2$ forces $j_1 = j_2 = 2$. A small component cannot supply the required near-full total. With two positive high surpluses, the two hole sizes sum to at most $s_1(s_1-1) + s_2(s_2-1) = r(r-1) - 2s_1s_2 < r(r-1)$, contradicting the required value. Again a full-to-full component makes the merger the identity on the constrained component.

For flags, use the merged history's own trigger, maintaining that a flagged merged state has one flagged component and the other full. The formation proof establishes this invariant. Its next merged current count is at most the flagged component's count, hence zero. Reset on nontriggering actions; a disjunction of component flags is not assumed. $\square$

Corollary 11.27 (The near-pronic band restriction preserves merging)**.** *Put $C = b(b-1)$ and $D = C_2(b)$. The merger remains valid when every transition with $j = 2$, $D < h - t \le C$, and original surplus $s \le b$ is required to satisfy $s = b$ and $(i, u, v) = (1, 2, 1)$. This necessary physical rule is Corollary 11.15. It already forbids consecutive band transitions through $u \le e$; no additional memory flag is needed.*

Proof. A merged event has $J = 2$, hence $j_1 = j_2 = 2$, and its component hole sizes sum to a number in $(D, C]$. Each is therefore at most C. A small component would have $x_l \le \mathcal{P}_{v_l,u_l,j_l}(\sigma_l) \le \sigma_l(\sigma_l - 1)$, hence $t_l \le \sigma_l^2 \le b^2$. Its holes would be at least $h - b^2 \ge b^2 > C$, impossible. The critical exception has outgoing count zero, so both components must use their high branches. Consequently

$$d_l \le \Phi_{0,2-v_l}(\sigma_l) \le s_l(s_l - 1).$$

If both original surpluses are positive, their sum at most b gives

$$d_1 + d_2 \le s_1(s_1 - 1) + s_2(s_2 - 1) \le (b - 1)(b - 2) \le D,$$

contradicting the event. Thus one component is original-full-to-full at cost zero. The merger is the identity on the constrained component, including its surplus and all three survivor/target features. □

The induction and reversed-walk argument of Corollary 11.23 now apply to this enriched model, including all pronic radii and the static near-pronic band restriction. If its finite solo clearing time is $\hat{\tau}$ and its smallest reachable size after t moves is f_t, the minimum joint total before solo clearing is $h + f_t$, and

$$T_k(w, n) \ge 2\hat{\tau} - 1, \qquad T_k(w, n) \ge 2\hat{\tau} \quad \text{if } 2f_{\hat{\tau}-1} > k.$$

This is an exact lower-model evaluator. Physical attainment and a fixed-size numerical evaluation of its clock remain separate questions.

11.10 Finite applications and a remaining gap

The geometric theorems apply uniformly, but the following numerical applications evaluate finite necessary models separately. Their matching upper bounds are inspection sequences replayed on the actual board. These distinctions matter: neither a necessary-model path nor an unmatched construction is an optimal physical strategy.

Theorem 11.28 (Certified finite capture times)**.** *The following values hold:*

$$T_5(8, 8) = 40, \qquad T_7(12, 12) = 72, \qquad T_4(7, 8) = 68, \qquad T_5(9, 10) = 96,$$
$$T_6(11, 12) = 128, \qquad T_6(11, 14) = 172.$$

For each equality, the accompanying artifact specifies a legal inspection sequence attaining the stated time.

Proof. For the two even squares, enumerate the size–corner model of Section 11.8, retaining every target allowed by the subcritical corner inequalities and every allocation of at most k inspections. Joint reachability first reaches a total of at most k after 39 and 71 unsuccessful days, respectively. The next inspection gives the lower bounds 40 and 72. The retained transitions include every physical strategy by Theorem 11.1. No square-compression assumption is needed for these lower certificates.

For 7×8 and 9×10, use the erosion-corner model with full reduced-surplus closure, maximal retention, and the critical pronic restrictions. Independent integer potentials on the joint state space give lower bounds 68 and 96: they vanish at capture and decrease by at most one on every permitted joint transition. The verifier checked 476 281 164 and 4 353 827 726 transition inequalities, respectively. These two lower certificates do not require the merger theorem.

For the two eleven-row boards, use instead the relaxed model of Theorem 11.25, including all pronic radii and Corollary 11.27. In particular, the two positive lower feature bounds omitted in that theorem are also omitted in the finite evaluator. Its pairs $(\widehat{\tau}, f_{\widehat{\tau}-1})$ are $(64, 4)$ and $(86, 4)$. As $2 \cdot 4 > 6$, the concentration and reversed-cut argument gives lower bounds 128 and 172, not merely 127 and 171. The forward-distance certificates check every edge from a reachable state and the minimum residual size before clearing.

The upper certificates supply a 20-day solo sequence on 8×8, a 36-day solo sequence on 12×12, and full-board sequences of lengths $68, 96, 128, 172$ on the four odd-width boards. Clearing the two cohorts successively in the first two cases gives 40 and 72 days. The recorded rooms are checked against the daily budget and the exact recursion $A_{t+1} = N(A_t \setminus S_t)$, starting from the appropriate full color class or full board and ending with no survivor. Thus the physical upper bounds match the lower certificates. □

The precise reproducibility records are as follows. Paths are relative to the research artifact; historical phrases such as "proof-pending" in an early receipt record the status at its generation, while the geometric premises are proved in the preceding sections.

- The even-square lower evaluator is `src/resumed_even_construction_formula_profile.py`; its receipts are `research/resumed-even-construction-formula-profile-8x8-k5.json` and `research/resumed-even-construction-formula-profile-12x12-k7.json`. The upper records are `research/resumed-even-construction-square-8x8-k5.json`, replayed by `src/resumed_even_construction_square_audit.py`, and `research/resumed-even-construction-transport-check.json`, replayed by `src/resumed_even_construction_transport.py`.
- The 7×8 and 9×10 joint lower records are in `research/resumed-odd-even-erosion-corners-full-closure-audited.json`, with evaluator `src/resumed_odd_even_erosion_corners.py` and independent joint-potential checker `src/resumed_odd_even_corner_history.py`.
- The two eleven-row lower records are in `research/resumed-erosion-merger-band-dual-solo-check.json`, generated by `research/resumed-erosion-merger-band-solo-check.py`. The receipt records the source hash and the deliberate feature relaxation; it uses the static band-dual rule, with no band-memory bit.
- All four odd-width upper sequences and their full-board replays are in `research/resumed-odd-even-minimum-upper.json`, generated by `src/resumed_odd_even_minimum_upper.py`.

For comparison, the present full-board bounds on the larger square are

$$206 \le T_{13}(24, 24) \le 207. \tag{71}$$

Corollary 11.19 proves the lower bound using the near-square physical obstruction and a separately verified time-sensitive joint equality certificate. Its reproducibility records are `research/resumed-even-construction-near-square112-verification.json` and `research/resumed-even-construction-joint-equality28-verification.json`. The latter verifier is `src/resumed_even_construction_joint_equality_verify.py`. The upper bound is the literal 207-day full-board replay in `research/resumed-even-construction-union-intersection-24x24-k13.json`, from `src/resumed_even_construction_union_intersection.py`. Its underlying solo sequence takes 104 days; reflection across the central capture inspection saves one day from simple concatenation. The underlying solo interval remains 103 to 104. The lower bound in (71) uses the additional joint equality argument; it is not inferred merely by doubling the physical solo lower.

These are ordinary geometric proofs combined with checked finite integer calculations and physical inspection replays. They have not been formalized in Lean. They establish the displayed individual values and bounds, rather than a new all-width classification or a uniform fixed-size numerical formula.

12 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 12.1 (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*

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

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

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

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. □

12.1 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 $\chi : V(Q) \to \{0, 1\}$. 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 $\mathrm{cl}(R)$ for the least such set containing R. These are ideals of a finite poset, since every predecessor lowers y. Put

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

The sharper, phase-dependent threshold needed for erosion is

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

Lemma 12.2 (Prescribed envelopes with arbitrary terminal supports)**.** *Let R be any set in physical color p, and let k satisfy*

$$|\,\mathrm{cl}(R) \cup Z_p| \le k \le |V_p|, \qquad Z_p = \{(v,0) : \chi(v) = p\},$$

where V_p is the full physical color class in the cylinder. In particular, admission holds if $|\,\mathrm{cl}(R)| \le k$ and $k > B_{Q,p}$, or if R is already a pyramid and $k = |R| + m \le |V_p|$ with $m \ge |Z_p|$. There are pyramids $A \supseteq R$ of color p and E of color $1-p$ such that

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

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

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

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.

Every bottom root is a minimal poset element, so $\mathrm{cl}(R) \cup Z_p$ is an ideal. Extend it to exactly k elements by repeatedly adding a minimal element of its complement; fix a linear extension to make this choice deterministic. Call it A. For the first sufficient admission condition, an arbitrary k-element ideal containing $\mathrm{cl}(R)$ already contains every root by the preceding distance bound. For the second, adjoining all roots adds at most $|Z_p| \le m$ elements. 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 $s_v = 0$. All these fibers are occupied by the chosen roots, 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 $n \geq \operatorname{diam}(Q)$: at a maximizing vertex v, choose $s_v = 1$ and heights $a_u = \operatorname{dist}(v, u) - 1$. 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,

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

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

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

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 $|\operatorname{cl}(R)| - |R| \leq m$. For a pyramidal R, deliberately adjoining the roots instead gives the same backward step whenever $m \geq |Z_p|$. The quadratic threshold is therefore unnecessary for this chosen upper construction; its role in the large-budget lower-bound argument is separate.

12.2 An all-budget upper bound and its exact range

Theorem 12.3 (Every even width: construction and large-budget equality)**.** *Let $b + 1 \leq m < h$, and write*

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

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

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

Equality holds whenever $m > b(b-1)$. At smaller budgets this theorem alone supplies an explicit upper bound. The boundary budget $m = b(b-1)$ is classified separately in Theorem 31.4, whose construction may improve the displayed bound. 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.

Lower bound. Assume in this part that $m > b(b-1)$. 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

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

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 12.1 gives precisely the claimed lower bounds. Its argument also covers the two extreme budget ranges. □

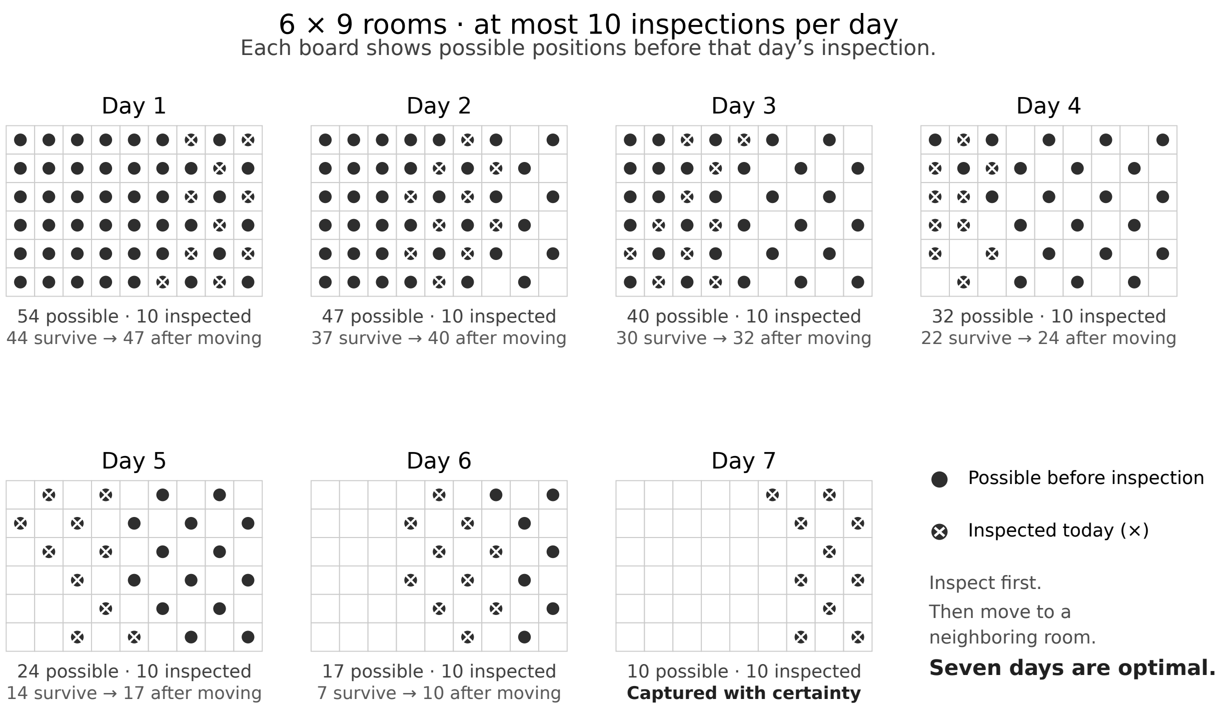

Figure 5: An optimal seven-day search on 6×9 with ten daily inspections, constructed backwards from the final survivor envelope. Every displayed possible-position set and inspection has been replayed on the physical room graph.

Physical upper construction. Apply Lemma 12.2 with $Q = P_{2b}$. Deliberately include its b bottom roots at each enlargement. Since $m \geq b+1$, this is always admitted, 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 adds enough rooms to include all roots, 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 $m \geq 2h$. $\square$

13 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

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

where k_v is the fiber cardinality. The values $-2, -1$ denote empty fibers. This is the height representation established in Lemma 12.2. All the results have reflected versions for upward pyramids.

Proposition 13.1 (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

$$e_v = \max\{b_v - 1, (1 - s_v) - 2\} + 2\mathbf{1}_{\{T_v\}}, \tag{72}$$

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

$$|P \setminus E(B)| = \frac{1}{2} \sum_v (a_v - e_v)_+ . \tag{73}$$

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 (72) 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. ☐

Lemma 13.2 (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*

$$\begin{gathered} z_v \equiv a_v + t \pmod 2, \qquad z_v \leq a_v + t \quad \textit{for every } v, \\ F := \frac{\sum_v a_v - \sum_v z_v + At}{2} \leq \sum_{j=1}^{t} p_j. \end{gathered} \tag{74}$$

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. ☐

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

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

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

$$|N(S)| \geq |S|, \qquad |S| > C_A \implies |N(S)| \geq |S| + \gamma_p. \tag{75}$$

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 13.3 (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*

$$\begin{gathered} t \leq \min_v a_v, \quad t \leq \min_v z_v, \qquad \max_v a_v,\ \max_v z_v \leq n-1-t, \\ |P| - P_* > C_A. \end{gathered} \tag{76}$$

An actual t-day block with these daily quotas and exact endpoint Z exists, without restrictions on intermediate supports, if and only if (74) 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 (75) bounds every survivor below by $|P| - P_* > C_A$. If the initial color is p, its jth expansion is therefore at least γ_{p+j-1}. For $s_v^p = (p - \chi(v)) \bmod 2$,

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

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 13.2. Its unshifted heights stay between $z - t\mathbf{1}$ 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 $r = \operatorname{diam}(Q)$ and $\mu(x) = A^{-1} \sum_v x_v$.

Proposition 13.4 (Balanced inspections in fixed physical margins). *Let Q be connected and bipartite with $A \geq 2$, and suppose the criterion* (74) *holds for a constant quota m and $t \geq 1$. If*

$$r + 2 \leq \mu(a), \mu(z) \leq n - r - 3, \tag{77}$$

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

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

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

$$t_0 \geq \max \left\{ 0,\ \max_v (z_v - a_v), \quad \left\lceil \frac{\sum_v a_v - \sum_v z_v}{D} \right\rceil \right\}. \tag{78}$$

Thus two maxima, an integer ceiling, and a parity adjustment give its cost. Under (77) 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 (76) with a large mt_0. The finite ports below handle those excursions explicitly.

13.1 Constructive transport without boundary margins

Lemma 13.5 (Upper transport without boundary margins). *Let Q be a connected bipartite transverse graph, and consider $Q \square P_n$. Let a_v and b_v be integer cutoff vectors of physical phases p and $p + t$, respectively, with adjacent heights differing by one. A belief is contained in the envelope a when every possible room (v, y) satisfies $y \leq a_v$. If*

$$F_t(a, b) := \frac{1}{2} \sum_{v \in V(Q)} (a_v - b_v + t)_+ \leq kt,$$

there is a directly prescribed t-day strategy, with at most k inspections per day, whose final belief is contained in the envelope b. There is no requirement that either envelope lie away from a board boundary.

Proof. Put $c_v = \min(a_v, b_v - t)$. The two arguments have the same parity; their minimum is edgewise 1-Lipschitz, so adjacent values differ by one. Starting from a, repeatedly lower by two a highest vertex still above c, with a fixed tie order. Every neighbor is one lower: a higher neighbor at target would violate the target's Lipschitz condition, and a higher neighbor above target would contradict the choice. The resulting legal flip word has exactly $F_t(a, b)$ letters and ends at c.

Partition the word into t consecutive groups of at most k letters. On each day inspect the old top room of every flip in that day's group when that room lies inside the board; distinct flips name distinct rooms. Then increase every virtual height by one. This dominates physical movement, both longitudinally and transversely, even when a virtual height lies below or above the board. Uniform additions commute with legal height flips. The final virtual envelope is $c + t = \min(a + t, b) \leq b$, proving the containment. When $t = 0$, the hypothesis says $a \leq b$ and no inspection is needed. $\square$

Corollary 13.6 (A bottom-boundary credit)**.** *Suppose a, b are canonical physical pyramid cutoffs and $t \geq 1$. Let $z(B)$ be the number of empty target fibers whose bottom cell has the target physical color. The same containment conclusion holds under the weaker sufficient condition*

$$F_t(a, b) \leq kt + z(B).$$

Proof. In the preceding legal flip word, a flip at unshifted old height $H \leq -t$ lies below the physical board on every day $0, \dots, t - 1$. There is exactly one such flip in each empty target bottom-root fiber: there $b_v = -2$, the target $c_v = -t - 2$, and the canonical initial height is at least $-t$. An empty nonroot target fiber has $b_v = -1$ and its last flip is at height $1 - t$; nonempty target fibers end higher still. Thus exactly $z(B)$ letters are guaranteed free. Partition the other letters into t groups of at most k, retaining all free letters in their original order. Actual visible inspections never exceed the charged count. The virtual endpoint remains $\min(a + t, b) \leq b$. $\square$

Corollary 13.7 (A linear-rate change of corner)**.** *Let the board have even width $w = 2r$. Let A and B be weightlex prefixes based at the two bottom corners, with $|A| \geq |B|$, and let their physical phases differ by $t \geq w - 1$. If*

$$|A| - |B| \leq (k - r)t + z(B),$$

there is a directly prescribed t-day strategy from A into B. In particular, the simpler condition $|A| - |B| \leq (k - r)t$ suffices.

Proof. Let a, b be the canonical cutoffs. Reflect B horizontally. If t is odd, the reflected prefix has the same phase as A and is contained in A. Thus $b_x \leq a_{w-1-x} \leq a_x + w - 1$. If t is even, the reflected prefix has the opposite phase. Since the board has a perfect matching, $|N(A)| \geq |A| \geq |B|$; the neighborhood of a weightlex prefix is a prefix in the same corner order. Consequently the reflected B lies in $N(A)$, and $b_x \leq a_{w-1-x} + 1 \leq a_x + w$. Since w is even and $t \geq w - 1$, we have $b_x - a_x \leq t$ in both cases. The canonical cutoff sums give

$$F_t(a, b) = |A| - |B| + rt.$$

The preceding bottom-boundary credit therefore proves the claim. $\square$

Corollary 13.8 (A size-only sufficient transport criterion)**.** *On an even-width board $w = 2r$, let A, B have canonical pyramid cutoffs and phases compatible with $t \geq 1$, and write $D = |A| - |B|$. If*

$$D + rt \geq r(w - 1), \qquad D \leq (k - r)t + z(B),$$

there is a directly prescribed t-day strategy from A into B.

Proof. The difference $\delta_x = b_x - a_x$ is 2-Lipschitz and has sum $-2D$. If its maximum is M, summing the bounds $\delta_x \geq M - 2|x - x_0|$ gives $-2D \geq wM - w(w-1)$, whence $M \leq w - 1 - D/r \leq t$. Thus $F_t(a, b) = D + rt$, and the bottom-boundary credit applies. □

14 Odd widths and even lengths: construction and large-budget equality

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

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

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

Theorem 14.1 (Odd width and even length: an all-budget construction)**.** *Suppose $b+1 \leq m < h$, and write*

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

Then

$$T_m(P_w \,\square\, P_n) \leq \begin{cases} 2q, & r = 0, \\ 2q + 1, & r > 0 \text{ and } 2J(r) \leq m, \\ 2q + 2, & r > 0 \text{ and } 2J(r) > m. \end{cases} \tag{79}$$

Equality holds for $m \geq b(b+1) + 1$. The proof of the construction uses every feasible budget; the lower-bound proof below retains this separate quadratic hypothesis. 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.

14.1 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 $U = \{0, 2, \ldots, 2b\}$ and $V = \{1, 3, \ldots, 2b-1\}$.

Lemma 14.2 (Critical orientation and a corner consequence)**.** *A phase-zero strictly middle critical survivor consists of prefixes of lengths t_x with*

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

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

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

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 $x \leq a^2$. □

14.2 The physical construction

Use the root-selected form of Lemma 12.2 with $Q = P_{2b+1}$. A physical-color-p envelope A containing all its $b + 1 - p$ bottom roots erodes to an opposite-phase pyramid E with

$$N(E) \subseteq A, \qquad |E| = |A| - (b + 1 - p). \tag{80}$$

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, including all phase roots, then apply (80). This backward step raises its rank by exactly D; the root choice is admitted since $m \geq b+1$. 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 (79).

14.3 A corner potential with at most one lost inspection

Throughout the lower-bound argument that follows, assume $m > B = b(b+1)$. Put $N = 2h$ and retain the function F of (36). Define

$$C = \min_{0 \leq x \leq B} \{x + F(2(h - x + \rho(x)))\}. \tag{81}$$

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

Lemma 14.3 (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 $U \in \{N, N-1\}$,*

$$x \leq a^2 \implies x + F(U - 2(x - a)) \geq F(U), \tag{82}$$

$$x + F(U - 2(x - a) + 1) \geq F(U). \tag{83}$$

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

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

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

$$\lambda - (d + c - 2)^2 \geq 2d(a - c + 2) + c > 0,$$
$$\lambda - (d + c - 1)(d + c - 2) \geq 2d(a - c) + 3d + c > 0.$$

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

$$\lambda - (d + c - 1)^2 \ge 2d(a - c + 1) + a - c + 3 > 0,$$
$$\lambda - (d + c)(d + c - 1) \ge 2d(a - c) + d + a - c + 2 > 0.$$

Here $c \le a$ in both rows. They prove (82). For the corrected target the discounts are $d + c - \rho_1(\lambda)$ and $d + c - \rho_0(\lambda + 1)$, controlled by

$$\lambda - (d + c - 1)(d + c - 2) \ge 2d(a - c) + 3d + 2 > 0,$$
$$\lambda + 1 - (d + c - 1)^2 \ge 2d(a - c + 1) + 2 > 0.$$

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 (83). Finally apply it at $U = N - 1$, with $a = \rho(x)$, to obtain $C \ge F(N - 1)$; $x = 0$ gives $C \le F(N)$. $\square$

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

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

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 14.2: a critical middle move has marks $0 \to 1$, 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 $b^2 + b + 1 \le m$, 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

$$T_m \ge \lceil 2C/m \rceil . \tag{84}$$

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

$$r > 0, \qquad J(r) = \lfloor m/2 \rfloor + 1, \qquad C = F(N) - 1. \tag{85}$$

14.4 Separating the two corners on longer boards

Suppose

$$h > m + 2b^2 + b. \tag{86}$$

It suffices to settle (85). Now mark a belief if its preceding survivor was either strictly middle critical, or co-small with deficiency x, surplus $a \le b$, and $x > a^2$. Use the common-cap potential

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

The new mark also forbids a critical next middle move by Lemma 14.2. For a co-small target, the corner estimates give

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

use (82) when $x \le a^2$ and (83) otherwise. Surplus greater than b already gives the interior rank inequality.

For a source deficiency d, monotonicity by two rank units gives $d + F(N-2d) \le F(N)$. A marked source has $d \ge 1$, and (85) together with $C \ge F(N-1)$ gives

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

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 (86) 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 $T_m \ge \lceil 2F(N)/m \rceil = 2q+2$ in the sole possible gap.

Every $q \ge 4$ satisfies (86), since

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

Also $q \ge 1$ 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 $t = x-a \le a^2$ costs $h+a$ when $t \ge s$, 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 (85), write $r = 2k+1$ and $c = \min(R(k+1), b)$. Then $k + c = \lfloor m/2 \rfloor$, so $k \ge b(b-1)/2$, while $h = 3m - 2b + k$. Hence

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

Only $q = 2$ still requires an argument.

14.5 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

$$D_* = h - k - c. \tag{87}$$

Use the symmetric inverse profile

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

whose superadditivity was proved in Lemma 12.1. 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 $a + a' \le b$, forcing $m \le a(a+1) + a'(a'+1) \le B$, 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 $h - Z = k + 1$,

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

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

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 14.2 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 $h - L_1(k+u) + u \le h - k - c$. If $s \ge h - B$, then $u \ge 2m - b - B > m - b \ge b^2 + 1$, 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 $h - k - b - 1 \le h - k - c$. This proves (87). 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 (87), 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 (84). The one- and two-day regimes follow by inspecting whole color classes. This completes Theorem 14.1.

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.

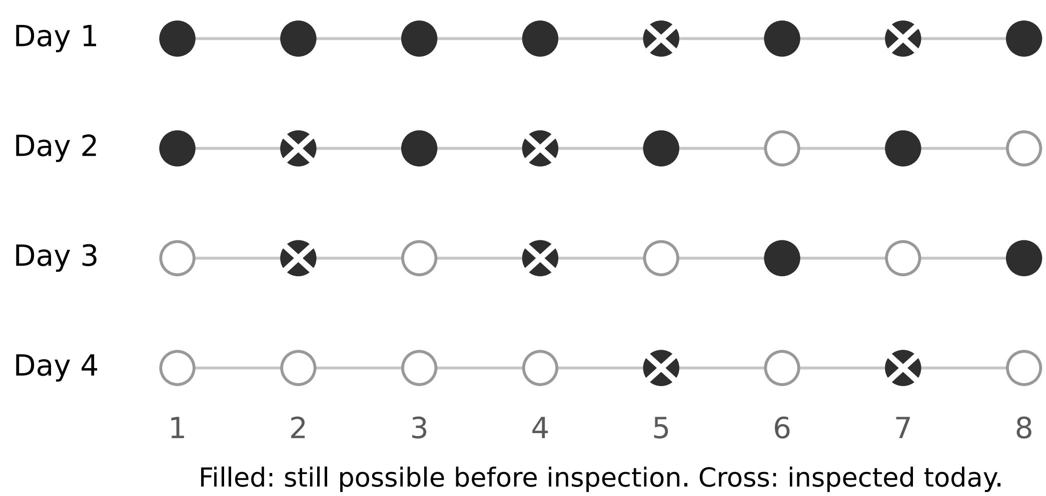

Figure 6: An optimal four-day search of eight rooms with two daily probes. The rows show every position still possible before that day's inspection.

15.1 Neighborhood inequalities

Number the vertices $0, \ldots, n-1$. Each set below is contained in one color class. If n is even, the edges $(0,1), (2,3), \ldots$ give a perfect matching and therefore $|N(R)| \ge |R|$. 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,

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

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. □

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

$$F_C(u,p) = \begin{cases} 0, & \lfloor u/2 \rfloor \leq p, \\ \min\{C, u - 2p + c\}, & \lfloor u/2 \rfloor > p, \end{cases} \tag{88}$$

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*

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

For $0 \leq p \leq m$, the constant-cap transition (88) *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

$u \le C$. Otherwise put $v = u - 2p + c$. It is positive and at most $u \le C$, so the cap does not act. Since $2p - c \le D$, $J(u) \le J(v) + 1$. Therefore

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

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

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

$$n = qD + r + 2, \qquad 1 \le r \le D.$$

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

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

The resulting lower bound is $2q + 1$ when $r \le m - 2$, 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

$$C(a) + L \le 2P + 2, \qquad P \le Lm.$$

Since $C(a) \ge n$, these imply $L \ge q + 1$ and $P \ge A := qm + m/2$. 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 $1 \le a \le q - 1$ 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 $n \ge 3$, 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 $1, \ldots, n$. Put

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

For $1 \le r < n$, $N(A(r)) = A(r + 1)$. 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

$$e = \left\lceil \frac{n-2}{d} \right\rceil, \qquad \rho = n - 2 - (e-1)d, \qquad c = \lfloor \rho/2 \rfloor + 1, \qquad \ell = m - c.$$

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

$$2e - 2 + \left\lceil \frac{2\rho + 2}{d} \right\rceil \quad (n \text{ odd}), \qquad 2e - 2 + \left\lceil \frac{\rho + 2c}{d} \right\rceil \quad (n \text{ even}).$$

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 $B \subseteq \{1, \ldots, n\}$. 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 $2 \le m < 2n$, 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 $T \ge \lceil 2N/M \rceil$.

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 $N = qM + s$, $0 \le s < M$. 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 $q \ge 1$, then $M \le n - 1$ and $M - s < n$ as well. Thus $b \ge 2$ 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 $n \ge 2$ 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 $n \ge 2$; 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

$$a_i = (2i, 0), \qquad b_i = (2i, 2), \qquad c_i = (2i + 1, 1) \quad (0 \le i < s).$$

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

$$a_i \leftrightarrow c_i \leftrightarrow b_i, \qquad a_i \to a_{i-1}, \quad b_i \to b_{i-1}, \quad c_i \to c_{i+1},$$

with out-of-range subscripts omitted. Consequently, for every survivor set R in this parity,

$$|N_G(R)| - |R| = |N_D^+(R) \setminus R|. \tag{89}$$

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, $|R| \le h-2$, and $|N(R)| = |R|+1$, then*

$$R = \{a_i, b_i, c_i : 0 \le i < j\} \cup U, \qquad U \subseteq \{a_j, b_j\}, \quad 0 \le j < s.$$

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 $|R'| \le h-2$, then $|R| = h-2$ and $|R'| = 1$.

Proof. By (89), 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. □

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 $\rho(\varnothing) = 0$. This label records history and imposes no shape restriction on A.

For a numerical rank $0 \le u \le 2h$ and an allocation p, put

$$F_h(u,p) = \begin{cases} 0, & \lfloor u/2 \rfloor \le p, \\ \min\{2h, u - 2p + 3\}, & \lfloor u/2 \rfloor > p. \end{cases} \tag{90}$$

It is nondecreasing in u and nonincreasing in p.

Lemma 17.2. *If $p \le m < h-2$ useful probes leave $R \subseteq A$, then $\rho(N(R)) \ge F_h(\rho(A), p)$.*

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 $1 \le r \le h-2$. 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

$$(u, v) \longmapsto (F_h(u, p), F_h(v, m-p)), \qquad 0 \leq p \leq m.$$

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

$$E :; (0,0), \quad ((2i-2,2),(2i-1,1),(2i,0))_{i=1}^{s-1}, \quad (2s-2,2),(2s-1,1),$$
$$O :; (0,1),(1,0), \quad ((2i-1,2),(2i,1),(2i+1,0))_{i=1}^{s-1}, \quad (2s-1,2).$$

Empty repeated blocks are omitted. Reading the neighbors of successive vertices in these lists gives, for every $k > 0$,

$$N(E_k) = O_{\min(h,k+1)}, \qquad N(O_k) = E_{\min(h,k+2)}.$$

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 (90), 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 (90), 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 14.1 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

$$P \geq \left\lceil \frac{W + 3\ell}{2} \right\rceil, \qquad \ell D \geq W. \tag{91}$$

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

$$J(r) = \lfloor r/2 \rfloor + 2, \qquad \delta(r) = r \bmod 2, \qquad 2J(r) + \delta(r) = r + 4 \quad (r > 0).$$

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 12.3. 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*

$$T_3(G) = 4n - 4, \qquad T_4(G) = 2n - 2.$$

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 $n \geq 5$,*

$$T_3(P_5 \,\square\, P_n) = 10n - 20. \tag{92}$$

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

$$g_0(k) = \begin{cases} h, & h \leq k \leq h+1, \\ k+1, & k = 1 \text{ or } h-3 \leq k \leq h-1, \\ k+2, & 2 \leq k \leq h-4, \end{cases} \tag{93}$$

$$g_1(k) = \begin{cases} h+1, & k = h, \\ k+2, & 1 \leq k \leq 2 \text{ or } h-2 \leq k \leq h-1, \\ k+3, & 3 \leq k \leq h-3. \end{cases} \tag{94}$$

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

$$W_h(b, z) = \begin{cases} 0, & b = 0, \\ 1, & 1 \leq b \leq 2, \\ b-2, & 3 \leq b \leq 5, \\ \min\{2b+z-8,\ b+h+z-10\}, & b \geq 6. \end{cases} \tag{95}$$

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

$$b = 0 \quad \text{or} \quad 3 - z \leq b \leq h+1-z. \tag{96}$$

Lemma 19.2 (Odd-length potential inequality)**.** *For $h \geq 12$, $z \in \{0,1\}$, a size satisfying* (96)*, and $0 \leq p \leq 3$,*

$$W_h(b, z) \leq W_h\big(g_z((b-p)_+), 1-z\big) + \mathbf{1}_{\{p \geq 2\}}. \tag{97}$$

For fixed h, z, the function $W_h(b, z)$ is nondecreasing in b.

Proof. Substitute (93)–(95), 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 (96). 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. □

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 (97) 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 $6 \le b \le h-2$, 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 $n = 2q \ge 6$, 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

$$c(A,B,C) = \left|\overline{B} \cap \left(V(B) \cup (A \cap \{1,3\}) \cup (C \cap \{0,2,4\})\right)\right|. \tag{98}$$

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 $s \le 2$, use states (A,B,e), $0 \le e \le s$. 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

$$B_q \cup V(B_q) \cup (B_{q-1} \cap \{1, 3\}) = \varnothing.$$

These checks cover 59 and 1057 reduced words, respectively. They are implemented directly from (98) 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 $q \geq 3$. □

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*

$$k - 1 \leq |R'| \leq k + 2, \qquad |R'| \leq h - 5,$$

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. □

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 $6 \le b \le h-3$. Every nonempty reachable support has size at least two. Define

$$V_h(b,z) = \begin{cases} 0, & b = 0, \\ 1, & 1 \le b \le 2, \\ b-2, & 3 \le b \le 5, \\ \min\{2b+z-8,\ b+h-10\}, & b \ge 6. \end{cases} \tag{99}$$

Here z records history, unlike the physical-parity variable in (95).

For $p \le 3$ 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 $z' = \mathbf{1}_{\{4 \le r \le h-5\}}$; if $z = 1$ and $3 \le r \le h-5$, this is forbidden by Lemma 19.4;
5. $\sigma \ge 3$, giving $b' \ge r+3$ and $z' = 0$.

In the forbidden case, the previous survivor size is $k = b-2$ and $r = k+2-p$, so $k-1 \le r \le k+2$ follows from $0 \le p \le 3$.

Substitution into (99) gives

$$V_h(b,z) \le V_h(b',z') + \mathbf{1}_{\{p \ge 2\}}. \tag{100}$$

This is a finite split of linear integer inequalities under $h \ge 15$, $b \le h$, $b = 0$ or $b \ge 2$, and $z = 1 \Rightarrow 6 \le b \le h-3$. 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 $h\to h-1\to h-2\to h-3$. Then $h-8$ pairs of a stationary round and a one-cell decrease reach size five. The last three rounds are $5\to 4\to 2\to 0$. 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*

$$T_4=2\left\lceil\frac{2h-8}{3}\right\rceil. \tag{101}$$

For $5\leq m<h$, write $2h-6=q(2m-5)+r$, $0\leq r<2m-5$, and set

$$J(0)=0,\quad J(1)=2,\quad J(2)=3,\quad J(3)=J(4)=4,\qquad J(r)=\left\lceil\frac{r+5}{2}\right\rceil\quad (r\geq 5).$$

Then

$$T_m=\begin{cases}2q, & r=0,\\ 2q+1, & r>0 \text{ and } 2J(r)\leq m,\\ 2q+2, & \text{otherwise.}\end{cases} \tag{102}$$

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 14.1 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*

$$(k,\ell)\in\{(h-6,3),(h-5,3),(h-5,4)\}. \tag{103}$$

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

$$e_0 = \min(d_0, d_1), \qquad e_1 = \min(d_0, (d_1 - 1)_+, d_2), \qquad e_2 = \min(d_1, d_2, d_3),$$
$$e_3 = \min(d_2, (d_3 - 1)_+, d_4), \qquad e_4 = \min(d_3, d_4).$$

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

$$\begin{array}{c|l} 3 & (2,1,0,0,0),\ (1,1,1,0,0),\ (1,0,1,0,1) \\ 4 & (2,1,1,0,0),\ (1,1,1,0,1) \\ 5 & (2,1,1,0,1),\ (1,1,1,1,1). \end{array}$$

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. □

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 (103).

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 $m \ge 5$, let F_m be the general potential (36) with $b = 2$. An explicit residue form is

$$D = 2m - 5, \qquad j(u) = \left\lfloor \frac{(u-7)_+}{D} \right\rfloor, \qquad s(u) = ((u-7)_+ \bmod D) + 1,$$

$$F_m(u) = \left\lfloor \frac{u + 5j(u)}{2} \right\rfloor - \mathbf{1}_{\{j(u) \ge 1,\ s(u) \in \{1,2,4\}\}}. \tag{104}$$

The common interior lemma gives the following properties:

1. F_m is nondecreasing and $F_m(u+2) - F_m(u) \le 3$;
2. if $\lfloor u/2 \rfloor \le m$, then $F_m(u) = \lfloor u/2 \rfloor$;
3. if $p \le m$ and $\lfloor u/2 \rfloor > p$, then $F_m(u) \le p + F_m(u - 2p + 5)$.

Indeed, Lemma 8.3 applies because $m \ge b^2 + 1 = 5$. Writing $u = 2k + z$, its target rank is at most $u - 2p + 5$, since $L_0(s) \le s + 2$ and $L_1(s) \le s + 3$. 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 $5 \le m \le h - 8$, set

$$C = \min\{F_m(2h),\ 2 + F_m(2h-2),\ 4 + F_m(2h-4)\}, \qquad \Phi(b,z) = \min\{C, F_m(2b+z)\}.$$

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 $p \ge h - 7 > m$ by (103). 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 $p \ge h - 5$. From an unmarked state, $b \le m + 1$; 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 $b \le m$ 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 $h-4, \ldots, h-1$. Rank increases need only monotonicity. The sole proper-source rank decrease is unmarked $h - 1 \to h - 2$ 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 $T_m \ge \lceil 2C/m \rceil$ throughout $5 \le m \le h - 8$. Direct substitution in the quotient and remainder of (102) gives, for $m \ge 6$,

$$C = qm + J(r) - \mathbf{1}_{\{r=2\}}.$$

Only $m = 5, 6$ remain after the general theorem, and both satisfy $m \le h - 8$ because $h \ge 15$. For $m = 5$, the physical value h is divisible by five: $q = n - 2$, $r = 4$, and $C = 5q + 4$, giving $T_5 \ge 2q + 2 = 2n - 2$. For $m = 6$, the values of $C - 6q$ for $r = 0, \ldots, 6$ 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

$$A_1 = \min\{H(h-1,0), 3+H(h-2,0)\}, \qquad C_4 = \min\{H(h,0), 2+A_1, 4+H(h-2,0)\}.$$

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 14.1. 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 $7 \to 5 \to 3 \to 0$, $8 \to 6 \to 4 \to 0$, and $6 \to 4 \to 0$. 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 (102).

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*

$$T_4 = 2\left\lceil \frac{5n-9}{3} \right\rceil + 2\mathbf{1}_{\{3|n\}}, \qquad T_5 = 2n - 2.$$

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

$$T_m = \begin{cases} 2q, & r = 0, \\ 2q+1, & r > 0 \text{ and } 2J(r) + \delta(r) \leq m, \\ 2q+2, & \text{otherwise.} \end{cases} \tag{105}$$

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 (93)–(94). 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 $0 \leq p \leq 4$,

$$V_h(b,z) \leq p + V_h(g_z((b-p)_+), 1-z). \tag{106}$$

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 (106) decrease by eight. For $b \le 12$, reduce only h by three: target sizes are at most fifteen, no high-end clause is reached, and the caps are inactive since $C(h-3) \ge 48 > H(b', z')$ for $b' \le 15$. Only $12 \le h \le 23$ 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 $T_4 \ge \lceil C(h)/2 \rceil$. For attainment start the majority cohort. The first two rounds give $(h+1,0) \to (h-2,1) \to (h-3,0)$. From majority size $b \ge 9$, two rounds give $(b-2,1) \to (b-3,0)$. The final majority sizes $6, 7, 8$ have tails $6 \to 4 \to 0$, $7 \to 5 \to 3 \to 0$, and $8 \to 6 \to 4 \to 0$. 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 $m \ge 5$, 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

$$L_0(k) = \begin{cases} 0, & k = 0, \\ 2, & k = 1, \\ k+2, & k \ge 2, \end{cases} \qquad L_1(k) = \begin{cases} 0, & k = 0, \\ k+2, & 1 \le k \le 2, \\ k+3, & k \ge 3. \end{cases}$$

Substituting in (33) gives precisely the function J in Theorem 19.5. Substitution in (34) gives

$$\delta(r) = \mathbf{1}_{\{r \in \{1,2,4\} \text{ or } (r \ge 5 \text{ and } r \text{ odd})\}}.$$

Finally $N - w - 1 = 5n - 6$. Thus the general theorem is exactly (105), 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 $n \ge 7$,*

$$T_5(P_7 \,\Box\, P_n) = 2 \left\lceil \frac{7n - 16}{3} \right\rceil.$$

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

$$g_0(k) = k + \min\{R(k), 3, R(H-k) - 1\}, \qquad g_1(k) = k + \min\{Q(k), 4, Q(H-1-k)\}$$

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

$$c_0 = (0, 0, \tfrac{1}{2}, \tfrac{3}{4}, 1, 1), \qquad c_1 = (0, 0, \tfrac{1}{4}, \tfrac{1}{2}, 1, 1).$$

They satisfy $c_0(r) + c_1(s) \le 1$ whenever $r + s \le 5$. 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

$$(p, k) \longrightarrow (1 - p, g_p(\max\{k - r, 0\})), \qquad \text{edge charge } c_p(r), \quad 0 \le r \le 5.$$

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

$$f_p(k) \le c_p(r) + f_{1-p}(g_p(\max\{k - r, 0\})). \tag{107}$$

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 (107) across the current physical colors bounds the total potential loss by one per day. Monotonicity handles neighborhoods larger than the profile minimum. Hence $T_5 \ge \lceil V \rceil$.

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

$$\left(\tfrac{3}{4}, -\tfrac{1}{4}, \tfrac{1}{2}, 0, 0, \tfrac{1}{4}\right). \tag{108}$$

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,

$$\begin{aligned} f_p(k + 3) &= f_p(k) + 2, \qquad && c - 10 \le k \le c + 10, \\ g_p(k) &= k + 3 + p, && c - 15 \le k \le c + 15. \end{aligned} \tag{109}$$

For $n = n_0 + 6j$, write $\Delta = 21j$, so $H = H_0 + \Delta$. Keep f_p unchanged through c. On the inserted interval define

$$\widetilde{f}_p(c + 3\ell + r) = f_p(c + r) + 2\ell, \qquad 1 \le r \le 3, \quad 0 < 3\ell + r \le \Delta.$$

Above $c + \Delta$, put $\widetilde{f}_p(k) = f_p(k - \Delta) + 2\Delta/3$. The period identity makes this a nondecreasing extension.

We check (107) 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 (109) 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 (108) 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

$$\begin{aligned}
&A_0 = B_0 = 0,\\
&\quad A_{t+1} = B_t + m - \min\{q(B_t+m), b\},\\
&\quad B_{t+1} = A_t + m - \min\{R(A_t+m), b+1\}.
\end{aligned}\tag{110}$$

Here $q(z)$ is the least nonnegative integer r with $z \le r(r+1)$. Let L_b be the first time $B_t = C$, with $L_1 = 0$.

Theorem 21.1. *For every odd $n \ge w = 2b+1 \ge 3$,*

$$T_{b+1}(P_w \,\square\, P_n) = 2wn - C_w, \qquad C_w = 8b^2 - 4b - 4L_b + 4. \tag{111}$$

The budget $b+1$ is the minimum feasible budget, and the recurrence (110) *terminates independently of n. Its clock can be evaluated in $O(b)$ arithmetic stages using Proposition 6.4. This width-dependent recurrence does not meet the absolute $O(1)$ numerical completion criterion when b varies. The following finite list does give fixed numerical constants for its stated widths. 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 (110) 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 $A \ge C+1$ requires the preceding $B \ge C$. 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 $b \ge 2$. 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

$$(s, s) \longmapsto (s+1, s) \longmapsto (s+1, s+1).$$

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

$$a_* = M - b^2 - 1, \qquad t_* = L_b + 2(a_* - C).$$

The inverse formulas give these affine transitions through the tie (a_*, a_*). It lies after $C+1$ since $a_* - (C+1) \ge 3b - 2 > 0$. 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

$$k \le m + D_p(T-1), \tag{112}$$

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 (111).

Finally, the rectangular-grid feasibility theorem [1, Theorem 2] gives minimum budget $b+1$. The clock also satisfies $L_b = \lceil \mu_1(C)/(b+1) \rceil$ by Proposition 7.6. Equivalently, (112) 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. □

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 The same expression as an upper bound at even lengths

Proposition 21.2 (All-length minimum-budget construction). *For $w = 2b + 1 \geq 3$, even $n \geq w$, and $k = b + 1$, the constant C_w of Theorem 21.1 gives*

$$T_{b+1}(w, n) \leq 2wn - C_w.$$

Together with that theorem, the construction is valid at every length. This upper bound can be strict: in width 25, sharing the final inspection day improves it from $50n - 924$ to the exact value $50n - 925$ for every even $n \geq 26$. Thus equality of the displayed expression throughout all even lengths and widths is false. The general exact minimum-budget classification remains open, and evaluating C_w in fixed size for arbitrary width remains a separate unresolved question.

Proof. Use coordinates $0 \leq x < n$, $0 \leq y < w$, and order each color by $(x + y, x)$, ascending. Write $I_p(a)$ for its a-room prefix and set $h = wn/2$. These particular prefixes have nested neighborhoods, with

$$g_0(a) = a + \min\{R(a), b, \rho(h-a)\}, \qquad g_1(a) = a + \min\{Q(a), b+1, R(h-a)\}$$

for $a > 0$, and $g_p(0) = 0$. Here $R(a) = \lceil \sqrt{a} \rceil$, $\rho(a)$ is the least $u \geq 0$ with $a \leq u(u+1)$, and $Q(a)$ the least $u \geq 1$ with $a \leq u(u-1)$, with $Q(0) = 0$. The two oriented profiles differ even though reflection makes the unrestricted physical profiles equal.

For completeness, let $\ell(d)$ be diagonal $x+y=d$'s length and define $J(0)=0,\ J(d+1)=\ell(d)-J(d)$. A prefix ending after j rooms on diagonal d has surplus

$$J(d)+1-\mathbf{1}_{\{d\geq 2b\}}-\mathbf{1}_{\{\max(0,d-2b)+j-1=n-1\}}.$$

Following the two forward neighbors of the terminal interval proves neighborhood nesting. The diagonal lengths increase, plateau, and decrease; the recursion gives

$$\begin{aligned}
J(d) &= \lceil d/2\rceil, && 0\leq d\leq 2b,\\
J(d) &= b+(d \bmod 2), && 2b\leq d\leq n-1,\\
J(n+u) &= b-\lfloor u/2\rfloor, && 0\leq u\leq 2b.
\end{aligned}$$

Substituting the terminal partial diagonal proves the two profiles, including the far-end correction. Their integer inverses satisfy $\iota_p(z)=h-g_p(h-z)$.

Inspect the final $\min(k,a)$ rooms of the current prefix each day. Its next size is $g_p((a-k)_+)$. Until capture, the deficit $h-a$ therefore follows the alternating inverse maps $\iota_p(z+k)$. Put $C=b(b-1)$, $B=b(b+1)$, and $a_*=h-B-1$. Through the tie (a_*,a_*), every inverse input is at most $a_*+k=h-b^2$, so the competing far-corner terms are inactive. The width-clock argument already proved above gives this tie at

$$t_*=L_b+2(a_*-C)=L_b+2h-4b^2-2.$$

Here $a_*\geq C+1$, since $h\geq(2b+1)(b+1)$. Both initial phases now have prefix size $B+1$, in opposite current colors.

Throughout the remaining search, sizes are at most $B+1$ and the far corner stays inactive, since $h-(B+1)\geq b^2+2b$. The same inverse-clock equivalence used in (112) says that a prefix of size $B+1$ clears in $s+1$ days when its fresh inverse clock reaches $B+1-k=b^2$ at time s. The first entry reaches this at L_b+2b-1, so choose the initial phase whose tail lasts L_b+2b days. Its full solo duration is

$$\tau=t_*+L_b+2b=wn-4b^2+2b+2L_b-2.$$

The inspection list followed by its reversal captures both initial colors in 2τ days: an avoiding walk in the second half, reversed, would contradict the first half's solo guarantee. This is a physical construction and uses no assumed optimality of the oriented prefixes. □

22 Certificate and affine-insertion tools

22.1 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; independent applications are retained in the companion.

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

$$c_0(r)+c_1(m-r)\leq 1. \tag{113}$$

Suppose nonnegative functions F_p, with $F_p(0)=0$, satisfy

$$F_p(k)-F_{1-p}\big(g_p((k-r)_+)\big)\leq c_p(r)\quad (0\leq k\leq H-p,\ 0\leq r\leq m). \tag{114}$$

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 (113). 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

$$T_m \geq \lceil F_0(H) + F_1(H-1) \rceil . \tag{115}$$

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 (113)–(114) directly with exact fractions.

22.2 Inserting an arbitrarily long affine interval

Lemma 22.1 (Affine insertion)**.** *Let $A = b(b+1)$. Suppose a certificate at majority size H has an integer cut c such that*

$$c \geq A + 2D, \qquad H - c \geq A + 2D + 2, \tag{116}$$

$$F_p(k) = 2k + \alpha_p \qquad (c - 2D \leq k \leq c + 2D,\ p \in \{0, 1\}). \tag{117}$$

For every integer $\Delta \geq 0$, replace H by $H + \Delta$ in the profile formulas and define

$$F'_p(k) = \begin{cases} F_p(k), & k \leq c, \\ F_p(c) + 2(k-c), & c \leq k \leq c + \Delta, \\ F_p(k-\Delta) + 2\Delta, & k \geq c + \Delta. \end{cases} \tag{118}$$

Then F'_p satisfies the same charge inequalities. The initial lower bound (115) *increases by exactly 4Δ.*

Proof. The pieces of (118) agree at their endpoints. For either the base or enlarged profiles, every successor $j = g_p((k-r)_+)$ satisfies

$$|j - k| \leq D. \tag{119}$$

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 (116). The successor is unchanged and is less than c by (119). 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

$$q \geq c - D - m \geq A + 1, \qquad H + \Delta - q \geq H - c - D \geq A + D + 2.$$

Both profiles are therefore on their affine plateaus: $j = k - r + \delta_p$, where $\delta_0 = b$ and $\delta_1 = b + 1$. By (119), 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 (118) adds 2Δ to the potential, giving the stated increase of 4Δ in the integer lower bound. □

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 22.2 (Insertion in a serial strategy)**.** *Under* (116)*, 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

$$g_0(k - m) = k - 1, \qquad g_1(k - m) = k.$$

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 (116) 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. □

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.

23 Eventual affine periods on odd rectangles

This is the rectangle specialization of the previously proved corner-profile and interval-insertion argument. Its canonical proof is shared with the parked all-odd-box application. An affine plateau alone would not justify changing the rectangle's length.

Let $Q = P_W$ be an odd path with $W \geq 3$ vertices. Put

$$S = W - 1, \qquad Z = W(S+1) = W^2, \qquad \beta = (W-1)/2.$$

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 compatible order of Theorem 5.1 applies. Equivalently, put the short coordinate first and use decreasing lexicographic order within a fixed weight layer. Write $H = (Wn+1)/2$ for its majority class size and g_p^n for its exact color-p profile.

23.1 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 $Q \times \{0, 1, \ldots\}$. 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 $r_p(v) \leq Z$.

Index the path vertices by $0, \ldots, W-1$. Its transverse prefix cost is $a_Q(v) = 1$ for $v < W-1$ and $a_Q(W-1) = 0$. Define

$$U_p(k) = \sum_{\substack{v \in Q_p \\ r_p(v) \leq k}} a_Q(v), \qquad B_p(k) = \left|\{v \in Q_p : r_p(v) \leq k\}\right|.$$

Both tables are constant for $k \geq Z$.

Lemma 23.1 (Stable corner decomposition)**.** *For $n \geq S+3$ and $k > 0$,*

$$g_p^n(k) = k + \mathbf{1}_{\{p=1\}} + U_p(k) - |Q_p| + B_p(H - p - k), \qquad g_p^n(0) = 0. \tag{120}$$

In particular, every surplus is at most $\beta + p$, and

$$k > Z, \quad H - p - k > Z \quad \Longrightarrow \quad g_p^n(k) = k + \beta + p. \tag{121}$$

Proof. The rectangle prefix-layer calculation in (22) assigns excess contribution $a_Q(v)$ on the bottom slice, zero on an interior slice, and -1 on the top slice. This includes the origin: its excess contribution is the transverse origin cost. The prefix-layer 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 (120). 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 (121). □

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

$$H-p-k > Z \implies g_p^{H+\delta}(k) = g_p^H(k), \tag{122}$$

$$k > Z \implies g_p^{H+\delta}(k+\delta) = g_p^H(k) + \delta. \tag{123}$$

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 odd rectangle has a spanning path whose two endpoints lie in its majority color: traverse successive rows alternately forward and backward. 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 rectangle contains these path edges, so

$$g_0(k) \geq k \ (k < H), \qquad g_1(k) \geq k+1 \ (k > 0), \qquad g_0(H) = H-1. \tag{124}$$

Corollary 23.2 (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 (121). Apply the rectangular-grid feasibility theorem [1, Theorem 2]. □

The same budget threshold holds for every $n \geq W$ on rectangles by the rectangle feasibility theorem. The conservative hypothesis here is kept for the shared insertion argument.

23.2 The period and an explicit threshold

Fix a feasible eventual budget $m \geq \beta+1$ and set

$$d = 2m - W > 0, \qquad g = \gcd(W, d), \qquad \Delta = \operatorname{lcm}(W, d) = Wd/g.$$

For the explicit threshold define

$$\begin{aligned}
&J = m+1, && A = Z+1, && R = \Delta + 2J,\\
&B = 8Z + 6d + 2, && L = Z+m+1, && \mathcal{E} = B + 4L(B+1),\\
&H_* = R + 2A + 4J + 2 + 2\mathcal{E}(R+2J+1).
\end{aligned} \tag{125}$$

Theorem 23.3 (Eventual affine period under the preceding profile hypotheses)**.** *For every odd n with $(Wn+1)/2 \geq H_*$,*

$$T_m(Q \,\square\, P_{n+2d/g}) = T_m(Q \,\square\, P_n) + 4W/g. \tag{126}$$

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.

23.3 A potential with a bounded total deficit

For $n \geq S+3$, write $H = (Wn+1)/2$ and $M = H-1$. Define

$$K = 2Z + 2m + 1, \qquad \Gamma = (2H - 4Z - 2m - 5)_+, \qquad F_p(k) = \min\{\Gamma, (2k + p - K)_+\}.$$

Lemma 23.4 (Uniform potential and time bounds)**.** *For $0 \leq r \leq m$,*

$$F_p(k) - F_{1-p}\big(g_p((k-r)_+)\big) \leq (2r - W)_+ . \tag{127}$$

Consequently

$$\lceil 2\Gamma/d \rceil \leq T_m(Q \,\square\, P_n) \leq 4 \left\lceil (M-m)_+ /d \right\rceil + 2. \tag{128}$$

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 (127).

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 $2\left\lceil (M-m)_+ /d \right\rceil + 1$ 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. □

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 $\Gamma = \big(2H - 4b^2 - 2b - 6\big)_+$. 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

$$\max\{0,\, 2wn - 2w^2 + 2w - 10\} \leq T_{(w+1)/2}(P_w \,\square\, P_n) \leq 2wn - 2w - 2. \tag{129}$$

Thus its leading term is $2wn$ for every fixed odd width, independently of whether a sharper correction has been determined.

Return to the preceding cylinder family, with $Z = W(S+1)$ and the potential of Lemma 23.4. Assume $H \geq H_*$, so $\Gamma > 0$ and $M > m$. Fix an optimal prefix strategy, using padded full quotas as in Section 2. 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 (128) gives

$$\sum_t \eta_t = dT_m - 2\Gamma \leq 8Z + 6d + 2 = B. \tag{130}$$

Here $d\lceil (M-m)/d\rceil \le M-m+d-1$ 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 (127) 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.

23.4 From flat potentials to empty or full cohorts

Zero inspections strictly increase every intermediate potential:

$$0 < F_p(k) < \Gamma \quad \Longrightarrow \quad F_{1-p}(g_p(k)) > F_p(k). \tag{131}$$

For a proper majority prefix, $g_0(k) \ge k$, 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, $g_1(k) \ge k+1$, 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, (131) 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

$$\mathcal{E} = B + 4L(B+1) \tag{132}$$

nonclean days.

23.5 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

$$A + 2J \le c \le H - R - A - 2J - 2. \tag{133}$$

There are $H - R - 2A - 4J - 1$ candidates. By (125), 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

$$c + \Delta < x_i \le c + \Delta + J. \tag{134}$$

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

$$\ell_i = x_i - \Delta, \qquad u_i = x_i.$$

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.

23.6 Removing and enlarging the crossing blocks

Proof of Theorem 23.3. 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

$$f_i(k) = \begin{cases} k, & k \le \ell_i, \\ k - \Delta, & k \ge u_i. \end{cases}$$

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

$$(H - \Delta) - \ell_i = H - u_i \ge H - (c + R) \ge A + 2J + 2,$$

so the low profile identity (122) 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 (123) 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

$$T_m(Q \,\square\, P_{n-2\Delta/W}) \leq T_m(Q \,\square\, P_n) - 4\Delta/d. \tag{135}$$

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 (135) at $H + \Delta$ for the reverse inequality. Since $\Delta/W = d/g$ and $\Delta/d = W/g$, this proves (126). $\square$

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

$$n_0, n_0 + 2, \ldots, n_0 + 2d/g - 2.$$

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 (129) give $2w + 2 \leq C(w) \leq 2w^2 - 2w + 10$. The offset recurrence and its parameter-dependent finite initialization are not a fixed-size numerical expression. Their reduction remains part of the rectangle objective when $Q = P_w$.

24 A height construction used by rectangle bounds

We retain this construction in its natural generality because its path-cross-section case supplies the upper bounds used in the rectangle interface proofs. Independent higher-dimensional applications are parked in the companion.

Theorem 24.1. *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 $0 \le z < n$. 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 $a'_u \le a'_v + 1$.

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

$$\max_v a_v(t) \le n - 1 + t - 2\lfloor mt/A \rfloor < n + 1 - t(2m - A)/A.$$

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. $\square$

25 A growth bound used by rectangle interfaces

The exact offsets remain sensitive to geometry, but the leading term of the optimal time has a uniform answer across all side parities.

Theorem 25.1. *Let H be a finite bipartite graph with $A \ge 2$ vertices and a Hamiltonian path. For each fixed integer $m > A/2$, as $n \to \infty$ through either parity,*

$$T_m(H \,\square\, P_n) = \frac{2An}{2m - A} + O_{A,m}(1). \tag{136}$$

The height strategy in Theorem 24.1 has a bounded additive excess over the optimal time. When A is even, the following bounds hold explicitly for every $n \ge 2$:

$$\frac{2\left(An/2 - A^2 - m\right)_+}{m - A/2} \le T_m(H \,\square\, P_n) \le 2\left\lceil \frac{A(n+2)}{2m - A} \right\rceil. \tag{137}$$

In particular the theorem applies to every Cartesian transverse box.

Lemma 25.2 (An interior expansion bound for even-width rectangles)**.** *On $P_{2b} \,\square\, P_n$, $b \ge 1, n \ge 2$, put $h = bn$ and $K = 2b^2$. Every one-color set R with $K < |R| < h - K$ satisfies $|N(R)| \ge |R| + b$.*

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

$$|V(D) \setminus C| \le b(2s + 1) < 2b^2 = K.$$

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 $|R| \le K$; in the second $h - |R| \le K$, proving the contrapositive. $\square$

Proof of Theorem 25.1. First let $A = 2b$ and consider the spanning rectangle $P_A \mathbin{\square} P_n$. Put $C = (h - 2K - m)_+$, $L = K + m$, and

$$\psi(k) = \min\{C, (k - L)_+\}.$$

If a cohort of size k receives $p \le m$ useful inspections, write $r = k - p$ and $k' = |N(R)|$ for its survivor size and actual next size. The matching gives $k' \ge r$ for every survivor set. If $r \le K$, the source potential is zero. If $r \ge h - K$, the target potential is C whenever $C > 0$, while $C = 0$ is trivial. In the remaining case Lemma 25.2 gives $k' \ge k - p + b$. Since clipping is monotone and has Lipschitz constant one, all cases give

$$\psi(k) - \psi(k') \le \max\{0, p - b\}.$$

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 $T_m \ge 2C/(m - b)$ on the rectangle. The rectangle is a subgraph of $H \mathbin{\square} P_n$, giving the same lower bound there. Theorem 24.1 gives the displayed upper bound, with the same coefficient of n.

For odd $A \ge 3$, take the largest odd $\ell \le n$. The spanning rectangle $P_A \mathbin{\square} P_\ell$ 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 $n - \ell \le 1$, this gives the required lower bound on the whole cylinder. The same height upper bound completes the proof. $\square$

The even-width expansion lemma and its potential proof will be used in the rectangle interface construction. General-cylinder applications are recorded in the companion.

26 Efficient fronts and eventual periods for rectangle proofs

This section retains the physical efficient-frontier and paired-column arguments used by the exact rectangle interface theorem. They are stated for Hamiltonian cross-sections because the proofs have that natural generality. The independent synthesis for every transverse Cartesian box is preserved in the parked companion. For the rectangle objective take $Q = P_w$ throughout. Effective periods and finite preprocessing remain intermediate results under the absolute $O(1)$ numerical completion goal.

Theorem 26.1. *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 $\Delta = 2d/\gcd(b, d)$,*

$$T_m(Q \,\square\, P_{n+\Delta}) = T_m(Q \,\square\, P_n) + \frac{2b\Delta}{d}, \qquad n \geq N. \tag{138}$$

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 25.2, 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 $|N(R)| - |R| = |B|$.

26.1 Equality in the interior expansion bound

Lemma 26.2 (A finite family of minimum-surplus frontiers)**.** *Put $K = 4b^2$. Suppose a one-color set R in $Q \,\square\, P_n$ satisfies*

$$K < |R| < h - K, \qquad |N(R)| = |R| + b.$$

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

$$R_i = \{0, \dots, c_i - 1\} \quad \textit{for every } i, \qquad \textit{or} \qquad R_i = \{c_i + 1, \dots, n-1\} \quad \textit{for every } i.$$

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 25.2 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 25.2 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

$$|R_i \setminus R_{i+1}| \leq 4, \qquad |R_{i+1} \setminus R_i| \leq 4.$$

An empty row would now imply $|R| \le 4b^2$, by propagating the bound through all rows. A row full except at its boundary similarly implies $h - |R| \le b(1 + 4(b-1)) \le 4b^2$. 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 $e_i \le c_i + 1$, and its useful inspection count is

$$p = \sum_{i=0}^{b-1} (c_i + 1 - e_i). \tag{139}$$

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.

26.2 Only boundedly many days can behave differently

Consider an optimal strategy on a sufficiently long cylinder, and set

$$C = h - 2K - m > d, \qquad \psi(k) = \min\{C, (k - K - m)_+\}.$$

The argument proving Theorem 25.1, with the larger cutoff $K = 4b^2$, gives for a cohort receiving p useful inspections and changing from count k to k',

$$\psi(k) - \psi(k') \le \max\{0, p - b\}. \tag{140}$$

The two cohorts start with total potential $2C$ and end with zero. Their combined daily decrease is at most d. The height upper bound $T_m \le 2\lceil b(n+2)/d \rceil$ consequently gives

$$dT_m - 2C \le B, \qquad B = 4K + 4m + 2b - 2. \tag{141}$$

Indeed the integer ceiling bound gives $dT_m \le 2b(n+2) + 2d - 2$.

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 (141), 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 (140) 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 26.2.

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

$$E = B + 2L(B+1). \tag{142}$$

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 (139), its average cut moves at the constant speed d/b toward clearance. Individual cuts need not move monotonically.

26.3 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

$$b(E+1) + mE + 2 \leq F, \qquad F = (b+m)(E+1) + 2$$

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

$$n > F(2\rho + 1) + (F+1)(W + 2v + 4) \tag{143}$$

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.

26.4 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 $a_u \le a_v + 1$, and reversing the edge gives the opposite inequality. Their parities differ, so

$$|a_u - a_v| = 1 \qquad \text{on every edge } uv \text{ of } Q. \tag{144}$$

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 (144) has neighborhood cutoffs exactly $a_v + 1$: longitudinal movement supplies them and transverse edges supply nothing larger. Right frontiers use reflected coordinates.

Lemma 26.3 (Height endpoints)**.** *Let Q be connected and bipartite, with $A \ge 2$ vertices and diameter R. Fix $m > A/2$ and put $D = 2m - A$. Let a, b be integer height configurations satisfying* (144)*. If*

$$t \ge \lceil AR/m \rceil, \qquad b_v \equiv a_v + t \pmod 2 \quad \textit{for every } v, \qquad \sum_v a_v - \sum_v b_v = Dt,$$

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

$$[\mu(b) - R,\ \mu(a) + R + 1], \qquad \mu(x) = A^{-1} \sum_v x_v.$$

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 13.2. 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

$$c_v - a_v \le -2mt/A + 2R \le 0.$$

Thus $b \le a + t\mathbf{1}$ 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 26.3 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,

$$\ell_0 = \frac{2b}{\gcd(b,d)}, \qquad \Delta = \frac{2d}{\gcd(b,d)}, \qquad K_0 = \left\lceil \frac{2b(2b-1)}{m} \right\rceil. \tag{145}$$

Here are conservative explicit margins that guarantee such a path well inside the protected interval. Set

$$\begin{aligned} U &= K_0 + \ell_0 + 4, \\ W &= 2(\Delta + v + m + 3) + \left\lceil \frac{d(U+4)}{b} \right\rceil + 2v + 2m + 10. \end{aligned} \tag{146}$$

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

$$N = 4 + \left\lceil \frac{2K + m + d + 1}{b} \right\rceil + F(2\rho + 1) + (F+1)(W + 2v + 4) + \Delta. \tag{147}$$

This ensures $C > d$, the protected-gap condition, and the same conditions after contracting by Δ. All constants depend only on b and m.

26.5 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 26.3. 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

$$T_m(Q \,\square\, P_{n-\Delta}) \leq T_m(Q \,\square\, P_n) - 2\ell_0. \tag{148}$$

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 26.3 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

$$T_m(Q \,\square\, P_n) \leq T_m(Q \,\square\, P_{n-\Delta}) + 2\ell_0. \tag{149}$$

Combining (148) and (149), and using $2\ell_0 = 2b\Delta/d$, proves Theorem 26.1 for every $n \geq N$ after renaming the shorter length n. The even displacement preserves either longitudinal parity throughout.

26.6 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 26.4. *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 26.5 (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 $s \le b$. 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,

$$R_y \subseteq R_x, \qquad R_x \setminus R_y \subseteq \{c_y\}.$$

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

$$R_y = \{0, \dots, c_y - 1\}, \qquad k \in \{c_y, c_y + 1\}.$$

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. □

Lemma 26.6 (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 $|S| \le A(b+1) < A^2$, 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 26.5, with X, Y exchanged, proves that the X fibers are prefixes as well, with the stated length relation. Reflection gives the other orientation. □

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.

26.7 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

$$(2k_t + e_t) - (2k_{t+1} + e_{t+1}) = 2p - 2s + e_t - e_{t+1}. \tag{150}$$

For a middle survivor with $s = b$, Lemma 26.5 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

$$L_0 = 2(K + m) + 1, \qquad C = 2h - 4K - 2m - 1 > D, \qquad \psi_t = \min\{C, (2k_t + e_t - L_0)_+\}.$$

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 (150) therefore give, in every case,

$$\psi_t - \psi_{t+1} \leq \max\{0, 2p - A\}. \tag{151}$$

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

$$DT - 2C \leq B, \qquad B = 8K + 8m + 2A.$$

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 (150). The only cases are

$$(s, e_t, e_{t+1}) = (b, 0, 1) \quad \text{or} \quad (b + 1, 1, 0). \tag{152}$$

The first survivor is a frontier by Lemma 26.5. In the second case the preceding survivor was a middle minimum-surplus frontier, so the actual containment in its neighborhood permits Lemma 26.6. 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 26.2 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

$$E = B + 2L_1(B + 1) \tag{153}$$

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 $2k + e$:

$$(2r - e) - (2r' - e') = D \tag{154}$$

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 (154) is therefore the same physical height drift used in Lemma 26.3.

Apply that lemma with $g = \gcd(A, D)$ and the conservative diameter bound $A - 1$. We may use

$$K_0 = \left\lceil \frac{A(A-1)}{m} \right\rceil, \qquad \ell_0 = 2A/g, \qquad \delta = D/g, \qquad \Delta = 2\delta. \tag{155}$$

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 26.3 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

$$\begin{aligned} U &= K_0 + \ell_0 + 4, \\ W &= 2(\delta + v + m + 3) + \left\lceil \frac{D(U+4)}{2A} \right\rceil + 2v + 2m + 10, \\ L_* &= 4 + \left\lceil \frac{4K + 2m + 1 + D}{2A} \right\rceil + F(2\rho + 1) + (F+1)(W + 2v + 4) + \delta, \end{aligned} \tag{156}$$

and taking $N = 2L_*$. The corrected mean moves by $D/(2A) \le m$ 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 26.3. 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

$$T_m(Q \,\square\, P_{n-\Delta}) \le T_m(Q \,\square\, P_n) - 2\ell_0.$$

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 26.5. Since $2\ell_0 = 2A\Delta/D$, this proves Theorem 26.4 on every sufficiently long even cylinder, with the displayed effective threshold.

27 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 27.1 (An exact finite-interface representation)**.** *Let Q be a fixed finite bipartite graph with a Hamiltonian path, with $A \ge 2$ 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 (78) *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 hidden dependence on Q, m, including finite preprocessing and shortest paths, means that they do not meet the absolute $O(1)$ operation goal for $T_k(a, b)$. 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

$$\begin{aligned} K &= A^2, & B &= 8A^2 + 8m + 2A, \\ E &= B + 2(K + m + 1)(B + 1), & L &= E + 1, \\ q &= (m + \alpha)L, & R &= 100(q + A^2 + L + A + 1), \qquad n_0 = 10R. \end{aligned} \tag{157}$$

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 26, and gives $An/2 > mL$.

27.1 Physical pyramids, including the end boundaries

Index Q by a Hamiltonian order $0, \ldots, A - 1$, so its bipartition is the parity of this index. A downward pyramid is a one-color set P such that

$$(v, y) \in P, \quad y > 0, \quad uv \in E(Q) \quad \Longrightarrow \quad (u, y - 1) \in P.$$

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

$$-2 \leq a_v \leq n - 1, \qquad a_v \equiv p - v \pmod 2, \qquad |a_u - a_v| = 1 \quad (uv \in E(Q)). \tag{158}$$

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

$$\eta(a)_v = \max(h_v, \{a_u : uv \in E(Q)\}), \qquad h_v = \begin{cases} (1-s_v) - 2, & k_v = 0, \\ \min(a_v + 1, t'_v), & k_v > 0. \end{cases}$$

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

$$0 \leq |N(P)| - |P| \leq \alpha, \qquad |P| \leq |N^r(P)| \leq |P| + \alpha r. \tag{159}$$

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 12.2 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.

27.2 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 26 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 26.2 and (144). For odd $A = 2b + 1$ and even n, the history-bit potential has daily bound $2m - A$. Its two efficient cases are exactly (152). Surplus b gives a pyramid by Lemma 26.5; surplus $b+1$ also gives a pyramid by the actual containment hypothesis in Lemma 26.6. Thus both efficient phases have the physical height condition (158) on every Q-edge.

The B in (157) 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.

27.3 Localizing a physical transition

Lemma 27.2 (Exact endpoint localization)**.** *Suppose an r-day block, $r \le L$, 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 $|D| \le (m+\alpha)r$. 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 $|P_1| \ge |P_0| - mr$, while (159) 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 $|P_1| \le A^2$, whence $|P_2| \le A^2 + (m+\alpha)r$; 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. □

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.

27.4 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 27.2 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 $r \le L$ 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 $n \ge 10R$ 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 $1, \ldots, L$. 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.

27.5 Finite ports and exact symbolic interior costs

Use the common homogeneous counter interval

$$R \le z \le \lfloor (n - 2 - 2R)/2 \rfloor .$$

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 $n \ge n_0$.

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 (78), 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 $r \le L$. The large collar in (157) implies the height margins in (76), 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 13.3 therefore bounds that edge against arbitrary intermediate supports by

$$z_v \le a_v + r, \qquad \sum_v a_v - \sum_v z_v + Ar \le 2mr.$$

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 (78). 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 (77) holds. Proposition 13.4 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 13.4, 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 27.1.

28 The inverse rectangle tradeoff at 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.

For rectangles, define $K_t(a, b) = \min\{k : T_k(a, b) \le t\}$. This is the inverse of the same constant-budget capture relation, so its rectangle consequences remain inside the main scope. For fixed a, t, the theorem below with $Q = P_a$ gives $K_t(a, n) = an/t + c_{n \bmod p}$ eventually, together with effective finite exceptions. The period, coefficients and exception table require parameter-dependent computation; this does not meet the absolute $O(1)$ numerical goal when all three parameters vary. The proof uses daily cost vectors as a tool, not as a new varying-budget research objective.

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 $n \ge 2$ and $T \ge 1$.

Theorem 28.1. *For fixed Q, T, there are effectively computable integers $n_0, p > 0$ and rational* *constants $c_0, \ldots, c_{p-1}$ such that*

$$b_T(n) = \frac{An}{T} + c_{n \bmod p} \qquad (n \geq n_0).$$

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.

28.1 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 $b = (b_1, \ldots, b_{T-1})$ records proposed survivor sets in one column. Put $b_T = \varnothing$, and let 0 denote the all-empty letter. For $a, b, c \in \mathcal{W}$ define

$$\begin{aligned} w_1(a,b,c) &= A - |b_1|, \\ w_t(a,b,c) &= \left|(N_Q(b_{t-1}) \cup a_{t-1} \cup c_{t-1}) \setminus b_t\right| \quad (2 \leq t \leq T). \end{aligned} \tag{160}$$

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 28.2. *A successful T-day strategy with respective daily capacities $m_1, \ldots, m_T$ 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 $R_T = \varnothing$. 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 (160).

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 $(a, b) \to (b, c)$ 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

$$F_n(\mathbf{x}) = u^{\mathsf{T}} M(\mathbf{x})^n v$$

has a positive coefficient at $\mathbf{x}^{\mathbf{c}}$ precisely when a certificate with exact cost vector $\mathbf{c}$ exists. Thus

$$\sum_{n \geq 0} F_n(\mathbf{x}) z^n = u^{\mathsf{T}} (I - zM(\mathbf{x}))^{-1} v$$

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.

28.2 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 $(n, c_1, \ldots, c_T)$ 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

$$E(n, m) \quad \Longleftrightarrow \quad n \geq 2 \text{ and } \exists \mathbf{c} \, ((n, \mathbf{c}) \in L \text{ and } c_t \leq m \text{ for all } t)$$

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 $(1, 2), (3, 4), \ldots$ 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

$$T b_T(n) \geq 2A \lfloor n/2 \rfloor \geq A(n-1).$$

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 28.1, 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.

29 Geometric memory and bounds at intermediate budgets

29.1 Exact backward frontiers with three or nine thresholds

The necessary models above need not be evaluated by listing their size states. Their entire backward reachable sets have the following finite dimensional descriptions. These are exact statements about the models; physical realization and evaluation without iteration remain separate.

Write $R(x) = \lceil \sqrt{x} \rceil$ and $\rho(x) = \lceil (\sqrt{1+4x} - 1)/2 \rceil$ for $x \geq 0$. For the corner capacities F_{ij}, G_{ij} in (49) and (60), let $S_{ij}(t)$ be the least surplus with $t \leq s + F_{ij}(s)$, and $H_{ij}(d)$ the least surplus with $d \leq G_{ij}(s)$. Their explicit values are

$$S_{ij}(t) = \begin{cases} i + 2j - 2 + \max(2, R((t - 2i - 4j + 4)_+)), & j > 0, \\ i - 1 + \max(1, \rho((t - 2i + 2)_+)), & j = 0 < i, \\ \min(\rho(2t), \rho(t) + 1), & i = j = 0 < t, \\ 0, & i = j = t = 0. \end{cases}$$

Put $A = 2 - i$, $B = 2 - j$. Then

$$H_{ij}(d) = \begin{cases} A + B + \max(0, R((d - B + 1)_+) - 1), & B > 0, \\ A + \rho(d), & B = 0 < A, \\ \min(\rho(2d) + 1, 1 + R(d + 1)), & A = B = 0 < d, \\ 0, & A = B = d = 0. \end{cases}$$

These formulas include the mandatory corner domains of their quadratics.

Theorem 29.1 (Three-threshold frontier)**.** *Let $w \leq n$, wn even, $w \geq 2$, $h = wn/2$, $r = \lfloor w/2 \rfloor$, and $k \geq r + 1$. Use the corner model of Theorem 11.20, optionally with its critical refinement. Its states capturable within t days are exactly*

$$\{(a, c) : c \leq a \leq \min(h - 2 + c, B_t(c))\}, \qquad c = 0, 1, 2.$$

Initialize $B_0 = (0, -1, -1)$. For a valid positive target size z of corner class j, set

$$s_{ij}(z) = \min(S_{ij}(z), H_{ij}(h - z), \kappa_{ij}), \qquad q_{ij}(z) = \min(h - 2 + i, z - s_{ij}(z)).$$

Set $Q_{ij}(z) = q_{ij}(z)$ when $q_{ij}(z) \geq \max(1, i)$, and zero otherwise; also set it to zero if $z < \max(1, j)$. In the original model $\kappa_{ij} = r$. In the refined model it is $r+1$, except that it equals r for $(i, j) = (1, 1)$ when a side equals $2r$, or $(i, j) = (2, 0)$ when a side equals $2r + 1$. Then

$$B_{t+1}(c) = \min\left(h - 2 + c, k + \max_{i \leq c,\ 0 \leq j \leq 2} Q_{ij}(B_t(j))\right). \tag{161}$$

Proof. For a fixed target and corner pair, admitted positive survivor sizes form the interval $[\max(1,i), Q_{ij}(z)]$. Increasing surplus weakens each component inequality, and their union with the unrestricted range is the interval starting at $s_{ij}(z)$. The static size profile is redundant: both component capacities are bounded by s^2, and the unrestricted surplus is at least r.

Moreover $Q_{ij}(z)$ is nondecreasing in z. The inverse S_{ij} increases by at most one at each integer increment, since $s + F_{ij}(s)$ increases by at least one. The inverse $H_{ij}(h-z)$ is nonincreasing. Thus $z - S_{ij}(z)$, $z - H_{ij}(h-z)$ and $z - \kappa_{ij}$ are nondecreasing. Taking their maximum and imposing the state cap and minimum positive survivor size preserves this property.

Induct on t. A nonterminal first day leaves at most $Q_{ij}(B_t(j))$ rooms, proving necessity of (161). Conversely choose indices attaining a positive maximum Q. Let $\ell = c - i \le 2 \le k$. For any valid source size $k < a \le B_{t+1}(c)$ take

$$q = \min(Q, a-\ell), \qquad p = a - q.$$

Then $q \ge \max(1,i)$ and $\ell \le p \le k$. The fixed-target interval therefore supplies the transition to $B_t(j)$. The survivor cap is consistent with all these sources because $(h-2+c) - \ell = h - 2 + i$. Sources of size at most k clear directly. This proves the entire interval assertion, including its model converse. □

The same proof works with the initial set consisting of all states of size at most a prescribed v, initialized by $B_0(c) = \min(v, h-2+c)$, or -1 if this is below c. It then describes reaching that set in at most t moves. In particular $v = \lfloor k/2 \rfloor$ tests the residual relevant to the central inspection. The graph audit in `research/bridge-corner-frontier-check.json` compares all state/deadline memberships in 232 cases, with 622731 agreements against an independent pre-existing graph implementation.

Theorem 29.2 (Nine-threshold erosion frontier)**.** *Let $w = 2b+1 \ge 5$, $n \ge w$ even, $h = wn/2$, and $1 \le k < h$. Use the weakened erosion model of Theorem 11.25, with maximal survivor retention and the all-radius bands of Theorem 29.3, but without triangle or near-square history flags. In each feature class (c,e) its positive states capturable within t days form the interval $1 \le a \le B_t(c,e)$. The nine thresholds have a fixed-size update described below.*

For completeness the update can be written using elementary roots. Let $A_m(x)$ be zero for $x \le 0$ and otherwise $\left\lceil (\sqrt{m^2+4x} - m)/2 \right\rceil$; let $Q(x)$ be zero for $x \le 0$ and otherwise $\lceil (1+\sqrt{1+4x})/2 \rceil$. The least closed surplus with $z \le \sigma + P(v,u,j;\sigma)$ is

$$L(z; v, u, j) = \begin{cases} v + 2j + A_4(z - 2v - 3j - u), & u > 0, \\ v + 2j + A_3(z - 2v - 3j), & u = 0 < v, \\ 2j + \min(A_5(2z - 6j), A_2(z - 3j + 2)), & u = v = 0 < j, \\ \min(A_1(2z), Q(z)), & u = v = j = 0. \end{cases}$$

The least surplus with $d \le F_{AB}(\sigma)$ is

$$K(d; A, B) = \begin{cases} A + 2B - 2 + \max(2, Q(d - A - 2B + 2)), & B > 0, \\ A - 1 + \max(1, R((d - A + 1)_+)), & B = 0 < A, \\ \min(Q(2d), 1 + R(d+1)), & A = B = 0 < d, \\ 0, & A = B = d = 0. \end{cases}$$

Fix target (z, j, v) and survivor corner counts (i, u), and put $\delta = v - i$, $d = h - z$. Set

$$\theta = \min\{L(z; v, u, j), K(d; 2 - j, 2 - v), b + 1\},$$

also including b in this minimum when $(v, j) = (2, 0)$. If $j = 2$, $2 \leq r = Q(d) \leq b$, $D_r < d \leq r(r-1)$, and $(i, u, v) \neq (1, 2, 1)$, put $\lambda = r + 1 - \delta$; otherwise put $\lambda = 0$. Here $D_r = \max(r(r-1)/2, r(r-2))$. Define, for $z > 0$,

$$W_{iujv}(z) = \max\{0, z - \delta - \max(\theta, \lambda, z - \delta - (h - 4 + i + u), 0)\},$$

and $W_{iujv}(0) = 0$. Initialize all nine thresholds at zero and update by

$$B_{t+1}(c, e) = \min\{h - 4 + c + e,\ k + \max W_{iujv}(B_t(j, v))\}. \tag{162}$$

The maximum is over the fixed set of indices satisfying

$$i \leq c, \quad u \leq e, \quad u \leq j, \quad i \leq v, \quad i \geq \max(0, c + v - 2), \quad c - i + e - u \leq k.$$

All corner and erosion counts range from zero to two.

Proof. Substituting the shifted variable in the four quadratic branches of P and F gives L and K. For a fixed target, closed surplus is $\sigma = z - q - \delta$; thus admitted positive survivors form precisely the interval $1 \leq q \leq W_{iujv}(z)$. The band restriction is the additional lower bound $\sigma \geq \lambda$. Smaller original surpluses were already excluded by the component capacities.

The upper endpoint W is nondecreasing with target size. Away from band entrances this follows from the same inverse monotonicity as in Theorem 29.1. At an entrance the preceding hole size is $r(r-1) + 1$. Original surplus at most r was already impossible there: its high capacity is at most $r(r-1)$, while its small branch, with outgoing corner count two, would have target size at most $r^2 < h - r(r-1) - 1$. The critical exception has outgoing count zero. Hence the extra constraint causes no downward jump. Band exits only remove a constraint, and the bands are disjoint.

Induct on the horizon. Necessity of (162) follows by taking the largest admitted target in its feature class. For sufficiency choose a positive maximum W, let $\ell = c - i + e - u \leq k$, and choose $q = \min(W, a-\ell)$, $p = a-q$ for any $k < a \leq B_{t+1}(c, e)$. Then $q \geq a-k \geq 1$ and $\ell \leq p \leq k$. The state cap is compatible because $(h-4+c+e)-\ell = h-4+i+u$. The initial-interval property supplies the required edge. The positive lower-cardinality constraints were explicitly removed from this merger model, so none is being assumed in this reconstruction. Terminal states clear directly. □

Each update above uses an absolute bounded number of standard numerical operations. Iterating to the first full-state threshold is still a recurrence, not the required fixed-size expression for physical capture time. The nine-threshold theorem deliberately excludes the stronger history flag below: its interval property has not been asserted for that flagged model. The independent all-state graph comparisons are retained in `research/bridge-lower-frontier-nine-check.json`.

29.2 All pronic radii and an observed outgoing corner

The localization constraints admit two uniform improvements. The first uses every subcritical radius and the corner information already present in an erosion. The second prevents a strict near-square localization from being followed immediately by another closed near-square event. Both retain the merger needed for a shared-budget lower bound. Throughout this subsection let

$$w = 2b + 1 \geq 5, \qquad n \geq w \text{ be even}, \qquad h = wn/2,$$

so that $h > 2b^2$. Keep the corner counts c, e of Section 11.4.

Theorem 29.3 (Pronic localization at every radius)**.** *For $2 \le r \le b$, put*

$$A_r = (r-1)(r-2), \qquad C_r = r(r-1), \qquad D_r = \max\{r(r-1)/2, r(r-2)\}.$$

Let U satisfy $c(U) = 0$, $|U| \le C_r$, and $|N(U)| - |U| \le r$. Suppose either

$$|U| > D_r, \qquad \text{or} \qquad c(N(U)) = 1 \ \text{ and } \ |U| > A_r.$$

Then U is connected under sharing a neighbor, its surplus is exactly r, and $N(U)$ contains exactly one corner o. Every room of U has Manhattan distance at most $2r-3$ from o.

Proof. For each nonempty sharing-neighbor component U_l, the finite corner-count dichotomy and its odd-width critical extension apply. The critical exception requires a current corner and is unavailable. The large branch is impossible because

$$h - |U_l| - s_l \ge h - r^2 > r^2 \ge \Psi_{0,j}(s_l).$$

Thus $|U_l| \le s_l(s_l - 1)$ and $s_l \ge 2$. Distinct components have disjoint neighborhoods. If there are at least two, convex allocation of their surpluses gives

$$|U| \le (r-2)(r-3) + 2 \le A_r \le D_r \qquad (r \ge 4).$$

For $r = 2, 3$, two components already require too much surplus. Hence U is connected. Surplus at most $r-1$ would imply $|U| \le A_r$, also impossible. If the outgoing count is not supplied by the hypothesis, the inequalities $\Phi_{0,0}(r) = D_r$ and $\Phi_{0,2}(r) \le D_r$ force it to equal one.

When $r < b$, the small side of Theorem 11.1 gives genuine corner quadrants. When $r = b$, the same separator argument as in Theorem 11.14 applies: its alternative with no adjacent unmarked columns and no current corner has outgoing count zero, and is excluded. Connectedness therefore places U in the quadrant based at its unique outgoing corner o. Its occupied odd diagonals are $1, 3, \ldots, 2L-1$ without gaps. Their disjoint upper shadows each contribute at least one extra room, and o contributes one more. Thus $|N(U)| \ge |U| + L + 1$ and $L \le r-1$, proving the radius bound for the original support. ☐

Corollary 29.4 (Dual intervals and mixed-radius exclusion)**.** *Let $B = N(S)$, put $d = h - |B|$, and suppose $c(B) = 2$, $d \le C_r$, and $|B| - |S| \le r$ for some $2 \le r \le b$. If $d > D_r$, or if $e(B) = 1$ and $d > A_r$, then*

$$|B| - |S| = r, \qquad (c(S), e(S), e(B)) = (1, 2, 1).$$

Two consecutive inspection transitions cannot both satisfy such an interval condition, even with different radii.

Proof. For $U = V_{\mathrm{opp}} \setminus B$, $N(U) \subseteq V_{\mathrm{src}} \setminus S$, and $c(N(U)) = 2 - e(B)$. The theorem applies and forces surplus exactly r. Equality follows in the complementary inclusion, so $N(U) = V_{\mathrm{src}} \setminus S$ and $c(S) = 1$. The second neighborhood of U lies within radius $2r-1$ of its corner. Both opposite-colored corners are at distance at least $n - 1 \ge 2b+1$, so their neighbor pairs lie in S and $e(S) = 2$. Also $E(B) = S$, hence $e(B) = 1$. After one such transition the erosion count is one; deletion cannot increase it to the two required by another. ☐

Corollary 29.5 (All the dual intervals preserve the merger)**.** *The guarded merger of Theorem 11.25 remains valid when all rules of Corollary 29.4 are imposed simultaneously.*

Proof. Repeat the identity argument of Corollary 11.27 at radius r. A triggering merged target has $J = 2$, so both component outgoing counts are two. Their holes sum to $d \leq C_r$. A small closed component would have target at most $\sigma_l^2 \leq r^2$ and holes at least $h - r^2 > C_r$; the critical exception has outgoing count zero. Both components are therefore high, with

$$d_l \leq \Phi_{0,2-v_l}(\sigma_l) \leq s_l(s_l - 1).$$

If both original surpluses are positive, their sum at most r gives $d_1 + d_2 \leq (r-1)(r-2) = A_r$. This contradicts either interval trigger. The remaining component is original-full-to-full at cost zero. The merger is the identity on the constrained component, transferring its original surplus and all three features. This holds for every r. $\square$

29.3 A forbidden successor of near-square localization

Write $\mathcal{Q}_r(o)$ for the r^2-room corner triangle consisting of relative diagonals $0, 2, \ldots, 2r-2$ about o. For successive hole sets U, V of one cohort and useful inspection set I, one always has

$$N(V) \subseteq U \cup I, \qquad |U \setminus N(V)| \leq |U| + |I| - |N(V)|.$$

This simple inclusion retains the cost of changing corner locations.

Lemma 29.6 (Opposite corner triangles leave a large remainder)**.** *Suppose $U \subseteq \mathcal{Q}_r(o_U)$ and $V \subseteq \mathcal{Q}_s(o_V)$ have opposite colors, where $2 \leq r, s \leq b$ and $|U| > r(r-1)$. Then*

$$|U \setminus N(V)| \geq \frac{r(r-1)}{2} + 1 \geq 2.$$

Proof. The anchors have opposite colors. If adjacent, they lie at opposite ends of the even side of length n. Reflect coordinates to put them at $(0,0)$ and $(0, n-1)$, with the second coordinate longitudinal. Set $z = n/2 - s \geq 1$. Then $N(V) \subseteq \{(x,y) : y - x \geq 2z\}$. On the diagonal $x + y = 2j$ of $\mathcal{Q}_r(0,0)$, exactly $\min\{2j+1, j+z\}$ rooms lie outside this half-plane. Their total is

$$C(r,z) = \begin{cases} r(r-1)/2 + rz - z(z-1)/2, & 0 \leq z \leq r, \\ r^2, & z \geq r. \end{cases}$$

In particular $C(r,z) \geq C(r,1) = r(r+1)/2$. At most $r-1$ rooms are missing from $\mathcal{Q}_r(o_U)$, giving the asserted remainder.

If the anchors are diagonally opposite, their distance is $w + n - 2 \geq 4b + 1$. The radii of $\mathcal{Q}_r(o_U)$ and $N(\mathcal{Q}_s(o_V))$ sum to at most $4b - 3$, so they are disjoint and the bound is stronger. The odd short side cannot join adjacent opposite-colored corners. These exhaust the orientations. $\square$

For a transition with survivor (q, i, u), target (t, j, v) and original surplus $\ell = t - q$, call it a *strict near-square event* when

$$i = 2, \quad (j, v) = (1, 2), \quad 2 \leq \ell \leq b, \quad \ell(\ell - 1) < h - t \leq \ell^2. \tag{163}$$

It sets a one-step history flag on the target. The flag is defined from the history's own preceding event, and is zero initially. Every nontriggering new action clears it. For a physical history, i is the actual corner count of its maximally retained survivor, not an optimization over unspecified supports.

Theorem 29.7 (Forbidden closed near-square successor)**.** *After a strict near-square event, the next transition cannot have $(j, v) = (1, 2)$ and closed surplus $\sigma = \ell + i - v$ satisfying*

$$2 \leq \sigma \leq b, \qquad \sigma(\sigma - 1) < h - t \leq \sigma^2. \tag{164}$$

This is a necessary rule for arbitrary physical histories after maximal survivor retention.

Proof. Let U be the holes following the strict event, with its surplus parameter r. Corner closure has no slack since $i = v = 2$. Thus $|N(U)| - |U| \leq r$, while the size profile forces the reverse inequality. The large corner branch is impossible because $h - r^2 > r^2$. Theorem 11.17 gives $U \subseteq \mathcal{Q}_r(o_U)$ and $|U| > r(r - 1)$.

For the proposed successor let V be its holes. The closed-surplus bound gives $|N(V)| - |V| \leq \sigma$. The size profile again forces equality, and the same localization gives $V \subseteq \mathcal{Q}_\sigma(o_V)$. Lemma 29.6 implies $|U \setminus N(V)| \geq 2$. The flagged source has $(c, e) = (1, 2)$, so maximality and target $v = 2$ force survivor $i = 1$. Hence $\ell - \sigma = v - i = 1$, and the complementary inclusion gives

$$|U \setminus N(V)| \leq |U| + |I| - |N(V)| = \ell - \sigma = 1,$$

a contradiction. □

Theorem 29.8 (The near-square flag preserves the guarded merger)**.** *The weakened, closed, maximal erosion model of Theorem 11.25 remains closed under its guarded merger after adding the flag* (163) *and the restriction of Theorem 29.7. Earlier static interval rules and pronic flags may be retained.*

Proof. First consider formation of a merged strict event of surplus r. Merged survivor count $I = 2$ and target erosion $V = 2$ force $i_1 = i_2 = v_1 = v_2 = 2$, so both closure slacks vanish. Target $J = 1$ forces $j_1 + j_2 = 3$, hence both outgoing counts are positive. The component holes sum to $D \in (r(r - 1), r^2]$. A small closed component would have target at most $s_l^2 \leq r^2$ and holes at least $h - r^2 > r^2$, impossible; the critical exception has outgoing count zero. Both components are high, and

$$d_l \leq \Phi_{2-j_l,0}(s_l) \leq s_l^2.$$

If both original surpluses were positive, their sum r would give $D \leq (r - 1)^2 + 1 \leq r(r - 1)$, a contradiction. Thus one component is original-full-to-full at cost zero. The merged event is the identity on the other component. Defining the merged flag from its own trigger maintains the invariant that one component is flagged and the other is full.

Now start from a flagged merged source and suppose a proposed successor satisfies (164). Target $V = 2$ forces both component target erosions to be two. Their source corner counts are one and two, so maximality forces $i_{\text{flag}} = 1$ and $i_{\text{full}} = 2$. Their closure slacks are one and zero; consequently the merged closed surplus σ is the sum of their closed surpluses. The same small-branch exclusion forces both high, now with

$$d_l \leq \Phi_{2-j_l,0}(\sigma_l) \leq \sigma_l^2.$$

Two positive closed surpluses would give total holes at most $(\sigma - 1)^2 + 1 \leq \sigma(\sigma - 1)$, contrary to the proposed event. One closed surplus is therefore zero. It cannot be the flagged one: its high bound would force a full target and survivor $q = h - 1$, whereas the flagged source has at least $r(r - 1) + 1 \geq 3$ holes. Thus the full component has closed surplus zero. Its closure slack also vanishes, so it is original-full-to-full. The successor is again an identity merger, and is excluded by the flagged component's own rule. Nontriggering actions reset the flag.

All new restrictions transfer by these identity arguments; they do not interpret a synthetic merged support as a physical set. The older pronic flags have target erosion zero, while the new flag has target erosion two, so their formation rules do not conflict. The original strict precapture guard is unchanged. ☐

The resulting clock and its preterminal residual give lower bounds by Corollary 11.23; they still need comparison with a physical construction. In particular, a backward interval recurrence for the *unflagged* erosion model is not being asserted for this richer history model.

29.4 Persistent corner radii and a clipped clock

A radius certificate of relative color $p \in \{0, 1\}$ says that the current support lies in a Manhattan ball of radius R about some corner of that relative color. Relative color zero means the support's color. Put $D = w+n-2$ and use D for an uninformative certificate. One inspection and movement sends (p, R) to $(1-p, \min\{D, R+1\})$. This assertion concerns the original support and requires no compression of the inspection sets.

Theorem 29.9 (Square interval localization, including the critical radius)**.** *Suppose* $2 \le r \le b$*,* $c(U) = 1$*,*

$$r(r-1) < |U| \le r^2, \qquad |N(U)| - |U| = r.$$

Then $U \subseteq \mathcal{Q}_r(o)$ *at its unique own-colored corner. Its target has* $(c, e) = (0, 1)$ *and relative-color-one radius at most* $2r-1$*. Its next target has* $e = 0$ *and* $c \le 1 + \mathbf{1}_{\{r=b\}}$*.*

Proof. For $r < b$, exclude the high branch and apply the near-square localization theorem. At $r = b$, the separator case of Theorem 11.3 uses the same quadrant argument. In its nonseparator case, reflect so that the empty endpoint is U_{2b}. Each even fiber U_{2q} is an actual prefix of length at most $b-q$. The odd fiber R_{2q+1} is contained in U_{2q+2}, hence in a prefix of length at most $b-q-1$. Under the physical map $(x, j) \mapsto (x, 2j + (x \bmod 2))$, both bounds imply $x + y \le 2b-2$. Thus the entire support lies in $\mathcal{Q}_b$; equality in the size bound is unnecessary. The high branch is excluded by $h - r^2 - r > r^2$. The target radius is $2r-1$. It omits both own-colored corners and contains the neighbor pair of o, giving $(c, e) = (0, 1)$. After another move the radius is $2r \le 2b < n$, so erosion is zero. The second own-colored corner is at distance $w-1 = 2b$, explaining the stated critical allowance. ☐

Retain two radii ρ_0, ρ_1. A pronic event from Theorem 29.3 with known outgoing corner tightens ρ_0 to $2r-2$, since its target is within that radius. The square event above tightens ρ_1 to $2r-1$. Every movement first swaps the two radii and adds one, capped at D; each applicable tightening is then a coordinatewise minimum. These two certificates may use different corners. Their exact feature consequences are

	c at most	e at most
$\rho_0 = R$	$1 + \mathbf{1}_{\{R \ge w-1\}}$	$\mathbf{1}_{\{R \ge n\}} + \mathbf{1}_{\{R \ge w+n-3\}}$
$\rho_1 = R$	$\mathbf{1}_{\{R \ge n-1\}} + \mathbf{1}_{\{R \ge D\}}$	$\mathbf{1}_{\{R \ge 1\}} + \mathbf{1}_{\{R \ge w\}}$.

An adjacent corner at distance d has neighbors at distances $d-1$ and $d+1$; the diagonally opposite corner has both at distance $d-1$.

Theorem 29.10 (The paired radius rules preserve the guarded merger)**.** *The two interval resets and all their later feature consequences preserve the guarded erosion merger. The coordinate invariant is*

$$\rho_p^{\mathrm{merged}} \geq \min\{\rho_p^{(1)}, \rho_p^{(2)}\}, \qquad p = 0, 1.$$

Proof. A merged square reset of surplus $r \leq b$ is an identity merger. If both component original surpluses were positive, each would be at most $r-1$. A low component would force merged survivor size at most $(r-1)^2 < r(r-1)$; two high components would bound target holes by r^2, contrary to $h - r^2 - r > r^2$. Thus one component has zero original surplus. Its empty alternative is excluded by the positive guarded survivor, leaving a full-to-full zero-cost partner. For a pronic reset, merged outgoing count one forces both component outgoing counts to be positive. A low component with positive surplus at most $r-1$ has at most $(r-1)(r-2)$ rooms, contradicting the strict pronic interval; two high components again contradict the hole bound. Formation is therefore an identity merger in both cases.

Propagation preserves the displayed coordinate inequality. A merged reset is an identical reset of the non-full component; any additional component reset only decreases its side of the inequality. The componentwise corner and erosion merger bounds then transfer each radius restriction. A common witness for both coordinates is not needed. ☐

Lemma 29.11 (The corner-ball lens)**.** *The largest intersection of balls of radii R_0, R_1 about opposite-colored corners, restricted to the color of the first corner, has size*

$$M(R_0, R_1) = \#\{(x, y) : 0 \leq x < w,\ 0 \leq y < n,\ x+y \equiv 0 \pmod 2, x+y \leq R_0,\ x+n-1-y \leq R_1\}.$$

Consequently a support with the two radius certificates satisfies $a \leq M(\rho_0, \rho_1)$. This inequality preserves the guarded merger. For the square triangle $\mathcal{Q}_r$, the corresponding intersection with an opposite-corner ball of radius R is

$$L_r(R) = r^2 - F_r\left(\left\lfloor \frac{R-n+1}{2} \right\rfloor\right), \quad F_r(t) = \mathrm{Tri}(r-1-t) - 2\,\mathrm{Tri}(-1-t) + \mathrm{Tri}(-r-1-t),$$

where $\mathrm{Tri}(z) = z_+(z_+ + 1)/2$ *and* $z_+ = \max\{z, 0\}$.

Proof. Reflect the first anchor to $(0, 0)$. On each horizontal row its ball is a left prefix. The aligned opposite anchor $(0, n-1)$ also gives a left prefix; the diagonal anchor gives the reflected right prefix with the same color cardinality, since $w-1$ is even. Aligned prefixes maximize the intersection. For $\mathcal{Q}_r$, on diagonal $x+y = 2j$ the aligned old ball omits $\max\{0, j-t\} - \max\{0, -j-t-1\}$ rooms. Summing for $0 \leq j < r$ gives $F_r(t)$, including negative t and radii beyond the triangle.

If one component witnesses both merged radius bounds, its intersection bound transfers since the merged size is no larger. If the witnesses are different, let B_0, B_1 be the two ball cardinalities. The merged size is at most $B_0 + B_1 - h$, which is at most the intersection of any such two balls by inclusion-exclusion, hence at most M. Monotonicity in the radii completes the transfer. ☐

Set $k = b+1$ and $K = b^2 + b + 1$. For positive $q \leq K$ define

$$\alpha_0(q) = \min\{\lceil \sqrt{q} \rceil, b\}, \qquad \alpha_1(q) = \min\{\min\{r : q \leq r(r-1)\}, b+1\},$$

with both zero at $q = 0$. Define $g_0(q, R) = \alpha_0(q)$. To define g_1, start with $\alpha_0(q)$ and impose the following improvements: if $R < n-1$, use $\alpha_1(q)$; if $R < D$ and $q > b^2$, use $b+1$; and for $1 \leq r \leq b$ with $r(r-1) < q \leq r^2$ and $q > L_r(R)$, use $r+1$. These conditions agree when they overlap.

Theorem 29.12 (A radius-aware clipped lower clock)**.** *Let $H_p(a, R)$, for $0 \le a \le K$, be the capture time of the recurrence which replaces positive a by*

$$q = \max\{0, a - k\}, \qquad a' = \begin{cases} 0, & q = 0, \\ q + g_p(q, R), & q > 0, \end{cases} \quad p' = 1 - p, \quad R' = \min\{D, R + 1\}.$$

Put $\overline{H}_p(a, R) = H_p(\min\{a, K\}, R)$ and

$$\mathcal{P}(a, \rho_0, \rho_1) = \max\{\overline{H}_0(a, \rho_0), \overline{H}_1(a, \rho_1)\}.$$

Every physical step, and every guarded merged step with the paired radius rules, satisfies

$$\mathcal{P}(\mathit{source}) \le 1 + \mathcal{P}(\mathit{target}).$$

Thus $\mathcal{P}$ is a lower bound for the remaining capture time, including histories whose sizes leave the interval $[0, K]$.

Proof. In this size sector the high branch cannot improve the square or corner-free pronic profiles. A strict square-band survivor with the cheaper surplus $r < b$ must have one own corner and is localized by Theorem 29.9; the degree-two argument handles $r = 1$. At $r = b$, an informative opposite radius $R < D$ excludes the second own corner, removing the critical exception. The critical interval localization then applies. A proposed cheap survivor larger than $L_r(R)$ contradicts the lens lemma. For $q > b^2$ and $R < D$, all low critical capacities are too small and the same exception is absent, forcing surplus at least $b + 1$. These prove the g_p profiles. They are nondecreasing in q and nonincreasing in R; the $q > b^2$ clause prevents a drop at the cap. The clock recurrence is finite and monotone by comparison with the anchored prefix maps.

For a physical step put $q_0 = \max\{0, \min\{|A|, K\} - k\}$. A q_0-room subset of the survivor retains its radius certificate, and its neighborhood lies in the actual target. Hence the target size is at least $q_0 + g_p(q_0, R)$ when $q_0 > 0$; this lower bound is at most K. Monotonicity of H proves its one-step inequality after clipping. The case $q_0 = 0$ has source clock at most one. Taking the maximum over the swapped coordinates preserves the inequality, and every later radius tightening only strengthens it.

For a guarded merged step, an actual survivor count $q \ge K$ gives target count at least q by nonnegative surplus (the even-board Hall inequality), so clipping already suffices. If $q < K$, use the same low profile and $q \ge q_0$. Each added lens exclusion transfers by the identity merger in Theorem 29.10: the full partner's opposite radius is D, so any informative merged opposite radius belongs to the identical non-full component. The corner-free and critical-exception exclusions transfer by the corner bounds. This proves the merged one-step inequality as well. ☐

This potential does not assert a matching clock from the unrestricted boundary K: initially both radii can be uninformative. Comparing its entrance value with the physical prefix clock remains a separate problem. None of these radius or lens rules asserts interval closure for a flagged backward capturable set.

29.5 A boundary-history obstruction independent of length

Theorem 29.13 (Transferred joint boundary bound)**.** *On every even-area rectangle with both side lengths at least 24, with budget 13, a full initial color class leaves at least $h - 111$ rooms after 28 unsuccessful inspection-and-move steps, where h is half the area. Full initial uncertainty leaves at least $2h - 111$ rooms.*

Proof. Use the old unrefined 24-square model and its independently verified forward and equality calculations from Corollary 11.19. Translate a physical feature (a, c) with deficit $d = h - a$ to $(288 - d, c)$. While source deficit is at most 112, at most 13 deletions leave a survivor of size at least $h - 125 \geq 163 > 11^2$. For surplus $s < 12$ the finite corner theorem is subcritical, and its small branch is impossible. Its high condition $h - t \leq G_{ij}(s)$ is unchanged by deficit translation. For $s \geq 12$ the old 24-square model is unrestricted. Inspection costs and all corner counts are preserved.

A perfect matching gives $|N(S)| \geq |S|$, so a day's deficit increase is at most 13. Thus the translation remains valid through a first hypothetical crossing of deficit 112: all translated sizes are at least 163 and satisfy the upper corner bounds. The old forward model has no state below 176 through day 28. Therefore a solo history cannot cross 112 erased rooms by that time, and equality has corner count one.

For two cohorts, apply the same translation until a hypothetical first crossing of total deficit 112. Each individual source deficit is then at most 112. The two exact inspection allocations still sum to at most 13, and the guarded merger has size $288 - d_1 - d_2 \geq 163 > 13$ through the crossing. The old forward bound excludes it. If equality 112 holds at day 28, the independently checked joint equality certificate forces terminal pair $((176, 1), (288, 2))$ after translation. Thus one physical cohort has deficit 112 and the other is full. It remains to exclude the former physical history, rather than imposing a solo restriction on a synthetic merged set.

Suppose its final belief is A and put $U = V_p \setminus A$, so $|U| = 112$. Reversing the 28 inspections clears $N(U)$ within 28 days. The old 24-square backward frontiers through 28 have size at most 123. This bound transfers to larger boards by backward induction: a move into target size $t \leq 123$ has survivor size at most 123 and source size at most 136, all valid 24-square feature sizes. A physical small branch remains the same branch; a physical high branch implies $288 - t \leq h - t \leq G_{ij}(s)$; and surplus at least 12 is unrestricted in the old model. Consequently $|N(U)| \leq 123$ on the larger board too.

The unconditional profile gives the reverse inequality, hence surplus exactly 11. The high branch is impossible since $h - 123 \geq 165 > 11^2$. Because $112 > 11 \cdot 10$, near-square localization places U in a 121-room corner triangle on diagonals at most 20. Its neighborhood lies on diagonals at most 21. Both therefore lie in a 24-square corner subboard with exactly the same neighborhood there.

Restricting the princess to that subboard can only make capture easier. The same strategy compression and eight-pattern certificate used in Proposition 11.18 show that its 123-room neighborhood requires at least 29 days. This contradicts reversal. It excludes solo deficit 112 and, by the separately transferred joint equality certificate, total joint deficit 112 as well. □

This proof reuses the old finite certificates uniformly; it requires no new shape enumeration for each length. The next consequence evaluates the resulting time-dependent lower bound without a length recurrence.

29.6 A prescribed corner change and its length dependence

Throughout this subsection the board is $2r \times n$, with $n \geq 2r$. Write $P_p(q)$ for the bottom-left weightlex prefix of physical color p: rooms are ordered by increasing $x + y$, breaking ties by decreasing x. Its neighborhood is an opposite-color prefix; put $g_p(q) = |N(P_p(q))|$. Canonical pyramid cutoffs in neighboring columns differ by one. Their uniformly increased cutoffs contain the neighborhood and add at most r rooms, so

$$g_p(q) \leq q + r. \tag{165}$$

In particular, the deterministic prefix strategy inspecting its final $\min(k, q)$ rooms terminates when $k > r$.

Lemma 29.14 (A constant-surplus interval)**.** *If $r \geq 2$, then, for either physical color,*

$$g_p(q) = q + r \quad \textit{whenever} \quad (r-1)^2 + 1 \leq q \leq r(n-r) + r - 1.$$

Proof. Represent the prefix by its last diagonal $x + y = d$ and the least occupied column c on that diagonal. The stated size interval puts d between $2r - 3$ and n. Count the full earlier diagonals and the last partial diagonal separately. When $d = 2r-3$ or $2r-2$, the earlier neighborhood count exceeds the earlier source count by $r-1$, and the last neighborhood diagonal has one extra room. For $2r-1 \leq d \leq n-2$, the earlier-count difference is r and the partial lengths agree. At $d = n-1$, the upper size bound forces $c \geq 1$; at $d = n$ it forces $c \geq r+1 \geq 2$. These are precisely the conditions preventing the top boundary from shortening the last partial neighborhood diagonal. The earlier-count difference is again r. These cases also cover $n = 2r$. □

Theorem 29.15 (An explicit minimum-budget corner change)**.** *Let $r \geq 5$ and $k = r+1$, and define*

$$Q_A = r(n-r) + 2r - 2, \qquad Q_B = r(n-r) + 3, \qquad t = 2r - 5,$$

$$p_A \equiv n + 1 \pmod 2, \qquad p_B \equiv n \pmod 2.$$

Starting from either full color class, run the deterministic bottom-left prefix strategy until the first belief in phase p_A with at most Q_A rooms. Its size q_A is Q_A or $Q_A - 1$. A directly prescribed t-day transport then leaves a belief contained in the bottom-right phase-p_B prefix of size Q_B. The deterministic bottom-right prefix strategy completes capture.

Thus the entrance stopping rule, bridge length, and target size require no optimization over possible splice indices. The entrance and final prefix clocks remain recurrences.

Proof. At source sizes Q_A+2, Q_A+1, Q_A, subtraction of k leaves sizes in Lemma 29.14. Their next belief sizes are therefore Q_A+1, Q_A, Q_A-1, respectively. Prefix neighborhoods are nested as the prefix size increases, so no larger belief can jump below Q_A+1. Equation (165) ensures eventual arrival at this level. The two subsequent visits have opposite phases, proving the assertion about q_A.

We next bound the endpoint cutoffs a, b. First take $n = 2r$. The phase-1 prefix of size $r^2 + 2r - 3 = Q_A - 1$ fills all diagonals through $2r-1$ and at least the first two rooms of diagonal $2r+1$: there are $r - 3 \geq 2$ such additional rooms. Every column consequently has cutoff at least 2. The larger prefix of size Q_A satisfies the same bound.

Reflect the target horizontally. It becomes the phase-1 prefix of size r^2+3. Its last diagonal is $2r-1$, with $r+3$ rooms and least occupied column $r-3$. Its earlier diagonals have height at most $2r-3$, and its last partial diagonal has height at most $r+2 \leq 2r-3$. Its first two rooms on that diagonal fill the only columns absent from the earlier diagonals, so every target column is nonempty.

Adding $n - 2r$ to all cutoffs gives the stated prefixes on the general board: their bottom truncations are inactive, their sizes increase by $r(n-2r)$, and their phases change by $n-2r$. Hence

$$\min_x a_x \geq n - 2r + 2, \qquad \max_x b_x \leq n - 3, \qquad \max_x (b_x - a_x) \leq t.$$

The positive-part penalty in the transport of Section 13.1 vanishes. The canonical cutoff sums give

$$F_t(a, b) = q_A - Q_B + rt \leq Q_A - Q_B + rt = (r+1)t = kt.$$

The prescribed descending local-maximum word therefore supplies the bridge, without a boundary credit. Its endpoint containment suffices for the target prefix's capture strategy. □

Corollary 29.16 (Two-row extension of the prescribed strategy)**.** *Fix $r \geq 5$ and the initial corner phase p_0. Let $L_{r,p_0}(n)$ be the prescribed solo length in Theorem 29.15. Then*

$$L_{r,p_0}(n+2) = L_{r,p_0}(n) + 2r.$$

The final solo inspection set has unchanged size. Consequently the full reversal construction, including the optional reflected central merger, has length increased by $4r$. Minimizing over the two initial corner phases gives an upper bound

$$U_r(n) = 2rn - C_{r,n \bmod 2},$$

where the two width constants are determined by the prescribed strategies at lengths $2r$ and $2r+1$.

Proof. Before the entrance stops, each survivor prefix has every column nonempty. Indeed its size is at least $Q_A - k = r(n-r) + r - 3$; the prefix counts show that this contains the bottom room of every column in either phase when $r \geq 5$. Adding two to every cutoff therefore commutes with its physical inspection-and-neighborhood steps, adding $2r$ rooms to every belief. The stopping condition and phase are unchanged, so the entrance time is unchanged. Both bridge endpoints also increase by two, preserving their differences and the prescribed bridge length.

The enlarged target has size $Q_B + 2r$. During its first $2r$ tail steps, the survivor sizes range from $Q_B + 2r - k$ down to $Q_B + 1 - k$. They lie in the constant-surplus interval of Lemma 29.14 on the $(n+2)$-long board. Thus each step removes exactly one possible room, reaching the original size-Q_B target in its original phase.

The remaining tail is unaffected by the taller board. The original target has maximum height at most $n-3$, and later sizes decrease by (165). In the target phase, each later prefix is contained in that target. In the other phase, it is contained in the target's neighborhood: the latter is a prefix of size at least Q_B, by the rectangle's perfect matching. Every later belief therefore has height at most $n-2$, so none of its neighborhoods can leave the old board at the top. Its inspections and neighborhoods are identical on the two boards. This proves the solo length identity and preservation of the final inspection set. The usual reversal construction doubles the length, or saves one central day when the two reflected final inspection sets fit the quota. □

For width 24, the directly replayed base strategies have

n	$L_{12}(n)$	final solo inspections	$U_{12}(n)$
24	104	2	207
25	116	13	232

and hence

$$T_{13}(24,n) \leq \begin{cases} 24n - 369, & n \text{ even}, \\ 24n - 368, & n \text{ odd}, \end{cases} \qquad n \geq 24. \tag{166}$$

The construction also gives the replayed bounds 112 on 16×16, 144 on 16×18, and 308 on 32×32. It is not asserted to be optimal throughout its parameter range. Its dependence on length has been evaluated, but the width constants still contain prefix-clock iteration; this is not an $O(1)$ evaluation in all parameters.

29.7 Exact scheduling of the prescribed transport word

The boundary credit of Section 13.1 counts flips invisible on every day. One might hope to improve it by scheduling other flips on days when they are invisible. For the prescribed descending word, the following result identifies exactly when that can help.

Proposition 29.17 (Physical windows and fixed-word exactness)**.** *Let a, b be canonical pyramid cutoffs and $t \geq 1$. Put $c_x = \min(a_x, b_x - t)$, and list the legal descending word's letters (x, H), with $H = a_x, a_x - 2, \ldots, c_x + 2$, in decreasing height order. For $0 \leq i, j < t$, define*

$$A_i = \#\{H \geq n - i\}, \qquad B_j = \#\{H \geq -j\}.$$

This fixed word can be scheduled over t days with at most k visible flips per day if and only if some integers $0 \leq i \leq j < t$ satisfy

$$B_j - A_i \leq k(j - i + 1). \tag{167}$$

The assertion includes empty words and $k = 0$. On width $2r$, if $k \geq r$, it is equivalent to the existing criterion

$$F_t(a, b) \leq kt + z(B).$$

Thus time-dependent visibility cannot strengthen that criterion for this prescribed word in the feasible even-width budget range.

Proof. On day d, a letter is visible exactly when $0 \leq H + d < n$. If (167) holds, leave the word untouched until day i. Its first A_i letters are then above the board and free. Partition its next $B_j - A_i$ letters among days $i, \ldots, j$, charging at most k per day. Its final letters are below the board and free on day j. This preserves the word order. When $k = 0$, the middle band is empty, so no division into nonempty paid groups is required.

For necessity, greedily execute each day's longest remaining prefix costing at most k visible letters, including all free letters before the next unaffordable letter. This rule dominates every other fixed-word schedule at every day boundary. Indeed, from a farther word position, the remaining part of a competing next group is its suffix and costs no more on that same day. This argument allows visibility to depend on the day.

Let q_j be the greedy position before day j. Because heights are nonincreasing, greedy completes that day exactly when $\max(q_j, A_j) + k \geq B_j$; otherwise $q_{j+1} = \max(q_j, A_j) + k$. Before the first completion, induction gives

$$\max(q_j, A_j) + k = k(j + 1) + \max_{0 \leq i \leq j}(A_i - ki).$$

The completion condition is precisely (167). It includes completion without any paid letter and immediate completion of an empty word.

For width $2r$, each integer height appears in at most r letters, since its parity fixes a column color. Canonical initial heights lie below n, so $A_i \leq ri \leq ki$ and $A_0 = 0$. Moreover $B_{j+1} - B_j \leq r \leq k$. Thus the best window is $i = 0, j = t - 1$, and the criterion becomes $B_{t-1} \leq kt$.

Writing $F = F_t(a, b)$, the difference $F - B_{t-1}$ counts letters with $H \leq -t$. There is exactly one in each empty target bottom-root fiber: there $b_x = -2$, $c_x = -t - 2$, and the last old height is $-t$. The canonical initial height is at least $-t$; for $t = 1$ this follows from its opposite bottom parity, and for $t \geq 2$ from $a_x \geq -2$. An empty nonroot target has $b_x = -1$ and last old height $1 - t$; nonempty target fibers end still higher. Hence $F - B_{t-1} = z(B)$, which proves the final equivalence. It concerns this word only and does not make the transport criterion necessary for all physical strategies. □

29.8 Buffered comparisons at linear budgets

Let $w \geq 4$, $n \geq w$, $wn = 2h$, and $w \leq k < h$. Write $m = \lfloor k/2 \rfloor$. Let $B_t(c)$ be the critical three-corner backward frontier for capture within t days, and let $C_t(c)$ be the same frontier for reaching size at most m within t inspection-and-movement steps. Their initial vectors are $(0, -1, -1)$ and (m, m, m), respectively. Write $P_t(p)$ and $Q_t(p)$ for the corresponding physical, fixed-corner prefix frontiers, with initial vectors $(0, 0)$ and (m, m). The two comparisons of interest are

$$B_t(2) \leq \max_p Q_t(p), \qquad C_t(2) \leq \max_p P_{t+1}(p). \tag{168}$$

The maxima can choose different initial phases: each inequality is used separately to construct its own reflected physical strategy.

These comparisons imply a one-day bracket. Let τ be the first capture time of the critical three-corner model. If its half-budget test at $\tau - 1$ fails, the merger lower bound is 2τ, while the first comparison gives a prefix palindrome of length at most $2\tau + 1$, sharing its two half-budget final inspections. If the test succeeds, the lower bound is $2\tau - 1$, and the second comparison gives a solo prefix of length at most τ, hence a full search of length at most 2τ. These are explicit prefix inspections and reflections; a path in the necessary model is never used as a physical inspection sequence.

Proposition 29.18 (Sufficient buffered comparison)**.** *The comparisons* (168) *hold in either of the following parameter regions:*

$$\begin{aligned} &w = 2r: && m - r \geq \left\lceil \frac{(r-1)^2}{2(k-r)} \right\rceil, \\ &w = 2b+1,\ n \text{ even}: && m - b - 1 \geq \left\lceil \frac{b^2}{2(k-b-1)} \right\rceil. \end{aligned}$$

In the second line the hypothesis $k \geq w$ makes the denominator positive. These are sufficient regions, not classifications of every partial physical support.

Proof. Put $R(x) = \lceil \sqrt{x} \rceil$ and $\rho(x) = \min\{u \geq 0 : x \leq u(u+1)\}$. For even width $2r$, direct inversion of the diagonal-prefix profiles gives, for $x < h$,

$$E_p(x) = x - \min\{L_p(x), r, H_p(h - x)\},$$

where $L_0 = R$, $L_1 = \rho$, and $H_p = R$ when $p \equiv n \pmod 2$, while $H_p = 1 + \rho$ otherwise. At $x = h$ set $E_p(h) = h$. The low roots come from inverting the two corner-triangle profiles; the plateau subtracts r; complement and half-turn reflection give the high roots and the indicated parity. The low and high portions are disjoint since $h \geq 2r^2$.

The unconditional profile has inverse

$$E(x) = \max(0, x - \min\{r, \rho(x), R(h - x)\}).$$

Consequently its scalar backward orbit $S' = \min(h, k + E(S))$ dominates the critical corner frontier from the same size allowance. Compare it with physical prefixes initially at $\min(h, a + z)$ in both phases, where the scalar starts at a and $r \leq z \leq k - r$. While $u = \rho(S) < r$, maintain

$$\min_p X_p \geq \min(h, S + z - u).$$

If the next scalar rank remains low, write $u' = \rho(S')$. The input $x = S + z - u = S' - (k - z)$ satisfies $x \le u'^2$, since $S' \le u'(u' + 1)$ and $k - z \ge r \ge u'$. Thus either physical cost is at most u', proving the invariant. At low exit the universal cost bound r leaves the common credit $z - r$, including direct jumps to capture. The common middle increment $k - r$ preserves this credit. At high entry $d = h - S \le (r - 1)^2$. Above S, one physical high cost is at most $R(d)$ and the other at most $R(d) + 1$. Each trajectory alternates the two, losing at most $\lceil j/2 \rceil$ credits in j steps. At most $\lceil (r-1)^2/(k-r) \rceil$ steps remain. This proves the even comparison for $z = m$ and $z = k - m$, the two required seeds. Clipping a physical rank at h preserves capture thereafter.

For odd width $2b + 1$, the exact physical backward maps are

$$G_0(x) = \min(h, k + x - \min\{\rho(x), b, R(h - x)\}),$$
$$G_1(x) = \min(h, k + x - \min\{R(x), b + 1, 1 + \rho(h - x)\}).$$

with $G_0(h) = G_1(h) = h$ and crossed update $X_0' = G_0(X_1)$, $X_1' = G_1(X_0)$. Write $a = \max\{0, B_2\}$, $\delta = k - b - 1$, $A = b(b + 1) + 1$, and $U = h - b^2 - 1$. The same assertions below apply to the half-budget frontier. After the first update its reachable coordinates are ordered, with adjacent gaps at most one. The initial capture vector has $a = 0$ and first largest coordinate k. Its largest update is bounded above by $\min(h, a + k - \min\{\rho(a), b, R(h - a)\})$, and is at least $\min(h, a + \delta)$.

Start both physical phases at $a_0 + z$, clipped at h, with $k - z \ge b + 1$. Below A the monotone invariant is

$$\min_p X_p \ge \min(h, a + z - \min\{b + 1, \rho(a)\}).$$

The even-width root argument proves this as well: for $s = a + k - \rho(a)$ and $v = \rho(s) \le b$, the lower input $x = s - (k - z)$ satisfies $x \le v^2$. On leaving this portion, use the cap $b + 1$. Thus entry supplies common credit $c = z - b - 1$.

When $A \le B_0 \le B_2 \le U$, all finite root branches at surplus at most b are excluded, except the critical pair $(i, j) = (2, 0)$. The exact middle update is

$$B_0' = \min(h - 2, a + \delta), \quad B_1' = \min(h - 1, a + \delta), \quad B_2' = \min(h, \max\{a, B_0 + 1\} + \delta).$$

It preserves the aligned bounds $X_0 \ge \min(h, B_2 + c)$ and $X_1 \ge \min(h, B_0 + c)$. At first entry one may have $B_0 < A$, but adjacent spread gives $B_0 \ge A - 2 > b^2$. The $i = 0$ low target capacities at surplus at most b are at most b^2, and all low capacities at surplus at most $b - 1$ are at most $b(b - 1) < A - 2$. Hence already $B_0' = \min(h - 2, a + \delta)$ and $B_2' \le \min(h, a + \delta + 1)$; these establish alignment without spending credit. Later middle steps spend none either.

At high exit $h - a \le b^2$. Follow the physical trajectory currently in phase zero, which has all c credits. Its next map is G_1; thereafter costly and noncostly maps alternate. On inputs at least a, their costs are at most $\min\{b, R(h-a)\}+1$ and $\min\{b, R(h-a)\}$, respectively. At most $\lceil b^2/\delta \rceil$ model steps remain, so the hypothesis supplies every possible credit expenditure. Take $(a_0, z) = (0, m)$ and $(m, k - m)$ to obtain the two comparisons. □

Theorem 29.19 (A common linear-budget one-day bracket)**.** *For every even-area rectangle with shorter side $w \ge 4$, longer side $n \ge w$, and $\lceil 4w/3 \rceil \le k < wn/2$, the critical three-corner lower bound and the physical prefix-palindrome upper bound differ by at most one day.*

Proof. For $w = 2r$, the first condition of Proposition 29.18 holds at $k = \lceil 8r/3 \rceil$: substituting $r = 3q, 3q + 1, 3q + 2$ directly proves $(r - 1)^2 \le 2(k - r)(\lfloor k/2 \rfloor - r)$. For $w = 2b + 1$, its second condition holds at $k = \lceil 8b/3 \rceil + 2$ by the same three residue cases. At the desired $k_0 = \lceil 4(2b+1)/3 \rceil$, the only missing pairs are

$$(b, k_0) = (2,7), (4,12), (5,15), (7,20), (8,23), (11,31), (14,39).$$

Indeed for $b = 3q$ the condition holds; for $b = 3q+1$ it reduces to $q^2 - 2q - 1 \geq 0$, and for $b = 3q+2$ to $q^2 - 4q - 4 \geq 0$. At $k_0 + 1$ it holds on all seven lines, and increasing k only improves the sufficient condition.

Here is a finite certificate with a uniform length proof for those seven fixed pairs. Put $D = 2k - 2b - 1$. In each of the two initial comparison modes, follow the five-coordinate state $(B_0, B_1, B_2, X_0, X_1)$ to an entry with

$$B_0 \geq b(b+1) + 1, \quad B_0 = B_1, \quad B_2 - B_0 \in \{0, 1\}, \quad \min X_p \geq b^2 + 1.$$

Low root branches are then inactive. In the common middle, two updates add D to all five coordinates. Record the largest coordinate M through entry and its first subsequent update. The margin $h \geq M + \max\{b^2 + 1, k + 2\}$ excludes high branches and state caps in these updates. The certificate gives the following stable starting lengths and even-length periods:

w	5	9	11	15	17	23	29
k	7	12	15	20	23	31	39
n_0	10	14	16	20	20	30	34
$2D$	18	30	38	50	58	78	98

For any $n \geq n_0$, increasing n by $2D$ increases h by wD. Insert w middle pairs at the fixed entry, adding $2w$ days and wD to all coordinates. Throughout insertion the deficits in the larger board are at least those covered by the stable margin. After insertion the entire old tail, including capture, translates by wD: low branches remain inactive, high-root arguments $h - x$ are unchanged, and every state cap translates. Even insertion length preserves phase. The comparison $B_2 \leq \max X_p$ is invariant under translation, and each inserted pair has the comparison margins of the first pair.

It therefore suffices to check the even lengths from $w+1$ through $n_0 + 2D - 2$, covering the shorter cases and one representative of every stable residue. The recorded certificate checks all 199 such lengths and all 12,393 deadlines in both modes; every comparison holds. It separately checks 107 translated tail states. The sources and exact entry coordinates are in `research/bridge-upper-transport-odd-exceptions.json`. This proves all lengths, and the terminal argument preceding the proposition proves the asserted one-day bracket. □

29.9 A sharper high-region estimate

Lemma 29.20 (Decreasing high-root deficit). *Let $v \geq 1$, $K > v$, $A = K - v$, and $d_0 \leq v^2$. Suppose that until capture the nonnegative deficits satisfy*

$$d_{j+1} \leq \max\{0, d_j - K + \lceil \sqrt{d_j} \rceil\}.$$

For $t \geq 1$, capture occurs within t steps if

$$tA + \sum_{j=0}^{t-1} \left\lfloor \frac{jA}{2v} \right\rfloor \geq v^2.$$

In particular the fixed-operation sufficient condition

$$4vtA + At(t-1) - 4v(t-1) \geq 4v^3 \tag{169}$$

suffices.

Proof. As long as capture has not occurred, $d_j \leq v^2 - jA$. Set $s = \lfloor jA/(2v) \rfloor$. Noncapture gives $s < v$, and $(v-s)^2 \geq v^2 - 2vs \geq v^2 - jA \geq d_j$. Thus the next decrease is at least $A + s$. Summing proves the first condition; if $jA \geq v^2$, the coarser bound has already proved capture. For the second condition use $\lfloor jA/(2v) \rfloor \geq jA/(2v) - 1$ for $j = 1, \ldots, t-1$ and sum. The hypothesis $t \geq 1$ is retained explicitly. □

Corollary 29.21 (A stronger linear-budget region)**.** *For even width $w = 2r \geq 4$, every $\lceil 13w/10 \rceil \leq k < rn$ satisfies the one-day bracket of Theorem 29.19. For odd width $w = 2b+1$, even length $n \geq w$, the same conclusion follows from $\lceil 13b/5 \rceil + 3 \leq k < wn/2$.*

More sharply, retain $k \geq w$, and require (169) *with $t = 2c$, $c \geq 1$, where*

	c	v	K
$w = 2r$	$\lfloor k/2 \rfloor - r$	$r-1$	k
$w = 2b+1$	$\lfloor k/2 \rfloor - b - 1$	b	$k-1.$

Either of these integer tests also proves both buffered comparisons.

Proof. The preceding low and middle arguments already provide c credits. In even width the high scalar deficit obeys the lemma with $K = k$, $v = r-1$. In odd width let $a = B_2 < h$, $d = h - a \leq b^2$. Choose survivor corner count one and target corner count two in the backward update. Its large inverse is $1 + \rho(d) \leq 1 + R(d)$; the survivor is admissible because $a \geq h - b^2 \geq b^2 + 3b + 1$ and $a - (1+\rho(d)) > 1$. Its upper cap is inactive since $a < h$. Consequently $B_2' \geq \min(h, a+k-1-R(d))$, giving the lemma with $K = k-1$. Capture within $2c$ steps requires at most c alternating credit expenditures, so both comparisons follow.

For the even clean bound, $t \geq 3r/5 - 1$ and $A = k - r + 1 \geq 8r/5 + 1$. The slack in (169) is increasing in A, t on this domain. Substituting these lower bounds and multiplying by 125 gives

$$52r^3 - 95r^2 - 25r + 250 = 52(r-2)^3 + 217(r-2)^2 + 219(r-2) + 236 > 0.$$

Also $c \geq 1$ for every $r \geq 2$. There are no finite exceptions. For the odd clean bound, $t \geq 3b/5$ and $A \geq 8b/5 + 2$; the same substitution gives positive slack $(52b^3 + 270b^2 + 350b)/125$. □

The new integer test is asymptotically sufficient at quota divided by half-width at least $(1 + \sqrt{17})/2$. It strengthens the previous sufficient gate; neither gate proves the entire proposed range $k \geq w$. The tests themselves use a fixed number of arithmetic operations, but the compared lower and physical upper clocks still require their stated evaluations. A uniform one-day bracket is not an exact fixed-operation formula for the minimum capture time.

30 Fixed numerical expressions for further odd rectangles

The number of stages in Proposition 6.4 depends on the width. The following region admits an absolute bound instead. This numerical simplification and the necessity of the scalar midpoint test are separate assertions.

Theorem 30.1 (A fixed number of solo evaluations). *Fix an integer $K \geq 1$, independently of all input parameters. Let $n \geq w = 2b+1$ be odd, $b \geq 2$, $m \geq b+1$, and put $D = 2m - w$. If $KD \geq b^2$, the exact solo capture time from any canonical prefix, and its remaining count at any prescribed day, can each be evaluated by at most $3K+1$ conditional paired-map blocks and one bulk floor division. Each block uses an absolute bounded number of elementary arithmetic operations, integer roots, and comparisons.*

In particular, fixing $K = 16$ gives 49 blocks whenever

$$16(2m - w) \geq b^2.$$

Four such evaluations give both solo times and both precentral residuals. They therefore evaluate the exact physical capture time in this region whenever $M \geq M_*$ in Theorem 7.5, and also wherever scalar necessity follows from Theorems 21.1 or 8.1, and in the larger domains of Theorems 30.4, 30.5 and Corollary 30.6 below.

Proof. Put $C_0 = M + 1$, $C_1 = M$, $H = b^2 + 1$, and $c_p = b + p$. Retain

$$f_p(x) = g_p((x - m)_+), \qquad P_p = f_{1-p} \circ f_p.$$

The profiles of Theorem 5.1 give $g_p(z) = z + c_p$ throughout $H \leq z \leq M - H$. Consequently, with

$$\ell_p = 2m + H - c_p, \qquad u = m + M - H,$$

both successive survivor inputs lie in that interval whenever $\ell_p \leq x \leq u$, and

$$P_p(x) = x - D. \tag{170}$$

The interval may be empty. The bounds for both inputs use $m \geq b + 1 \geq c_p$.

For every source, the stronger global progress bound

$$P_p(x) \leq \max(0, x - D) \tag{171}$$

holds. If either constituent step captures the cohort, this is immediate. Otherwise the two profile surpluses sum to at most $b + (b+1) = w$, giving $P_p(x) \leq x - 2m + w$. A positive output is therefore impossible when $x < D$.

The initial excess above the affine interval satisfies

$$C_p - u = H + 1 - p - m \leq b^2 \leq KD.$$

Write K conditional paired blocks, each acting only while the source exceeds u and recording capture after either constituent day. By (171), these blocks capture the cohort or leave $x \leq u$. If now $\ell_p \leq x \leq u$, perform

$$q = 1 + \lfloor (x - \ell_p)/D \rfloor$$

pairs at once, replacing x by $x - qD$ and adding $2q$ days. Every source of these pairs is in the affine interval, so their outputs are exact and positive. Otherwise omit this bulk operation.

Unless already captured, the remaining source obeys

$$x \leq \ell_p - 1 = D + b^2 + b + 1 - p \leq D + 2b^2 \leq (2K+1)D.$$

Thus another $2K + 1$ conditional paired blocks suffice, with the last block again distinguishing capture after its first or second day. There are $3K + 1$ blocks in total. For a prescribed day, truncate the bulk count by the available number of complete pairs, and stop a block after one constituent day if required. Subsequent blocks then do nothing. The same bound holds for every prefix size and every prescribed day.

Compute the two full-cohort durations, take their minimum τ, and evaluate both full-cohort residuals at $\tau - 1$, with their physical colors determined by time parity. Their sum is the left-hand side of (24). The cited scalar-necessity theorems now select $2\tau - 1$ or 2τ. $\square$

For $K = 16$, the entire scalar calculation uses at most 196 paired blocks and four bulk divisions. This is an absolute operation bound; integer bit lengths are not bounded. Selecting K afresh as $\lceil b^2/D \rceil$ would not give this conclusion. Outside a proved scalar-necessity domain, the same calculation gives the one-day interval and its sufficient central test, rather than an exact physical time. The implementation and comparisons with the independent band evaluator are recorded in `research/bridge-odd-clock-quadratic-check.json`.

30.1 The remaining seven-row odd strips

Theorem 30.2 (Seven-row formulas; finite certificate and ordinary period proof)**.** *Let $n \geq 7$ be odd. For $m \in \{6, 8, 9\}$ put $D = 2m - 7$ and write $n = 7 + 2Dq + 2j$, where $0 \leq j < D$. Then*

$$T_m(P_7 \,\square\, P_n) = 28q + c_m(j),$$

where the lists below are indexed from zero:

m	$(c_m(0), \ldots, c_m(D-1))$
6	$(16, 21, 26, 32, 38)$
8	$(10, 12, 16, 18, 22, 25, 28, 31, 34)$
9	$(8, 10, 13, 16, 18, 20, 23, 26, 28, 30, 33)$.

The remaining intermediate budget has $T_7(P_7 \,\square\, P_n) = 2n - 2$.

Proof. First establish scalar necessity at every length. Substitution of $b = 3$ into (26) and Theorem 7.5 gives the following complete finite complements of its onset:

m	M_*	odd lengths below onset	number
6	327	$7, 9, \ldots, 93$	44
7	393	$7, 9, \ldots, 111$	53
8	657	$7, 9, \ldots, 187$	91
9	777	$7, 9, \ldots, 221$	108.

For these 296 cases the finite certificate `research/bridge-odd-clock-width-seven-check.json` verifies scalar necessity by the following exhaustive calculation. Invert the exact forward profiles to obtain the two integer inverse tables. Starting from $(0, 0)$, apply every quota split in (27), retaining the complete coordinatewise maximal frontier at each layer. Monotonicity makes this pruning exact; no low-total mixed state is otherwise omitted. At $L = \tau - 1$, minimize

$$(E - x - u)_+ + (M - y - v)_+$$

over every pair $(x, y), (u, v)$ of frontier states. In every case its comparison with m agrees with the scalar test. The certificate retains all final frontiers and the exact central minima; its independent implementation in `src/bridge_odd_clock_width_seven.py` also agrees with the retained-frontier implementation on every central minimum. There are 27 746 689 quota transitions and at most 581 retained states in any layer. The next odd length in each row has $M \geq M_*$, so the eventual theorem covers every remaining length. This establishes scalar necessity without extrapolating a finite pattern.

We next prove the numerical period for all lengths. Put $D = 2m - 7$, $E = M + 1 = (7n+1)/2$. The two profile surpluses equal 3 and 4 on

$$5 \leq z \leq M - 9, \qquad 7 \leq z \leq M - 7,$$

respectively. Requiring both successive survivor inputs to be in these intervals gives $P_p(x) = x - D$ on

$$\ell_0 = 2m + 4 \leq x \leq u_0 = m + M - 9, \qquad \ell_1 = 2m + 1 \leq x \leq u_1 = m + M - 7.$$

Starting from the full color classes, the first paired step leaves

$$x_0 = M - D, \qquad x_1 = M - D - \epsilon, \qquad \epsilon = \begin{cases} 1, & m = 6, \\ 0, & m = 7, 8, 9. \end{cases}$$

Indeed the first phase-zero surplus is 2, followed by 4; the first phase-one surplus is 3 for $m = 6$ and 4 otherwise, followed by 3. The increasing roots have reached their caps already at $n = 7$, while the reflected arguments depend only on m. No capture occurs in this first pair: its smallest output is $13 > m$.

Since $M \geq 24$, one has $\ell_p - D \leq x_p \leq u_p$ for every allowed budget and length. Thus

$$q_p = 1 + \lfloor (x_p - \ell_p)/D \rfloor \geq 0, \qquad r_p = x_p - q_p D$$

give the exact number of further affine pairs and the terminal source. In particular

$$11 \leq r_0 \leq 2m + 3 \leq 21, \qquad 8 \leq r_1 \leq 2m \leq 18.$$

Every subsequent profile argument is at most 12. Both reflected root terms are inactive for these arguments when $M \geq 24$, and counts strictly decrease because the surpluses are at most $4 < m$. The complete terminal evolution from (p, r_p) is therefore independent of n.

Replacing n by $n + 2D$ increases M and each x_p by $7D$, increases each q_p by 7, and preserves each r_p. Both solo durations consequently increase by exactly 14. We must also check that the common precentral horizon lies in this invariant terminal evolution. The two numerators defining q_p differ by $3 - \epsilon$, which lies strictly between 0 and D. Hence $q_1 - q_0 \in \{0, 1\}$. Put $S_p = 2 + 2q_p$, the terminal entry time. Then $S_1 \in \{S_0, S_0 + 2\}$. The phase-zero terminal evolution takes at least two days since $r_0 > m$, and the phase-one evolution takes at least one. Therefore

$$L = \min(\tau_0, \tau_1) - 1 \geq S_0, \qquad L \geq S_1 - 1.$$

Phase zero is already in its invariant tail at L. Phase one is also in its tail, except possibly at the intermediate day of its last affine pair. In that exception $q_1 = q_0 + 1 \geq 1$; the pair starts from $r_1 + D$, so its intermediate count also repeats. The added 14 days preserve physical-color parity. Thus the scalar central decision repeats, including this possible endpoint, and

$$T_m(P_7 \square P_{n+2D}) = T_m(P_7 \square P_n) + 28.$$

The displayed constants are the base values at $n = 7, 9, \ldots, 7 + 2(D-1)$, all included in the finite certificate. For $m = 7$ the seven values are $12, 16, 20, 24, 28, 32, 36$; the same period simplifies to $2n - 2$. $\square$

These are fixed lists of constants and ordinary floor-and-remainder expressions. The attaining inspections are the compatible prefix preparations and central inspection, or the doubled solo schedule, from Theorem 6.1. Their precentral residuals may be evaluated by Theorem 30.1 with $K = 16$ (already $K = 2$ suffices here). Listing every inspection is, of course, allowed to take time proportional to the schedule.

30.2 A shorter ancestry argument and larger scalar domains

Lemma 30.3 (A short-window coordinate bound)**.** *Put $\Gamma = m-1$. If integers $q \geq 1$ and $1 \leq r \leq b$ satisfy $q(2r-b-1) \geq \Gamma$, every attainable exceptional mixed deficit pair obeys*

$$\min(x, y) \leq K_{q,r} := (q-1)m + r(r-1).$$

Proof. Before solo capture, every inverse output is at most its input, and the joint input sum is at most M. Indeed the larger solo capacity plus m is at most M at every proper transition, including the majority inverse's special saturation threshold. If fewer than q days have elapsed, the total deficit is at most $(q-1)m$, which proves the claim. Otherwise suppose the final minimum exceeds $K_{q,r}$. The minimum increases by at most m in a day, so both outputs in each of the last q transitions are at least $r(r-1)+1$.

Writing the inverse inputs as X, Y, both X, Y and their complements $M-X, M-Y$ are at least $r(r-1)+1$. The exact inverse formulas therefore lose at least r on each coordinate. The joint total S increases by at most $m-2r$ per day. The larger solo capacity A increases by at least $m-b-1$, since every proper inverse is at least its input minus $b+1$. Hence $A-S$ increases by at least $2r-b-1$ on each of these q days. Concentration gives $A-S \geq 0$ initially, whereas exceptionality requires $A-S < A-B \leq \Gamma$ finally. This contradicts the assumed inequality for q, r. $\square$

Theorem 30.4 (A shorter scalar onset)**.** *Let $K \geq m$ be any valid minimum-coordinate bound for every exceptional state. Put $H = b^2+1$, $\Gamma = m-1$, and*

$$\begin{aligned}
J_0 &= \max(2K+m+1, K+H),\\
W_0 &= \left\lceil \frac{(\Gamma-b^2)_+}{b} \right\rceil + 2\min(b, \lceil\sqrt{\Gamma}\rceil),\\
M_0 &= J_0 + \Gamma + m(W_0+1).
\end{aligned}$$

If $M \geq M_0$, the exact full time is given by the scalar criterion (30).

Proof. At any layer whose smaller solo capacity is $B \geq J_0$, an exceptional pair has a unique larger coordinate $a > K$ and a smaller coordinate $k \leq K$. Its larger coordinate satisfies $a \geq B+1-K \geq H+1$. Follow its initial-cohort label. The other coordinate is at most $K+m$ on the next day. If the labelled coordinate fell to at most K, the successor total would be at most $2K+m < J_0$, so the successor would not be exceptional. Thus a surviving exception cannot swap its numerical orientation. New mixed births from a pure endpoint are correctly oriented.

Suppose an exception is wrongly oriented: its numerical larger coordinate is its ancestry-secondary. By Lemma 7.3, its corresponding solo capacity is the smaller one, B. Set $\ell = B-a$ and $F = \ell - k = B-(a+k)$. Coordinate domination, concentration and exceptionality give

$$\ell \geq 0, \qquad -\Gamma \leq F < 0.$$

Allocate u of the next day's probes to the smaller coordinate. The larger inverse input is at least H; the expansive inequality in Lemma 7.8 gives $\ell' \geq \ell + u$. The smaller input is at most $M - H$, so its inverse is the bottom inverse. Its positive output k' loses at least $\min(\lceil \sqrt{k'} \rceil, b)$ from its input. Consequently

$$F' \geq F + \min(\lceil \sqrt{k'} \rceil, b).$$

While the exception remains wrongly oriented, $v = -F$ and $v' = -F'$ are positive integers, $v \leq \Gamma$, and $k' \geq v'$. Hence

$$v' \leq v - \min(\lceil \sqrt{v'} \rceil, b).$$

Above b^2 this removes at least b per transition. Once $v \leq r^2$ with $r \leq b$, it cannot remain above $(r-1)^2$ for two more transitions: two decrements of at least r would give $v'' \leq r^2 - 2r < (r-1)^2$. Thus no wrong exception survives W_0 such transitions. A change of faster initial cohort erases the history by fastest ancestry, so it cannot evade this argument.

Let t_0 be the first layer with both solo capacities at least J_0. Before that layer their maximum is at most $J_0 + \Gamma - 1$, so no capture can bypass it under the displayed bound on M. At entry the maximum is at most $J_0 + \Gamma + m - 1$. For the next W_0 days it is at most $J_0 + \Gamma + m - 1 + ms \leq M - 1$. These layers are proper: at a putative first captured layer the same estimate bounds every input by $M - 1$, below either saturation threshold. Both solo capacities remain at least J_0.

At $t_0 + W_0$ every surviving exception is correctly oriented, and the no-swap argument preserves this until the precentral deadline. Its common ancestry-secondary is therefore at most K. Two such states leave at least $M - 2K > m$ rooms in that color. Pairs with a nonexceptional state cannot improve the scalar test by Lemma 7.2; the pure endpoints give sufficiency. This proves the assertion. $\square$

The original bound $K = b^2 + 2m(\lfloor (m-1)/(2b+1) \rfloor + 1)$ is always available. An explicit alternative is obtained from

$$q = \max\left(\lceil \sqrt{b} \rceil, \left\lceil \frac{m-1}{b-1} \right\rceil \right), \qquad r = \left\lceil \frac{q(b+1) + m - 1}{2q} \right\rceil.$$

Here $q \geq 2$, $r \leq b$, and $q(2r - b - 1) \geq m - 1$, so Lemma 30.3 applies. Taking the minimum of a fixed number of such explicit bounds introduces no variable optimization. The theorem also applies directly if, for $L = \tau - 1 \geq W_0$, both capacities at $L - W_0$ are at least J_0.

Theorem 30.5 (Every odd rectangle up to a linear budget)**.** *For every odd $n \geq w = 2b + 1 \geq 3$ and $b + 1 \leq m \leq 2b$, the scalar criterion* (30) *gives the exact full time, attained by the canonical prefix construction.*

Proof. For $b \geq 160$ choose

$$q = 8, \qquad r = \left\lceil \frac{8(b+1) + m - 1}{16} \right\rceil, \qquad K_8 = 7m + r(r-1).$$

The preceding lemma applies and $K_8 \geq m$. Every term of the onset is nondecreasing in m, so it suffices to consider $m = 2b$. There $r \leq (5b + 11)/8$ and

$$K_8 \leq \frac{25b^2 + 966b + 33}{64} \leq \frac{b^2}{2}.$$

The last inequality follows from $7b^2 - 966b - 33 \ge 0$ for $b \ge 160$. Thus $J_0 \le 3b^2/2 + 1$. Also

$$\left\lceil \sqrt{2b-1} \right\rceil \le \sqrt{2b-1} + 1 \le b/8 - 1, \qquad W_0 \le b/4 - 2;$$

squaring the positive sides reduces the second inequality to $b^2 - 160b + 320 \ge 0$. It follows that $M_0 \le 2b^2$, below the minimum board value $M = 2b(b+1)$. For $b = 140, \ldots, 159$, direct substitution at $m = 2b$ gives the positive margins

$$184, 205, 226, 427, 450, 473, 96, 117, 324, 347,$$
$$370, 583, 608, 633, 852, 879, 1102, 1131, 1160, 1389.$$

This proves the theorem for every $b \ge 140$.

The minimum-budget theorem covers $m = b + 1$. For the remaining pairs $2 \le b \le 139$, $b + 2 \le m \le 2b$, take the minimum of K_8 and the original coordinate bound and form M_0 above. Writing $n = 2b+1+2j$ gives $M = 2b(b+1) + (2b+1)j$. The complete finite complement is therefore

$$0 \le j < \max\left(0, \left\lceil \frac{M_0 - 2b(b+1)}{2b+1} \right\rceil\right).$$

It contains exactly 19 852 triples in 3 171 nonempty budget–width intervals. A complete independent integer certificate checks all of them by Lemma 7.2: start at $(0, 0)$, apply every quota split, retain the two pure endpoints and all mixed exceptions, prune only coordinatewise dominated pairs, and examine every pair of final frontier points in (28). Its exact central decision agrees with the scalar criterion on every input. No temporal acceleration or scalar filtering is used.

The source is `src/bridge_odd_clock_linear_verify_full.cpp`, with coverage regenerated by its companion Python runner. The receipt `research/bridge-odd-clock-linear-budget-whole-verified.json` records every interval, full input, exact central cost and source hash. It contains 6 443 749 161 quota transitions and has maximum frontier size 422. The inverse lookup and per-coordinate maxima in the implementation are exact rearrangements of the displayed recurrence and Pareto pruning. This finite complement, together with the ordinary onset, covers every remaining length. □

Corollary 30.6 (A superlinear budget region without a width cutoff)**.** *The same exact scalar criterion holds on every odd rectangle when $m \ge b + 1$ and $27m^3 \le b^4$.*

Proof. If $m \le 2b$, use the preceding theorem. Otherwise the cubic inequality forces $b > 216$, so $b^{1/6} \ge 12/5$ and $\sqrt{b} \ge 14$. In particular $m \le (5/36)b^{3/2}$. Take $q = \left\lceil \sqrt{b} \right\rceil$ and $r = \lceil (q(b+1) + m - 1)/(2q) \rceil$. Then $r \le b$, $q(2r - b - 1) \ge m - 1$, and

$$r \le 41b/72 + 3/2 \le 3b/5, \qquad K_{q,r} \le \frac{449}{900}b^2 < \frac{b^2}{2}.$$

Also $m \le 5b^2/504 < b^2/100$. Hence $J_0 \le 3b^2/2 + 1$, $W_0 \le 2\sqrt{m} + 2$, and $2m^{3/2} \le 2b^2/\sqrt{27} < 2b^2/5$. Therefore

$$M_0 \le J_0 + 2m^{3/2} + 4m - 1 < \frac{97}{50}b^2 < 2b^2.$$

Every board satisfies this onset, proving the corollary. □

These theorems settle the scalar decision and provide optimal strategies on the stated families. Their solo clocks still require $O(b)$ arithmetic stages in general. An absolute bounded expression follows only in a separately proved numerical region, such as the intersection with Theorem 30.1; the decision theorem alone does not remove that numerical distinction.

31 Fixed numerical expressions for further even-area rectangles

31.1 Uniform numerical consequences in widths 24 and 25

The three- and nine-threshold recurrences can be evaluated uniformly in the longer side for the next two widths. A fixed entrance and exit surround an arbitrarily long plateau whose transfer is explicit. The joint physical restriction in Theorem 29.13 is essential to these bounds.

Lemma 31.1 (A backward frontier after the joint restriction)**.** *Consider a rectangle with* $2h$ *rooms, both sides at least* 24*, and budget* 13*. Use either necessary model of Theorems 29.1 and 29.2, with its guarded merger. Let* B_t *be its capture frontier, and let* R_t *be its backward frontier for reaching a state of size at most* 6*. Suppose that, for an integer* $N \geq 1$*,*

$$\max B_{29} < h, \qquad \max B_N \leq h - 112, \qquad \max R_N \leq h - 112.$$

Then every full-board belief after $L = N + 28$ *steps has at least* $h + 7$ *rooms, and* $T_{13} \geq 2N + 58$*.*

Proof. Write a_t for the physical joint belief size minus h. The guarded merger maps the history to a model history while $a_t > 13$, including the transition at its first crossing to $a_t \leq 13$. Such a crossing by time 28 would permit a model capture from the full size h by time 29, contradicting the first hypothesis.

Theorem 29.13 gives $a_{28} \geq h - 111$. If the first crossing to at most 13 occurred at some $t \leq L - 1$, its merged continuation and one clearing inspection would give a model capture from that time-28 feature within $t - 28 + 1 \leq N$ days. This contradicts $\max B_N \leq h - 112$. The merger is therefore valid through time L. A value $a_L \leq 6$ would contradict the bound on R_N, so $a_L \geq 7$.

In any proposed schedule of length $2L + 1$, the first L inspections and the last L inspections read backward each leave a belief of size at least $h + 7$ at the middle day. Undirected walk reversal justifies the second belief. Their intersection has at least $2(h + 7) - 2h = 14$ rooms, and the middle inspection cannot cover it. An avoiding walk exists, proving $T_{13} \geq 2L + 2 = 2N + 58$. $\square$

Corollary 31.2 (A one-day bracket for even lengths at width 24)**.** *For every integer* $n \geq 24$*,*

$$T_{13}(P_{24} \,\square\, P_n) \geq 24n - 370.$$

For even $n \geq 24$ *this gives the numerical bracket*

$$24n - 370 \leq T_{13}(P_{24} \,\square\, P_n) \leq 24n - 369.$$

Proof. Set $h = 12n \geq 288$. Use the original, unrefined corner model, so the unrestricted surplus in Theorem 29.1 is 12 for every corner pair. For its capacities, direct substitution gives

$$(11 + F_{ij}(11))_{ij} = \begin{pmatrix} 110 & 121 & 85 \\ 132 & 102 & 70 \\ 112 & 85 & 57 \end{pmatrix}, \qquad (G_{ij}(11))_{ij} = \begin{pmatrix} 65 & 81 & 90 \\ 82 & 100 & 110 \\ 101 & 121 & 99 \end{pmatrix}.$$

Thus $S_{ij}(z) \geq 12$ for $z \geq 133$, and $H_{ij}(d) \geq 12$ for $d \geq 122$, for every i, j.

Start the capture and half-budget frontiers at

$$B_0 = (0, -1, -1), \qquad R_0 = (6, 6, 6).$$

The fixed first 36 substitutions in (161) give

t	B_t	R_t
29	$(124, 125, 125)$	$(125, 125, 125)$
35	$(133, 134, 134)$	$(134, 134, 134)$
36	$(135, 135, 135)$	$(135, 135, 135)$.

These substitutions are independent of $h \geq 288$: every target used before step 36 has size at most 134, hence at least $154 > 121$ holes. The high branch cannot improve surplus 12, and all state caps are inactive. The entries follow by substituting the small quadratic capacities at the fixed surpluses $0, \ldots, 12$. Both complete entrance traces are retained in `research/bridge-boundary-numeric-24-check.json`.

If all three thresholds equal a, with $135 \leq a \leq h - 122$, both capacity bounds apply. Every corner pair has minimum surplus exactly 12, and (161) becomes

$$(a, a, a) \longmapsto (a + 1, a + 1, a + 1).$$

Both frontiers therefore reach the all-$(h - 121)$ vector at time $36 + (h - 121 - 135) = h - 220$.

Write a frontier as $(h - d_0, h - d_1, h - d_2)$. From this point the small branch is inactive below surplus 12, since every target is at least $h - 121 \geq 167 > 132$. With the state caps inactive, the exact update is the h-independent expression

$$d'_c = \min_{i \leq c,\, 0 \leq j \leq 2} (d_j + \min(H_{ij}(d_j), 12)) - 13.$$

Its next seven substitutions are

t	d	t	d
0	$(121, 121, 121)$	4	$(115, 115, 115)$
1	$(120, 120, 119)$	5	$(114, 114, 113)$
2	$(118, 118, 118)$	6	$(112, 112, 112)$
3	$(117, 117, 116)$	7	$(111, 111, 110)$.

All deficits remain at least 110, so the stated absence of state caps holds throughout. Consequently, for $N = h - 214$,

$$B_N = R_N = (h - 112, h - 112, h - 112).$$

The initial table also gives $\max B_{29} = 125 < h$. Lemma 31.1 proves $T_{13} \geq 2(h - 214) + 58 = 24n - 370$. The upper bound for even lengths is Corollary 29.16. $\square$

The arithmetic certificate for this corollary derives the entrance and exit directly from the quadratic inequalities, independently of the inverse-root formulas. It also compares every state membership in 76 backward layers for both seeds with the full original model graph on the 24-square: all 130 872 memberships agree. The varying-length conclusion follows from the plateau identity, not from testing several lengths.

Corollary 31.3 (An exact value at width 25)**.** *For every even $n \geq 26$,*

$$T_{13}(P_{25} \,\square\, P_n) = 50n - 925.$$

Proof. Put $h = 25n/2 \geq 325$ and use the nine-feature model of Theorem 29.2. Its indices are (c, e), and its closed surplus is σ; the original surplus is $\sigma + \delta$, where $\delta = v - i \geq 0$. The unrestricted closed surplus is 13, with the critical value 12 allowed for $(v, j) = (2, 0)$. Direct substitution in the capacities gives

$$\max_{v,u,j}(12 + P(v, u, j; 12)) = 156, \qquad \max_{A,B} F_{AB}(12) = 144,$$
$$\max_{j,\, v \leq 1} F_{2-j,2-v}(12) = 132, \qquad \max_{v} F_{0,2-v}(12) = 132,$$
$$\max_{A,B} F_{AB}(11) = 121.$$

All corner indices range from zero to two, with $u \leq j$. Monotonicity extends these bounds to smaller surpluses.

Initialize $B_0(c, e) = 0$ and $R_0(c, e) = 6$. The fixed first 70 substitutions of (162) give

$$B_{70}(c, e) = R_{70}(c, e) = 157 \qquad \text{for all } (c, e).$$

Every target used in this entrance has size at most 157, hence at least $168 > 144$ holes. The high branch is inactive below closed surplus 13, all state caps are inactive, and no all-radius band applies: each such band has at most $12 \cdot 11 = 132$ holes. Thus these fixed substitutions are the same for every $h \geq 325$. The two complete traces, obtained directly from P and the fixed surpluses $0, \ldots, 13$, are recorded in `research/bridge-boundary-numeric-25-check.json`.

For $157 \leq a \leq h - 134$, let $E(a)$ be the all-a vector, and let $C(a)$ equal a when $c = 0, 1$ and $a + 1$ when $c = 2$, independently of e. We claim the exact two-step identity

$$E(a) \longmapsto C(a) \longmapsto E(a + 1). \tag{172}$$

For the first step, the low branch excludes every $\sigma \leq 12$ since $a > 156$, and the high branch excludes every $\sigma \leq 11$ since there are at least $134 > 121$ holes. If $v \leq 1$, the high branch also excludes $\sigma = 12$ by the bound 132. A source with $c \leq 1$ therefore always pays original surplus at least 13: when $v = 2$, it has $\delta = v - i \geq 1$; otherwise its closed surplus is at least 13. Its threshold cannot increase. A source with $c = 2$ attains an increase of one by choosing $i = v = 2$, $u = j = 0$, and critical surplus 12. Every source attains the unchanged threshold by choosing $i = v = u = j = 0$ and surplus 13. These choices satisfy survivor maximality and the deletion charge.

For the second step, targets with $j = 0, 1$ still have size a, so surplus at least 12 bounds their contribution by $a + 1$. Targets with $j = 2$ have size $a + 1$ and at least 133 holes. Their high capacity at $\sigma = 12$ is at most 132, so they require closed surplus at least 13 and likewise contribute at most $a + 1$. Equality in every source class is attained with $(j, v) = (2, 0)$, $i = u = 0$, and surplus 13. All deficits in these two steps are at least 133, so no band filter intervenes; all state caps are inactive and the deletion charges are at most $4 \leq 13$. This proves (172).

There are $h - 290$ such pairs from $E(157)$ to $E(h - 133)$. Both frontiers reach the latter vector at time $2h - 510$. Subsequent target sizes are at least $h - 133 \geq 192 > 156$, so the low branch remains inactive below closed surplus 13. In deficit coordinates every update now depends only on the fixed high capacities, critical exception, and all-radius bands, independently of h. Nineteen fixed substitutions starting from the all-133 deficit vector give, at steps 18 and 19,

$$d_{18} = (113, 113, 113, 112, 112, 112, 112, 112, 112),$$
$$d_{19} = (112, 112, 112, 111, 111, 111, 110, 110, 110),$$

in feature order $(0,0),(0,1),(0,2),(1,0),\ldots,(2,2)$. Every earlier deficit is at least 112. The full exit table is in the same certificate; it checks all radii $2,\ldots,12$ directly. Deficits remain at least 110, so no state cap is active. It follows that for $N = 2h - 492$,

$$\max B_N = \max R_N = h - 112.$$

Also $\max B_{29} \leq \max B_{70} = 157 < h$, by monotonicity of the capture sets. Lemma 31.1 would give the preliminary bound $50n - 926$. We now use the persistent radius information of Theorem 29.10 to gain one day.

The following fixed terminal certificate is the additional input. In the persistent model with the 25-by-26 radius thresholds, none of the three features of size 157 with $c = 0$ and $e = 0,1,2$ clears in 70 days, even when both incoming radii are unknown. Their minimum time is 71. The independent verifier `src/bridge_upper_transport_verify_reach.py` retains the complete grouped-radius reachability set, without radius dominance; its receipt is `research/bridge-upper-transport-terminal-feature-verified.json`. The old 70-day frontier already bounds every successful suffix state by size 157. At those sizes there are at least 168 holes, so the high branches and their radius bands are inactive. Freezing the larger-board radius thresholds at $n = 26$ and replacing radii at least 49 by an unknown radius only weakens the necessary model. Thus the same finite low graph and exclusion apply at every $h \geq 325$.

The persistent 70-day capturable set is consequently contained in the numerical hull D, with threshold 156 for $c = 0$ and 157 for $c = 1,2$, independently of e. This is a containing hull, not an assertion of interval closure for flagged states. One old nine-frontier predecessor maps this hull to $E(157)$. Indeed the maximum low target capacity at closed surplus 11 is 132, and at surplus 12 it is 156. Targets of size 157 require closed surplus 13; the sole critical exception $(v,j) = (2,0)$ instead uses the hull's target 156, with original surplus at least 12. Every predecessor contribution is therefore at most 157. Equality for every source feature is attained using target $(j,v) = (1,0)$ of size 157, $i = u = 0$, and surplus 13.

We may therefore start the necessary numerical hull $E(157)$ at time 71, apply the same $h-290$ plateau pairs and then the first 18 exit steps. At the model deadline $N = 2h - 491$ its maximum is $h - 112$. No persistent state of size at least $h - 111$ can clear within N days, whatever its incoming radii.

Put $L = N + 28 = 2h - 463$ and again write a_t for joint belief size minus h. An early first crossing to $a_t \leq 13$ by time 28 would contradict $B_{29} < h$. The joint boundary theorem gives $a_{28} \geq h - 111$. A first crossing by any $t \leq L - 1$ would give a persistent-model capture from that state in at most $t - 28 + 1 \leq N$ days, a contradiction. Hence $a_{L-1} > 13$. The next merged survivor is nonempty, and the perfect matching ensures $a_L > 0$.

A proposed $2L$-day full strategy now fails on an even cut. Its first L inspections and moves leave more than h rooms. Read its last L inspections backward: after $L-1$ moves the belief has more than $h + 13$ rooms, and the last inspection still leaves more than h. These two beliefs occupy the same cut in a graph of $2h$ rooms, so their avoiding walks intersect and concatenate. Thus $T_{13} \geq 2L + 1 = 50n - 925$.

For attainment use the construction underlying Proposition 21.2. At $b = 12$, its constant is $L_{12} = 46$, and both initial prefix phases reach 157 rooms after

$$t_* = 46 + 2h - 4 \cdot 12^2 - 2 = 25n - 532$$

steps, in opposite current phases. All subsequent counts are at most 157 and have at least 168 holes, so the two terminal profiles are independent of n. Their fixed evaluations take 70 and 71 days; the faster phase's last shot has size 2. These evaluations are recorded in `research/bridge-boundary-numeric-25-exact-check.json`. Choose the faster initial phase. Its solo duration is $25n-462$. Reflection along the even side supplies the opposite cohort, and the two central residuals have union at most $4 \le 13$. The resulting full strategy has length $2(25n-462)-1 = 50n-925$. $\square$

The width-25 certificate evaluates both fixed portions directly through the quadratic inequalities and independently compares each substitution with the inverse-root nine-threshold update. The unbounded part is the two-step identity (172). The additional persistent terminal exclusion makes width 25 exact; the width-24 statement remains a one-day bracket.

31.2 The quadratic-edge budget curve

Theorem 31.4 (The first budget below the strict quadratic threshold)**.** *Let $w = 2r \ge 6$, $n \ge 2r$, $h = rn$, $k = r(r-1)$ and $s = r(r-2)$. Define*

$$j = \max\left(0, \left\lfloor \frac{h-k-s-3}{s} \right\rfloor\right), \qquad z = h-(j+1)s-k-1, \qquad c = z + \left\lceil \sqrt{z} \right\rceil .$$

Then, for every such r, n,

$$T_k(P_{2r} \,\square\, P_n) = 2(j+3) - \mathbf{1}_{\{2c \le k\}}.$$

This is an absolute fixed-size numerical expression with a directly specified optimal strategy.

Proof. The scalar first step from h is $h-s-1$: its reflected pronic root is $r-1$, and its increasing root is saturated. All later reflected roots remain saturated, so the scalar profile is $L(q) = q + \min(\lceil \sqrt{q} \rceil, r)$. Its full-budget map is $x \mapsto x - s$ for $x \ge k+s+2$. Exactly j such steps follow the first one; the remaining source is $k+z$. Here

$$z \equiv -1 \pmod r, \qquad r-1 \le z \le s-1.$$

The next count is $c \le k-2$ and the following inspection clears it. No earlier capture was omitted: the first source after the initial step is at least $k+3r-1$, and the last affine source is still above k. Thus the scalar duration is $\tau = j+3$, and the symmetric scalar lower bound is the displayed expression.

Choose the initial anchored-prefix phase $(n+1) \bmod 2$. After deleting k, its survivor is the full prefix through diagonal $n-1$, of size $h-k$. The preceding opposite diagonals have r fewer rooms and the new last diagonal has $2r-1$, so the first actual neighborhood is $h-s-1$. The next j steps lie in the common anchored plateau and decrease by s. The final survivor z has near-corner neighborhood either $c = z + R(z)$ or $c' = z + Q(z)$, with $R \le Q \le R+1$ and $c' \le k-1$. This strategy therefore has the same duration τ.

If $2c < k$, then k is even and $c' \le k/2$, so reflection attains $2\tau - 1$. If $2c > k$, the doubled solo strategy attains 2τ. Only $c = k/2$ needs attention. The congruence for z rules this out when r is odd. For $r = 2u$ it forces $R(z) = u+1$ and $z = 2u(u-1)-1$. The root inequalities hold only for $u = 3, 4$, hence $r = 6, 8$. At $r = 6$, $Q(11) = R(11) = 4$. At $r = 8$, $z = 23$, $R = 5$ and $Q = 6$; its defining length is $n = 6j + 16$. The final phase is $j \bmod 2$, so only positive odd j needs a change of shape.

For that class use coordinates $0 \le x < 16$, $0 \le y < n$. After j prefix steps the same color-zero set A of size 127 is reached on every board:

$$A = \{(x,y) : x+y \text{ even},\ x+y \le 22\} \setminus \{(0,22)\},$$

with intersection with the board understood. Its column heights are $20, 21, 20, 19, \ldots, 7$. The set

$$U = \{(x, y) : x + y \text{ even}, \; y \leq 8\} \setminus \{(0, 8)\}$$

has 71 rooms and lies in A. Inspect $A \setminus U$, using 56 probes. Its neighborhood is the 79-room set

$$B = \{(x, y) : x + y \text{ odd}, \; y \leq 9\} \setminus \{(0, 9)\}.$$

Inside it take the 25-room opposite-corner triangle $Q = \{(x, y) : x + y \text{ odd}, \; (15 - x) + y \leq 8\}$ and delete $(14, 7), (15, 8)$ to leave V of size 23. Inspect $B \setminus V$, again using 56 probes. The nonempty column heights of V at $x = 7, \ldots, 15$ are $0, 1, 2, 3, 4, 5, 6, 5, 6$; those of its neighborhood at $x = 6, \ldots, 15$ are $0, 1, 2, 3, 4, 5, 6, 7, 6, 7$. Thus $|N(V)| = 28$. The two-day transition is $127 \to 79 \to 28$, and the next shot clears. Reflection across the short side switches colors, so the central union has at most 56 rooms. The resulting full duration $2j + 5$ equals the scalar lower bound, closing the last residue. $\square$

The arithmetic audit is `research/bridge-boundary-numeric-quadratic-edge-check.json`. The exceptional physical transition is independently replayed in `research/bridge-quadratic-edge-exception-splice.json` and `research/bridge-upper-transport-quadratic-boundary-verified.json`. The ordinary plateau and fixed-set construction supply the unbounded claims; the finite examples are verification receipts.

31.3 Fixed-block numerical bounds at arbitrary even-width budgets

Theorem 31.5 (Absolute block bounds for the model and physical clocks)**.** *Fix $K \geq 1$ independently of the inputs. Let $w = 2r \geq 4$, $n \geq w$, $h = rn$, $r + 2 \leq k < h$, $s = k - r$ and $\delta = s - 1$. If $K\delta \geq r^2$, a capture or half-budget backward frontier of the critical three-threshold model, at any deadline or at its first admission of the full state, uses at most $3K + 3$ ordinary three-coordinate updates and one bulk division. A physical anchored prefix clock, including its last shot or a prescribed-day residual, uses at most $2K + 1$ ordinary updates and one bulk division under the weaker hypothesis $Ks \geq r^2$.*

Thus $K = 16$ gives explicit lower and physical upper expressions using at most 168 such blocks and four divisions in total. Equality of the resulting bounds specifies an exact numerical subfamily and an optimal physical strategy; a strict gap remains a bracket.

Proof. Use Theorem 29.1 with unrestricted cap $r + 1$ and critical cap r at $(i, j) = (1, 1)$. The extra cap at $(2, 0)$ when $n = 2r + 1$ is retained. After the first update the thresholds obey

$$B_0 \leq B_1 \leq B_2, \qquad B_1 \leq B_0 + 1, \qquad B_2 \leq B_1 + 1.$$

For the spread bound, delete one current survivor corner when moving to the next smaller source class and increase surplus by one. Direct substitution in the capacities gives $F_{i-1,j}(z+1) \geq F_{ij}(z) - 1$ and $G_{i-1,j}(z + 1) \geq G_{ij}(z)$, preserving both alternatives. Removing either critical corner pattern permits the raised unrestricted surplus $r + 1$. The state caps differ by one. If the survivor had size one, deletion leaves zero and the bound is immediate.

Choosing survivor corner class zero and target threshold B_2 with unrestricted surplus $r + 1$ gives

$$B_2' \geq \min(h, B_2 + \delta).$$

The first update has $B_2 \geq k$. Therefore, until full admission, after $t \geq 1$ updates one has $B_2 \geq t\delta + r + 1$ and $B_0 \geq t\delta + r - 1$. When $B_2 = h$, the spread forces the absorbing cap vector $(h-2, h-1, h)$.

Set $A = r^2 + r + 1$, $H = r^2 + 1$ and $U = h - \max(H, s+1)$. On a source vector (x, y, z) whose coordinates all lie in $[A, U]$, the low and high capacity roots exceed r, and the state caps are inactive. The exact update is

$$(x, y, z) \longmapsto (z + \delta, \max(z, y+1) + \delta, \max(z, y+1) + \delta).$$

The extra $(2, 0)$ exception does not alter this formula: for $n \leq 2r + 1$, $U \leq r^2 + r - 1 < A$, so the interval is empty. Every nonempty common interval has only the $(1, 1)$ exception. In at most two updates the vector is normalized to $V(a) = (a-1, a, a)$, unless it already exits the interval, and $V(a) \mapsto V(a+s)$ there.

At most $2K$ conditional updates suffice to reach $B_0 \geq A$ or admit the full state. After at most two normalizations, use $1 + \lfloor (U-a)/s \rfloor$ translations at once when the source is $V(a)$ in the common interval. Upon exiting, the remaining distance $h - B_2$ is at most $\max(r^2, s)$, so at most $K+1$ further updates suffice. The total is $3K+3$. An exhausted deadline stops the conditional blocks and truncates the bulk division; an empty interval simply omits that bulk operation.

For physical prefixes, Lemma 29.14 gives $g_p(q) = q + r$ in both phases for $(r-1)^2 + 1 \leq q \leq h - r^2 + r - 1$. Lemma 12.2 also gives $g_p(q) \leq q + r$ globally: the height-plus-one envelope adds at most one room in each of the r transverse columns. Thus a positive full-budget step decreases size by at least s, and is exactly $x \mapsto x - s$ on

$$k + (r-1)^2 + 1 \leq x \leq k + h - r^2 + r - 1.$$

The full starting excess above this interval is at most r^2, so K ordinary steps suffice before a single bulk division. The remaining size is at most $s + r^2$, so $K+1$ more steps suffice. Retain the parity of the number of bulk steps and the last preinspection count. This proves the physical block bound, also for arbitrary prefix sizes and prescribed deadlines.

Let τ be the model's first admission time of $(h, 2)$ and $R_{\tau-1}$ the backward frontier with seed $(\lfloor k/2 \rfloor)^3$. The guarded-merger clock gives

$$L = 2\tau - 1 + \mathbf{1}_{\{R_{\tau-1}(2) < h\}}.$$

For each physical starting phase let t_p and a_p be the prefix duration and last shot. Reflection across the even short side gives

$$U = \min_{p=0,1} \left(2t_p - \mathbf{1}_{\{2a_p \leq k\}}\right).$$

The minimum with the explicit envelope upper of Section 13.1 is also valid. Two model and two prefix evaluations use $2(3K+3) + 2(2K+1) = 10K + 8$ blocks, hence 168 at $K = 16$. Whenever $L = U$, the corresponding physical construction is optimal. No model path is asserted to be physically attainable when the bounds differ. □

The ordinary update, normalization and phase arithmetic are audited in `research/bridge-boundary-numeric-budget-check.json`: 6 811 deadline comparisons on 140 finite inputs agree with the direct recurrences, along with large-length checks. The absolute block count is the theorem above, rather than an extrapolation from those comparisons.

32 Formal verification, reproducibility, and open problems

The shared canonical development comprises 83 source modules using Lean 4.33.1 and Std. Its main and parked application roles are listed separately in the accompanying source manifest; these are not 83 complete rectangle classification theorems. The rectangle package contains 43 modules; the parked package contains 51, with 11 shared between them. The split checks source hashes and import closure against the existing compilation receipt; it is not a new compilation or formalization. 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.

32.1 Complete classifications and reusable formal theorems

Two families in the rectangle scope have complete physical-game classifications:

- `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.

The complete three-cube and four-cube classifications are preserved in the parked companion and use the same authoritative Lean source tree. The two rectangle classifications each include 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 and finite stabilization have reusable formal proofs. The concrete square operators used for the four-cube theorem are documented in the companion. 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.

32.2 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; the rectangle applications of the cylinder time and eventual-period arguments; and fixed-deadline semilinearity. The companion separately states the ordinary proof boundary for higher-dimensional applications. 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.

32.3 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 numerical target is an absolute $O(1)$ number of standard numerical operations for $T_k(a, b)$, independent of all three parameters, together with directly specified optimal inspections. The width-only minimum- budget clock is an evaluation gap under this criterion, even though it proves exact time and its dependence on length. The bounded odd frontier, eventual solo criteria, effective affine periods and finite boundary representations also retain parameter-dependent evaluation obligations. On shorter odd rectangles, the simpler proposed middle-day decision is not established in all intermediate budgets. On even-area rectangles, the boundary optimization is not uniformly evaluated. A proof of an optimal full-board strategy family would suffice; arbitrary partial-state optimality is not required.

The split companion at `https://angelraychev.com/princess/extensions/` is parked. Its open higher-dimensional, general-cylinder and independent partial-state questions are not requirements for completing the rectangle paper. The dated combined release remains available as a historical archive.

32.4 Research and writing provenance

The new finite-board corner and erosion bounds, critical extensions, near-extremal localization, and merger theorems are ordinary proofs. Their finite applications use independently checked integer certificates and physical inspection replays. None of these newly added arguments has a complete physical-game Lean formalization in this release. The unchanged formal sources retain their source-specific earlier compiler evidence.

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.

[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.

[13] 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.

[14] 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.