

TIMED GUARDING IN DEGREE-REDUCTION CLOUDS

LEVI NEUWIRTH

ABSTRACT. We prove a conditional cop-number transfer for the actual all-vertex Hosseini–Mohar–Gonzalez Hermosillo de la Maza degree-reduction tower. If its connected simple base has order N and satisfies the fixed-parameter five-layer multistage hypotheses, the stopped tower H satisfies $c(H) \leq C\sqrt{N}$, with C independent of cloud order and tower scale. Each ancestry cloud is a retract. Predetermined port placement improves the deployment bound by exactly $\lambda - 1$ moves, paying for one cop to establish an exterior shadow guard. During synchronized setup the robber visits at most two clouds; a reusable internal-pursuit team handles those exceptions during cleanup and finishing. The proof retains the full remaining-ball Hall demand and establishes finite capture, without the previous occupation-based capture-time bound. A 104-vertex stopped tower obstructs five-turn handoffs released only from the old side. Separately, a Hall-capacity inequality shows why full-region demand remains incompatible with the proposed slow-growth model even without thin-sphere bounds. These results concern specified replacement clouds over qualifying bases and do not establish an all-graphs bound.

1. INTRODUCTION

A region may require many cops to occupy and only one cop to guard. For a graph retract, a pursuit team can catch the retracted image of the robber and leave one cop tracking that image. Turning this structural fact into a timed strategy requires two further ingredients: a legal route to the guard’s starting position, and a setup schedule that works while the robber continues to move. This paper supplies both ingredients for the ancestry clouds of the degree-reduction construction of Hosseini, Mohar, and Gonzalez Hermosillo de la Maza [1].

The companion occupation paper [5] transfers the multistage strategy of Prałat and Wormald [3] through abstract coarse projections. With base order N and cloud order at most P , it uses squads to fill fibers and obtains $CP\sqrt{N}$. Its deployment bound at base radius t is $\lambda(t + 2) - 2$. Here a cop initially placed at a designated port reaches a target port in at most $\lambda t + (\lambda - 1)$ moves. The remaining $\lambda - 1$ moves suffice to establish a guard while the robber stays outside that target cloud. Thus

$$\underbrace{\lambda t + (\lambda - 1)}_{\text{port travel}} + \underbrace{(\lambda - 1)}_{\text{exterior setup}} = \lambda(t + 2) - 2.$$

This identity consumes the old deployment allowance. It does not exhaust the entry deadline: the proof retains four moves before earliest entry, and at least two after root-dependent activation is included.

The exterior condition during setup requires a separate argument for finishing balls. In a k -round tower, let $\lambda = 3^k$ and $L = \lambda - 1$. A walk of at most L edges can visit at most two original ancestry clouds. Simultaneous setup therefore guards every unvisited cloud. At the finishing deadline the boundary is still unvisited, so the at most two exceptions are interior. Simplicity of the base leaves one joining edge between two adjacent exceptions; one reserve cop blocks it while an internal-pursuit team captures on the observed side.

Date: September 9, 2026.

2020 Mathematics Subject Classification. 05C57, 05C12, 05C40.

Key words and phrases. Cops and Robber; graph retractions; degree reduction; timed guarding; accessible reservoirs.

The quantitative conclusion before absorbing the reserve is

$$c(H) \leq C_0 \sqrt{N} + K + 1, \quad K = \max_{v \in V(G)} c(H[F_v]) \leq 2\Delta(G).$$

The same fixed-parameter growth hypotheses imply $\Delta(G) = O(\sqrt{N})$, including degrees in the exceptional set. Consequently $c(H) \leq C\sqrt{N}$, and hence also $c(H) \leq C\sqrt{|V(H)|}$ on this class. The constant does not depend on P , λ , or $\Delta(G)$. The formal base hypotheses are stated in full below.

1.1. Attribution and scope. The degree-reduction gadget and its lower cop-number transfer are due to Hosseini–Mohar–Gonzalez Hermosillo de la Maza [1]. Theorem 2.7 gives a self-contained proof of that comparison for the entire tower, with the cop and robber turns explicit. The multistage radius schedule, fresh random teams, contracting uncovered frontiers, and finishing Hall estimate are inherited from Pralat–Wormald [3]. The five-layer accessibility hypothesis and its fixed-degree random-regular application are developed in the companion paper [5], using intermediate levels in [4]. We keep that hypothesis unchanged. The retract setup principle is standard; the proof is included to make its resource account explicit [2]. The two-cop bound inside one gadget is proved directly.

The cloud retractions, port travel, synchronized exterior setup, and their use in this timed transfer are the derivations studied here. The theorem supplies neither a simulation of every base pursuit strategy nor a stronger bound for arbitrary coarse projections. The branching parameter remains fixed. No uniformity for a growing branching parameter or shrinking expansion is asserted. The full remaining-ball Hall condition is retained; Section 9 proves a concrete obstruction to one proposed weakening. No novelty claim follows merely from the bounded source audit.

1.2. Game convention and notation. All graphs are finite, undirected, and simple. The cops choose their vertices before the robber chooses hers; multiple cops may share a vertex. Each cop may move along one edge or stay, followed by one robber move of the same kind. Capture occurs when a cop and the robber occupy the same vertex. The least winning cop count is $c(G)$. For $U \subseteq V(G)$, $B_G(U, r)$ and $S_G(U, r)$ are the closed radius- r neighborhood and the set at distance exactly r from U . A singleton argument is written as $B_G(v, r)$ or $S_G(v, r)$. A lazy graph homomorphism maps every edge to an edge or one vertex. A retraction onto an induced subgraph fixes that subgraph pointwise.

A retained guard need not occupy every vertex of its cloud. It remains inside the cloud and tracks a retracted robber image, so entry by the robber is captured no later than the next cop move. Phase roots and surviving exits are considered after this compulsory capture check.

2. THE REPLACEMENT GEOMETRY

We use the construction of Hosseini, Mohar, and Gonzalez Hermosillo de la Maza, with the all-vertex convention made explicit [1, Section 2]. All graphs in this section are finite and simple. Distances in an induced subgraph are intrinsic unless an ambient graph is specified.

Definition 2.1 (All-vertex HMGHM tower). For $r \geq 2$, form a graph $A_r(m)$ from r independent vertices, called *ports*, partitioned into m nonempty classes $X_1, \dots, X_m$ whose sizes differ by at most one. For each $1 \leq i < j \leq m$, add a vertex y_{ij} adjacent to every port in $X_i \cup X_j$, and add no other internal edges. Use

$$m(2) = 2, \quad m(3) = 3, \quad m(4) = 3, \quad m(r) = \lceil \sqrt{2r} \rceil \quad (r \geq 5).$$

Thus the degree-two gadget is a three-vertex path, the degree-three gadget is a six-cycle, and the degree-four gadget has class sizes 2, 1, 1 and seven vertices. A vertex of degree zero or one is replaced by a singleton; in the degree-one case that singleton is its port.

Starting from a connected base graph, in one round replace every vertex v by its gadget, associate its ports bijectively with its incident edges, and retain each old edge as one edge between the

corresponding ports. This includes vertices of degree two or three whenever a round is performed. Let

$$G = G_0, G_1, \dots, G_k = H$$

be k such rounds. The *ancestry cloud* F_v consists of all final vertices descended from $v \in V(G)$. Each edge incident with v determines one edge from F_v to another ancestry cloud and hence one designated port of F_v . Put

$$\lambda = 3^k, \quad L = \lambda - 1, \quad D_j = \Delta(G_j), \quad D = D_0.$$

The *stopping tower* stops at the first k with $D_k \leq 3$; in particular $k = 0$ if the base is already subcubic. Statements about a specified number of rounds do not require this stopping rule.

Every cloud is connected, and exactly one edge remains between the ancestry clouds of each adjacent pair of base vertices. These assertions follow inductively from the connectedness of the gadgets and the retention of every old edge. For $k = 0$, all ports associated with one base vertex coincide at that vertex.

Lemma 2.2 (Equilateral designated ports). *Suppose $k \geq 1$. In a cloud F_v , ports associated with two distinct base edges have intrinsic distance exactly $L = 3^k - 1$.*

Proof. In one gadget, ports in different classes share y_{ij} , and two ports in the same class share a pair vertex involving that class and another nonempty class. Ports are independent, so their distance is exactly two. This proves the assertion for one round.

Suppose the assertion holds after j rounds, with distance $L_j = 3^j - 1$. Consider two designated final ports after one further round and a shortest path between them within their enlarged ancestry cloud. Its sequence of edges between the new gadgets projects to a walk between the old designated ports. If that sequence uses q edges, then $q \geq L_j$. The first new designated port differs from the port used by the first internal old edge: the former represents an edge leaving the old ancestry cloud, whereas the latter represents an edge within it. The analogous statement holds at the other endpoint. Each endpoint traversal therefore costs at least two. Each intermediate traversal also joins distinct ports; otherwise the path immediately reverses an external edge and is not shortest. The new path has length at least

$$q + 2(q - 1) + 4 = 3q + 2 \geq 3L_j + 2.$$

Conversely, lift an old intrinsic geodesic and use length-two port paths at both endpoints and every intermediate gadget. This realizes $3L_j + 2 = 3^{j+1} - 1$, completing the induction. $\square$

For reference, the usual normalized metric bounds also hold without any stopping assumption. We give the details needed to distinguish them from the stronger endpoint statement used below.

Lemma 2.3 (Cloud metric bounds). *For every $v \in V(G)$,*

$$\text{diam}_{H[F_v]}(F_v) \leq 2L, \quad \text{dist}_{H[F_v]}(x, p) \leq \frac{3}{2}L \quad (x \in F_v, p \text{ a designated port}).$$

For distinct $u, v \in V(G)$, $x \in F_u$, and $y \in F_v$,

$$(1) \quad \lambda(\text{dist}_G(u, v) - 1) + 1 \leq \text{dist}_H(x, y) \leq \lambda(\text{dist}_G(u, v) + 2) - 2.$$

Proof. In a one-round gadget every vertex is within three of each specified port, and the diameter is at most four. For a pair vertex and a port outside its two classes, a three-edge route uses one intermediate port and a second pair vertex. Two pair vertices with intersecting class pairs have a common port; disjoint class pairs can be joined in four edges. Singleton gadgets satisfy the same upper bounds.

Lifting an old path of length q between arbitrary vertices costs at most $3q + 4$, including both endpoint costs; when both old endpoints coincide, the one-round diameter bound applies. If one new endpoint is a designated port for an edge leaving the ancestry cloud, its endpoint cost is at

most two instead of three, so lifting costs at most $3q + 3$. Consequently the intrinsic diameter and designated-port eccentricity satisfy the respective recurrences

$$d_{j+1} \leq 3d_j + 4, \quad e_{j+1} \leq 3e_j + 3, \quad d_0 = e_0 = 0.$$

Their solutions are $d_j \leq 2(3^j - 1)$ and $e_j \leq \frac{3}{2}(3^j - 1)$.

Let $r = \text{dist}_G(u, v)$. A shortest ambient path between F_u and F_v uses $q \geq r$ external edges. Every intermediate segment joins distinct designated ports, so Lemma 2.2 gives length at least $q + (q - 1)L \geq \lambda(r - 1) + 1$. For the upper bound follow a base geodesic. Its endpoint costs sum to at most $3L$, its intermediate costs sum to $(r - 1)L$, and it uses r external edges. The resulting length is at most $r + (r + 2)L = \lambda(r + 2) - 2$. When $k = 0$ the assertions reduce directly to ordinary base distances. $\square$

Lemma 2.4 (Designated-port travel). *Let $p \in F_z$ and $a \in F_w$ be designated ports, and suppose $\text{dist}_G(z, w) \leq t$. Then*

$$(2) \quad \text{dist}_H(p, a) \leq \lambda t + L = \lambda(t + 1) - 1.$$

This includes $z = w$ and $t = 0$. Choosing one common anchor per cloud makes the same-cloud travel cost zero.

Proof. If $z \neq w$, follow a base geodesic of length $q \leq t$. Each endpoint port traversal costs either zero or L , and each of the $q - 1$ intermediate traversals costs L . Including the q external edges gives

$$q + (q + 1)L = \lambda q + L \leq \lambda t + L.$$

When $z = w$, two designated ports coincide or have intrinsic distance L . A degree-one cloud is a singleton and has zero endpoint cost. When $k = 0$, all clouds are singletons and the assertion is the ordinary distance bound. If the connected base consists of one isolated vertex, designate that vertex as its anchor and handle its capture directly. $\square$

The hypothesis in Lemma 2.4 concerns the cops' actual positions when they are dispatched. Unused sampled teams may wait at their initially chosen ports through earlier rounds. Cops moved for a previous guard or pursuit cannot be reused under this bound without paying for their redeployment.

Lemma 2.5 (Clouds visited by a walk). *A walk in H of $T \geq 1$ edges visits at most*

$$2 + \left\lfloor \frac{T - 1}{\lambda} \right\rfloor$$

distinct ancestry clouds. In particular, at most $j\lambda$ edges visit at most $j + 1$ clouds for every integer $j \geq 0$, and at most $L + 1$ edges visit at most two clouds.

Proof. Record the successive directed external crossings. If there are a external crossings and b consecutive pairs that are not opposite, the walk visits at most $b + 2$ clouds. The first crossing visits at most one new cloud; a crossing immediately reversing its predecessor visits none, and every other crossing visits at most one. Each of the b intervals joins distinct ports within a single cloud and costs at least L , so $T \geq a + bL \geq (b + 1) + bL = b\lambda + 1$. Thus the number of visited clouds is at most $2 + \lfloor (T - 1)/\lambda \rfloor$. With no external crossing it is one. For $T \leq j\lambda$, this gives $j + 1$; the zero-length case is immediate. $\square$

Lemma 2.6 (Isometric lifts of base geodesics). *Let $v_0, \dots, v_q$ be a base geodesic with $q \geq 1$. Start at the designated port of F_{v_0} toward v_1 , end at the designated port of F_{v_q} toward v_{q-1} , and join the appropriate ports in each intermediate cloud by an intrinsic geodesic. The resulting path is ambient-isometric and has length $\lambda(q - 1) + 1$.*

Proof. Its length is $q + (q - 1)L = \lambda(q - 1) + 1$, which equals the lower bound in (1) for its endpoints. Every subpath of a shortest path is shortest. The endpoint ports are prescribed by the first and last edges; arbitrary designated endpoints can cost more. $\square$

2.1. The lower cop-number comparison.

Theorem 2.7 (HMGHM lower transfer, with explicit turns). *For every tower in Definition 2.1, including an arbitrary fixed number of rounds,*

$$c(G) \leq c(H).$$

No growth, stopping, or accessibility assumption is needed.

Proof. The case $k = 0$ is immediate. Fix $h < c(G)$ and a winning robber strategy against h cops on G . For any initial placement of h cops on H , use their ancestry vertices as the initial base cop positions. The base strategy chooses a vertex v at distance at least two from each of them: otherwise the first base cop move could capture it. Place the actual robber at any designated port of F_v . The metric lower bound puts every actual cop at distance at least $\lambda + 1$ from F_v . The first actual cop move is therefore safe and projects to a legal first base cop move. The singleton base, for which the conclusion is trivial, may be excluded when choosing ports.

We now take checkpoints immediately after cop moves. At a checkpoint the actual robber is at a designated port $p \in F_v$, and the cop shadows and v form the corresponding uncaptured base position on the robber's turn. Let $w \in N_G[v]$ be the base strategy's response. Its distance from every current cop shadow is at least two, since the base strategy must survive the next base cop move. Write z for any actual cop's current position. Then

$$(3) \quad \text{dist}_H(z, F_w) \geq \lambda + 1.$$

If $w = v$, let the actual robber wait for λ moves; the displayed bound makes all those waits safe. Otherwise let $a \in F_v$ be the designated port toward w and let $b \in F_w$ be its external neighbor. If $p \neq a$, follow a port path of length L from p to a ; if $p = a$, wait L times. Then cross ab . Denote the resulting positions by $r_1, \dots, r_\lambda$, with $r_\lambda = b$. The move order in this block is

$$R_1, C_1, R_2, C_2, \dots, R_\lambda, C_\lambda.$$

In particular, the robber's j th move precedes the j th cop reply. Since a has a neighbor in F_w , (3) implies $\text{dist}_H(z, a) \geq \lambda$. For $1 \leq j \leq L$, the remaining path from r_j to a has length at most $L - j$, and hence

$$\text{dist}_H(z, r_j) \geq \lambda - (L - j) = j + 1.$$

Thus neither the $j - 1$ preceding cop moves nor the j th reply can capture r_j . The last position b is safe through the λ th reply by (3). This checks every intermediate position against arbitrary cop walks, including those through shared hubs.

During the block each cop takes λ moves. By Lemma 2.5, its endpoint shadows are equal or adjacent, so these endpoints give a legal next base cop move. The actual robber is now at a designated port over w , safely after that move. This restores the checkpoint invariant and permits indefinite iteration. Therefore h cops cannot win on H . $\square$

The three-move argument in [1, Theorem 3(a)] is read with these robber-turn checkpoints. An interception starting with a cop outside the departure cloud and giving it an additional move before R_1 changes the phase. If that extra move enters the departure cloud, its current shadow already coincides with the base robber, contradicting the checkpoint invariant. Shared transit hubs do not invalidate the comparison; the safety estimate above does not require private tunnels.

3. RETRACTIONS AND EXTERIOR TREE SHADOWS

A *lazy graph retraction* from a graph J onto an induced subgraph R is a map $\rho : V(J) \rightarrow V(R)$ fixing R such that the images of the endpoints of each edge are equal or adjacent. Thus every robber walk, including waiting and reversal, has a legal projected walk in R .

Lemma 3.1 (Exterior port separation). *Let x_i, x_j be distinct designated ports of F_v , and let a_i, a_j be their respective neighbors outside F_v . Then*

$$\text{dist}_{H-F_v}(a_i, a_j) \geq 2L + 1,$$

where the distance is infinite when the two vertices lie in different components.

Proof. The exterior neighbors lie in different ancestry clouds because the base is simple. Any path between them avoiding F_v uses at least one external edge. In its first ancestry cloud, it starts at the designated port toward v and exits through a different port, at intrinsic cost at least L . The same holds in the last cloud. The external edge or edges contribute at least one more. Additional intermediate segments only increase the length. For $k = 0$ the asserted lower bound is one. $\square$

Theorem 3.2 (Retractions onto ancestry clouds). *Every ancestry cloud F_v is an induced retract of H . The same holds for descendant clouds at intermediate tower levels. Restricting such a retraction to an enclosing ancestry cloud is again a retraction.*

Proof. Assume first that $k \geq 1$ and v has at least two incident base edges. Fix a designated anchor port x_0 . For each other port x_i choose an intrinsic geodesic

$$p_i(0), p_i(1), \dots, p_i(L), \quad p_i(0) = x_i, \quad p_i(L) = x_0,$$

and let p_0 be the constant path at x_0 . Write a_i for the exterior neighbor of x_i . Fix F_v pointwise. For $z \notin F_v$, set

$$(4) \quad \rho_v(z) = \begin{cases} p_i(j), & j = \text{dist}_{H-F_v}(z, a_i) < L \text{ for some } i, \\ x_0, & \text{otherwise.} \end{cases}$$

The index in the first case is unique by Lemma 3.1: two such sources would have distance at most $2L - 2$.

No edge joins two different regions on which the first case applies, since it would give source distance at most $2L - 1$. Within one region, distance from its source changes by at most one along an edge, so its images are equal or consecutive along p_i . An edge leaving such a region has its inner endpoint at distance $L - 1$ from the source; the other endpoint maps to $x_0 = p_i(L)$. Finally an edge $x_i a_i$ crossing the boundary of F_v has both endpoints mapped to x_i . Therefore ρ_v is a lazy graph homomorphism fixing F_v .

For a degree-one base vertex, map the exterior constantly to its unique port. If $k = 0$, each cloud is a singleton and its constant map is a retraction. A connected singleton base is immediate. The construction applies with any intermediate graph as the base of its remaining tower. If F' is a descendant cloud contained in F_v , the restriction to $H[F_v]$ of a retraction $H \rightarrow H[F']$ still fixes F' and preserves edges, proving the final assertion. $\square$

Lemma 3.3 (Exterior tree shadows). *For any designated anchor $x_0 \in F_v$, the retraction in Theorem 3.2 can be chosen with a tree $T_v \subseteq H[F_v]$, rooted at x_0 and of height at most L , such that its restriction to $H - F_v$ is a lazy graph homomorphism into T_v . A cop initially at x_0 aligns with this retracted robber within L cops moves whenever the actual robber stays outside F_v throughout that setup interval. Once aligned, the cop can maintain a guard of F_v against all subsequent robber walks.*

Proof. Choose a breadth-first spanning tree of $H[F_v]$ rooted at x_0 and let T_v be the union of its root paths to the designated ports. By Lemma 2.2, those paths have length L , except for the constant path to x_0 . Use their reversals as the paths p_i in (4). The proof of that formula shows more than containment of the exterior image in T_v : the images of every exterior edge are equal or consecutive along a tree path. Thus an exterior robber has a legal lazy shadow in this tree.

Starting at its root, a cop follows the unique tree path toward the current shadow. Until alignment occurs, each cop move goes one level deeper. The shadow cannot leave the rooted subtree below

the cop without first coinciding with it. Since the height is at most L , alignment occurs within L cops moves. If it occurs immediately after a robber move, the next cops move can wait to establish the same alignment after a cops move; this does not exceed the bound, since a nonaligned move to depth L is impossible.

After alignment, on every subsequent cops move follow the robber's current image under the full retraction ρ_v . This move is legal even when the robber has entered F_v . In that case its image is its actual position, so the cop captures it on the next cops move. The setup hypothesis that the robber remains outside is needed only for the bounded-time tree chase, not for maintaining an established guard. Singleton clouds have a height-zero tree and require no setup. $\square$

3.1. Guarding two adjacent clouds.

Theorem 3.4 (Exact adjacent-cloud criterion). *Suppose $k \geq 1$ and $uv \in E(G)$. The induced union $R = H[F_u \cup F_v]$ is a lazy retract of H if and only if uv lies in no triangle of G . In that case its exterior shadow can be chosen in a tree of height at most $L + 1$, rooted at either endpoint of the unique edge joining the clouds. One cop at that root establishes a guard within $L + 1$ cop moves if the robber stays outside R throughout setup.*

Proof. Write ab for the joining edge, with $a \in F_u$ and $b \in F_v$. First assume that uv lies in no triangle. Root breadth-first trees of the two clouds at a, b . For each boundary edge $p_i z_i$ of R , let P_i be the tree path from p_i to a if $p_i \in F_u$, or the tree path from p_i to b followed by ba otherwise. Its length h_i is L or $L + 1$, respectively. These paths lie in one actual tree T rooted at a , of height at most $L + 1$.

The exterior sources z_i lie in distinct ancestry clouds: a shared neighbor of u and v would complete a triangle. As in Lemma 3.1, any path outside R between distinct sources pays two endpoint traversals of length L and an external edge. Hence their exterior distance is at least $2L + 1$. Fix R pointwise and define outside it

$$\rho(z) = \begin{cases} P_i(j), & j = \text{dist}_{H-R}(z, z_i) < h_i \text{ for some } i, \\ a, & \text{otherwise.} \end{cases}$$

The index is unique since two candidate distances would sum to at most $2L$. If an exterior edge joins different source regions, its endpoint distances j, j' satisfy $2L + 1 \leq j + 1 + j' \leq 2L + 1$. Thus $j = j' = L$, both sources attach to F_v , and both images are b . Within a source region images follow its tree path; an edge to the default region maps to the last tree edge toward a . Every boundary edge $p_i z_i$ maps to the single vertex p_i . This proves the retraction property and the exterior homomorphism into T . The root chase of Lemma 3.3 gives the setup bound.

Conversely, suppose uvw is a triangle, and choose the ports $p \in F_u, q \in F_v$ toward w . Their intrinsic distance in R is $2L + 1$, since every such path uses ab . The route through F_w has length $L + 2$. For $k \geq 1, L \geq 2$ makes this a strict shortcut. A retract is ambient-isometric, so no retraction onto R exists. $\square$

For $k = 0$ any two-cloud union is an edge and is a lazy retract, regardless of triangles. The exterior condition in the setup bound is essential: a retract need not be cop-win internally. Nor does girth five suffice for a union over a closed base neighborhood. In C_5 , that neighborhood is a three-vertex path. Between its two ports toward the complement, the three-cloud union has intrinsic distance $3L + 2$, whereas the exterior route has length $2L + 3$. For $k \geq 1$ this excludes a retraction. Attaching a degree-four tree at a base vertex outside the chosen neighborhood preserves this obstruction and forces a nontrivial stopped tower. These facts do not change the two-exception cleanup used below.

4. INTERNAL PURSUIT AND REUSABLE SETUP

We distinguish the cops needed to establish a guard from those retained after establishment. Retraction bounds for cop number are standard; the retract-cover bound is recalled, with attribution to earlier work, in [2]. The following elementary game argument gives the specific accounting used here.

Lemma 4.1 (Reusable retract setup). *Let R be a nonempty connected induced retract of a finite connected graph J . A team of $c(R)$ cops initially anywhere in J can, after finite setup, capture the robber or leave one cop maintaining a guard of R ; the other setup cops may then be reused. In particular, for nonempty connected induced retracts $R_1, \dots, R_s$, put $K = \max_i c(R_i)$ and let C range over the components of $J - \bigcup_i V(R_i)$. Then*

$$(5) \quad c(J) \leq s + \max \left\{ K - 1, \max_C c(C) \right\},$$

with $\max_C c(C) = 0$ when the complement is empty.

Proof. Deploy the setup team to a winning initial placement in R . This takes finite time because J is connected. The robber may move arbitrarily during deployment. Thereafter pursue its image under a retraction $\rho : J \rightarrow R$ using a winning strategy in R . That image is a legal robber walk. After finite time some cop reaches $\rho(r)$ immediately after a cops move, where r is the actual robber position. If $r \in R$ this is capture. Otherwise retain that cop and let it follow $\rho(r)$ on each later cops move. The retraction property makes every such move legal, and any actual entry into R is captured on the next cops move.

For (5), establish the retracts one at a time and retain one guard after each successful setup. Before the final setup there are at least K uncommitted cops. After all s guards are established, at least $\max_C c(C)$ cops remain for the component containing the robber. They may deploy there from arbitrary positions while all guards remain active. A surviving robber cannot change components. If the complement is empty, the final setup or the next cops move captures it. No uniform bound on the setup or cleanup time is asserted. $\square$

Lemma 4.2 (Independent twin blowups). *Let Q be connected with at least two vertices. Replace each vertex by a nonempty independent set, joining two such sets completely exactly when the corresponding vertices of Q are adjacent. The resulting graph $B(Q)$ satisfies*

$$c(B(Q)) \leq \max\{c(Q), 2\}.$$

Proof. Use a winning strategy on Q , representing cop positions by vertices in the corresponding twin classes. Every simulated move is legal because adjacent classes are joined completely. If virtual capture occurs without actual capture, the robber and the capturing cop occupy different vertices of the same independent twin class. Keep that cop fixed. Any move of the robber out of the class can be captured by it on the next cops move. A second cop travels to a neighbor of the robber and captures if the robber keeps waiting. Such a neighbor exists since Q is connected and nontrivial. Thus a virtual capture can be completed using an existing second cop when $c(Q) \geq 2$, and one extra cop otherwise. $\square$

Lemma 4.3 (One-round internal pursuit). *Every one-round HMGHM cloud has cop number at most two. Singleton and star clouds have cop number one.*

Proof. The graph $A_r(m)$ is the independent twin blowup of the graph obtained by subdividing every edge of K_m once: the original vertex i becomes X_i , and each subdivision vertex remains the singleton y_{ij} . Here is a direct two-cop strategy on that subdivided clique for $m \geq 2$. Start at distinct original vertices a, b . A robber on a subdivision vertex incident with either is captured immediately. For a robber on the subdivision of rs with $r, s \notin \{a, b\}$, move the cops to the subdivisions of ar, bs . Moving to r or s then loses on the next cop move; if the robber waits, advance to r, s and capture

next. For a robber starting at an original vertex $r \notin \{a, b\}$, move the first cop to the subdivision of ar and keep the second at b . Waiting or entering the subdivision of ar or br loses immediately or on the next cop move. If the robber moves to the subdivision of rs for another s , move the cops to r and the subdivision of bs ; waiting, moving to r , and moving to s all lose by the following cop move. Lemma 4.2 preserves this two-cop upper bound. The singleton and $m = 2$ HMGHM cases are trees. $\square$

Proposition 4.4 (Internal cloud budget). *For a tower of $k \geq 1$ rounds,*

$$(6) \quad \max_{v \in V(G)} c(H[F_v]) \leq 2 + \sum_{j=0}^{k-2} D_j.$$

If the stopping rule is used and $D \geq 4$, then

$$(7) \quad K := \max_{v \in V(G)} c(H[F_v]) \leq 2D.$$

If the base is already subcubic, $k = 0$ and $K = 1$.

Proof. The case $k = 1$ is Lemma 4.3. For more rounds, consider the first gadget of an ancestor v of degree $r \geq 2$. Its r port vertices and its pair vertices each have a descendant cloud in the final graph. By Theorem 3.2, every child cloud is a retract of the parent cloud. Suppose their internal cop numbers are at most K' . Guard the r port-descendant clouds using Lemma 4.1. Removing them leaves each pair-descendant cloud as a separate component, since pair vertices are independent in $A_r(m)$. Therefore the parent cloud needs at most $r + K'$ cops. An ancestor of degree zero or one has just one child and needs no separator charge. Induction on the remaining number of rounds gives (6).

For the stopping bound, the gadget degrees give

$$D_{j+1} \leq f(D_j), \quad f(4) = 3, \quad f(x) = 2 \left\lceil \sqrt{x/2} \right\rceil \quad (x \geq 5).$$

Indeed a port has total degree $m(r)$ after its external edge is attached, and a pair vertex has degree at most $2\lceil r/m(r) \rceil$; these are bounded by the displayed expression when the current maximum degree is $x \geq 5$. The special degree-four gadget has maximum degree three. This is the one-round bound of [1, Theorem 3].

Define $T(x) = 2$ for integers $x \leq 4$ and $T(x) = x + T(f(x))$ for $x \geq 5$. The sequence is well defined because $f(x) < x$ for $x \geq 5$, and T is nondecreasing by induction. The actual degree sequence in (6) is bounded by this envelope. For $5 \leq x \leq 8$ one has $T(x) = x + 2$, and for $9 \leq x \leq 15$ one has $T(x) = x + 8$, so $T(x) \leq 2x$ in these cases. For $x \geq 16$,

$$f(x) \leq \sqrt{2x} + 2 \leq x/2.$$

Strong induction consequently gives $T(x) = x + T(f(x)) \leq x + 2f(x) \leq 2x$. This proves (7). The $k = 0$ statement follows from the singleton fibers. $\square$

The geometry and retractions above hold for any specified number of rounds. The simplification to $2D$ uses the stopping rule: gratuitous rounds on an already subcubic graph are not covered by that degree-sum estimate.

5. A LATE HANDOFF CANNOT BE REPAIRED BY FIVE LOCAL MOVES

The timing gains from port placement are necessary hypotheses of the prepared strategy, rather than consequences of adjacency alone. The next example quantifies that distinction while retaining the old guard.

Proposition 5.1 (A late-trigger obstruction). *There is a 104-vertex subcubic stopping tower with $\lambda = 9$, adjacent ancestry clouds F_A, F_B , a legal retained guard of F_A , and an allowed starting side for an arbitrarily large transition team, with the following properties. No transition cop can align with the robber's image under any retraction onto F_B within five moves. Moreover, after a deadline at the end of round five, a legal entry into F_B in round six remains uncaught on the next cops move, in round seven.*

Proof. Take the nine-vertex tree formed from the path $A - B - C$, a second leaf C' adjacent to B , an edge AD , and four leaves adjacent to D . The degree of D is five. The all-vertex stopping tower has maximum degrees 5, 4, 3, so it performs exactly two rounds. Its A cloud has $3 \cdot 3 = 9$ vertices. In the first B gadget, three ports have total degree three and three pair vertices have degree two, giving a final B cloud of order $3 \cdot 6 + 3 \cdot 3 = 27$. The first D gadget has five ports of total degree four, three pair vertices of degree three, and three pair vertices of degree two. Its final cloud therefore has order $5 \cdot 7 + 3 \cdot 6 + 3 \cdot 3 = 62$. The six original leaves remain singletons, so the total order is $9 + 27 + 62 + 6 = 104$.

Write b_-, b_+ for the ports of F_B toward A, C , respectively, and a for the port of F_A adjacent to b_- . Let c, c' denote the singleton descendants of the two leaves at B , and set

$$W = V(H) \setminus (F_B \cup \{c, c'\}).$$

The edge ab_- is the only edge from W to its complement. Each exterior component of F_B attaches at just one port, so leaving the cloud cannot shorten a path between its ports. Lemma 2.2 gives

$$\text{dist}_H(W, b_-) = 1, \quad \text{dist}_H(W, b_+) = 9, \quad \text{dist}_H(W, y) \geq 8 \quad (y \in N_H[b_+] \cap F_B).$$

Equality in the last minimum is attained by the neighbor of b_+ preceding it on an intrinsic geodesic from b_- .

Place all transition cops anywhere in W , and retain an additional old guard at a . Choose the robber initially at c . This placement is legal with the cops chosen before the robber. The old guard can remain stationary: every surviving walk from the robber's component into F_A would enter at the occupied vertex a .

For any retraction $\rho : H \rightarrow F_B$, the edge cb_+ and the identity $\rho(b_+) = b_+$ imply

$$\rho(c) \in N_H[b_+] \cap F_B.$$

Every such image is at distance at least eight from every allowed transition-cop starting position. Therefore five moves cannot establish alignment while the robber waits at c , regardless of the team size or the chosen retraction.

Number each round as a cops move followed by a robber move. Let the robber wait at c through round five, enter b_+ on its move in round six, and wait there in round seven. No transition cop reaches b_+ in fewer than nine moves. The entry is thus safe both when it occurs and on the next cops move. The old guard remains valid throughout. $\square$

Remark 5.2 (Scope and larger scales). This is a starting-location and deadline obstruction, not a cop-number lower bound. The robber in the example can eventually be captured. Earlier deployment, cops prepositioned outside W , and prepared source anchors evade the specified hypothesis. In particular, the proposition does not contradict a schedule that pays for anchor travel and the L -move exterior-tree setup before the deadline.

For any prescribed $k \geq 1$, replace the degree-five star center in the construction by one of degree $2^{k-1}+1$. Its descendants realize the maximum-degree sequence

$$2^{2^{k-1}+1}, 2^{2^{k-2}+1}, \dots, 4, 3,$$

so the stopping rule performs exactly k rounds. For each degree $r = 2^{2^j+1}$ above four, $m = \sqrt{2r}$ is an integer, all port classes have size $m/2$, and both port and pair degrees are exactly m ; this verifies the sequence. The target's ports have intrinsic distance $\lambda - 1$, so retraction alignment from the

old side needs at least $\lambda - 1$ moves. After a deadline at the end of round T , entry in robber round $T + 1$ remains safe through the following cops move whenever $\lambda > T + 2$. These are exact distance lower bounds. The upper bound $\text{diam}(F_v) \leq 2(\lambda - 1)$ alone would not establish any such delay.

6. MULTISTAGE BASES AND THE TRANSFER THEOREM

For a vertex set A in a graph G , write $B_G(A, r) = \{x : \text{dist}_G(x, A) \leq r\}$ and $S_G(A, r) = \{x : \text{dist}_G(x, A) = r\}$. For a singleton source, we omit its braces. An empty source has empty neighborhoods. A set $U \subseteq V(G)$ is (t, c_1, c_2) -*accessible* if there are pairwise disjoint sets $W(w) \subseteq B_G(w, t)$, one for each $w \in U$, satisfying

$$|W(w)| \geq c_1 \min\{d^t, c_2 N/|U|\}.$$

Here $N = |V(G)|$ and d is the branching parameter specified below.

Definition 6.1 (Five-layer multistage base). Fix an integer $d \geq 2$, a constant $a_0 \geq 0$, and positive constants $\delta, J, a_1, \dots, a_5$. A connected simple graph G on N vertices is a *five-layer multistage base* if there is a set $X \subseteq V(G)$, with $|X| \leq a_0 \sqrt{N}$, such that:

(H1) For every $v \notin X$, all integers $r, r' \geq 1$ with $d^r, d^{r'} < N^{1/2+\delta}$, and every $V' \subseteq B_G(v, r) \setminus X$ of order m satisfying $md^{r'} \leq N/(\log N)^J$,

$$a_1 md^{r'} \leq |S_G(V', r')| \leq a_2 md^{r'}.$$

(H2) For all integers r, r' such that

$$N^{1/4-\delta} < (d+1)d^r < N^{1/4+\delta}, \quad N^{1/4-\delta} < (d+1)d^{r'} < N^{1/4+\delta},$$

every $v \notin X$, and every $A \subseteq S_G(v, r) \setminus X$ with $|A| > N^{1/4-\delta}$, put $U = \bigcup_{a \in A} S_G(a, r')$. If $d^{r+r'} < a_3 N/|U|$, there is a set $Q \subseteq V(G)$ such that

$$|S_G(a, r') \cap Q| < N^{1/4-2\delta} \quad (a \in A),$$

and $U \setminus Q$ is $(r + r' - 4, a_4, a_5)$ -accessible.

(H3) $G - X$ is contained in one component of G .

These are exactly the $s = 5$ hypotheses of the occupation transfer [5, Section 9]. They are the hypotheses of Prałat–Wormald [3, Theorem 4.1], with accessibility required five layers earlier. We keep (H3) to make the correspondence explicit; it is automatic for a connected base. The target spheres in (H1) are *not* trimmed by X , although the sources are. All parameters in the definition are fixed as N tends to infinity.

Theorem 6.2 (Timed transfer through actual clouds). *For the fixed parameters of Definition 6.1, there are constants C, C_*, N_0 with the following property. Let G be a five-layer multistage base of order $N \geq N_0$, and let H be its actual all-vertex HMGHM tower, stopped when its maximum degree first becomes at most three. Write $D = \Delta(G)$, let F_v be the original ancestry clouds, and put $K = \max_v c(H[F_v])$. Then*

$$(8) \quad c(H) \leq (a_0 + 2C + 3)\sqrt{N} + K + 1 \leq (a_0 + 2C + 3)\sqrt{N} + 2D + 1,$$

and, in fact,

$$(9) \quad \boxed{c(H) \leq C_* \sqrt{N}.}$$

The constants depend only on the listed multistage parameters, and not on the cloud order, tower scale, or D . All cops are placed before the robber chooses her initial vertex. The strategy captures every lazy robber walk after finitely many moves.

The theorem retains the full remaining-ball Hall requirement of the occupation argument. It does not assert a transfer for an arbitrary coarse projection or for a base satisfying only a lower expansion bound. Nor do we retain the occupation theorem's explicit capture-time bound. The constant is not asserted uniform when the branching parameter d grows. The maximum degree D may grow within (H1); the distinction between these two parameters is used below.

7. SYNCHRONIZATION AND LOCAL CLEANUP

Throughout this section the tower has at least one round. Set $\lambda = 3^k$ and $L = \lambda - 1$, and choose a designated original-edge port a_v as a fixed anchor in each cloud F_v . Fresh sampled cops will start at these anchors and remain there until their team's dispatch.

Lemma 7.1 (Exit lower bound). *If a robber starts in F_u and first reaches a cloud over $S_G(u, r)$ after m moves, then*

$$m \geq \lambda(r - 1) + 1.$$

Before that move her base shadow lies in $B_G(u, r - 1)$.

Proof. The base shadow moves along an edge or stays. A path joining clouds whose base distance is r uses at least r inter-cloud edges. After erasing immediate reversals from a shortest such path, its intermediate cloud traversals join distinct designated ports and cost at least L each. Thus its length is at least $r + (r - 1)L = \lambda(r - 1) + 1$. The assertion about earlier shadows follows from their lazy motion. $\square$

Lemma 7.2 (Synchronized two-round deadline). *Suppose a round of radius r is followed by a round of radius $r' \geq 2$, and $r + r' \geq 4$. At the start of the first round, a cop at a fixed source anchor is assigned to F_w , with source base distance at most $t = r + r' - 4$ from w . She travels to a_w and waits there if she arrives before the next base root u is observed. Once both arrival and root observation have occurred, activate this cop if $w \in S_G(u, r')$, using the exterior-tree pursuit. Every assigned cloud of this actual sphere can be guarded before the first sphere entry, simultaneously for all assigned targets and every robber walk.*

Proof. Number cop and robber moves after dispatch so that cop move j immediately precedes robber move j . Lemma 2.4 gives anchor arrival by

$$B = \lambda t + L = \lambda(r + r' - 3) - 1.$$

Let q be the robber move ending the first round. Let m be the first subsequent sphere-entry move. Lemma 7.1 gives

$$q \geq \lambda(r - 1) + 1, \quad m \geq q + \lambda(r' - 1) + 1 \geq \lambda(r + r' - 2) + 2.$$

Use the common activation time $\sigma = \max\{B, q\}$, after robber move σ , and use cop moves $\sigma + 1, \dots, \sigma + L$ for setup. The two relevant bounds are

$$(10) \quad B + L = \lambda(r + r' - 2) - 2 \leq m - 4, \quad q + L \leq m - 2.$$

The second uses $r' \geq 2$. Hence $\sigma + L \leq m - 2$. Every target cloud remains unentered throughout its setup, so Lemma 3.3 establishes every assigned guard. The guards then track their retracted shadows permanently. $\square$

The destinations and source anchors are chosen before the unknown root u is revealed; only activation depends on u . Earlier visits to a target, before the first round ends, cause no problem because the cop waits at the anchor. The saving in Lemma 2.4 is essential to this calculation: arbitrary fiber representatives have only the generic bound $\lambda t + 2L$, which would consume a sixth layer for the same setup uniformly in λ .

Lemma 7.3 (Cleanup with a reusable reserve). *Let $R \subseteq V(G)$ be a specified set of base vertices. A team of $|R|$ anchor cops and $K + 1$ reserve cops, initially anywhere in the connected ambient graph, can force capture or an exit from $\bigcup_{v \in R} F_v$ after finitely many moves.*

Proof. Assign one cop to each target anchor, wait until all have arrived, and start their exterior-tree pursuits simultaneously. Record the clouds visited by the robber during the next L robber moves, including the starting cloud. By Lemma 2.5, at most two distinct original clouds are visited. If the robber exits the prescribed region, the exit alternative has occurred. Otherwise all unvisited target clouds acquire guards within L cop moves by Lemma 3.3. Every visited target is marked exceptional, whether or not its assigned cop has already aligned. If it has not aligned, that cop can remain idle, so no pursuit toward an image outside its tree is prescribed. Every successful guard remains active.

If one exceptional cloud remains, the reserve can deploy K cops to a winning initial configuration of its induced graph and play there. If two remain, they are adjacent and their induced union has exactly one joining edge, since the base is simple. Send one reserve cop to an endpoint of that edge. Once it is occupied, the surviving robber is in one observed side, where the other K reserve cops deploy and pursue. During this deployment, or during pursuit, leaving the exceptional union either enters a guarded cloud or exits the prescribed region. All travel is finite. The known cop bound on each induced cloud gives finite capture under continued confinement. No retraction onto the union of the exceptional clouds is needed. $\square$

The setup cops in this lemma can be reused after an actual round exit; the continuing-round guards belong to different teams and are retained. For finishing we will apply the lemma's reserve argument only after proving that the two exceptional clouds are interior to a fully guarded boundary. That will remove the exit alternative.

8. PROOF OF THE MULTISTAGE TRANSFER

We reproduce the radius, reservoir, and Hall estimates of Pralat–Wormald [3, Theorem 4.1 and Claim 4.2], with the five-layer reservoirs and remaining-ball target of [5, Section 9]. These are probabilistic estimates on the base. The deterministic replacement of occupation by timed guards is proved above; we do not infer it from the earlier cop-number conclusion.

8.1. Placement, radii, and finite rounds. Put $F = J + 2$ and $i_f = \lceil F \log \log N \rceil$. Independently for each $1 \leq i \leq i_f$, sample $Z_i \subseteq V(G)$ by including each vertex with probability c_i/N , where

$$c_i = C e^{-i} \sqrt{N} \quad (i < i_f), \quad c_{i_f} = \sqrt{N}.$$

Team i consists of one cop at a_z for each $z \in Z_i$. There are also $|X|$ exceptional-guard cops, $\lceil N^{1/3} \rceil$ cleanup anchor cops, and $K + 1$ reserve cops. All are placed before the original robber start. The last three groups may start anywhere.

First establish a guard on every X cloud using the known sequential retract setup of Lemma 4.1. The $|X|$ assigned cops and $K + 1$ reserve suffice: each established cloud retains one cop, and at least K mobile cops remain available for the next setup. At the end there are still $K + 1$ mobile reserve cops. Sampled teams stay at their chosen source anchors throughout this prephase and all earlier rounds. The prephase either captures the robber or finishes with her outside all X clouds. Let v_1 be her actual base shadow at that time.

By Chernoff's inequality, with probability $1 - o(1)$ every sampled team has at most $2c_i$ members. On this event the total number of cops is at most

$$(11) \quad |X| + 2\sqrt{N} \left(C \sum_{i \geq 1} e^{-i} + 1 \right) + \lceil N^{1/3} \rceil + K + 1 \leq (a_0 + 2C + 3)\sqrt{N} + K + 1$$

for sufficiently large N .

Choose a sufficiently small fixed $\varepsilon_0 > 0$. Define

$$r_1 = \left\lfloor \frac{1}{4} \log_d(\varepsilon_0 N) \right\rfloor, \quad r_{i+1} = \left\lfloor \log_d(\varepsilon_0 e^{2i} \sqrt{N}) \right\rfloor - r_i.$$

Thus the source radius condition is

$$(12) \quad \varepsilon_0 e^{2i} \sqrt{N} / d < d^{r_i + r_{i+1}} \leq \varepsilon_0 e^{2i} \sqrt{N}.$$

Each parity subsequence is nondecreasing, since $r_{i+2} - r_i$ is the difference of two successive increasing floor terms. Moreover $r_1, r_2 = \frac{1}{4} \log_d N + O(1)$, so, uniformly for $i \leq i_f + 1$,

$$(13) \quad d^{r_i} = \Omega(N^{1/4}), \quad d^{r_i} = O(e^{2i} N^{1/4}).$$

All implicit constants here may depend on the fixed parameters. In particular every needed radius is at least two, and successive pairs lie in the window of (H2), for sufficiently large N . The ceiling in i_f causes only the harmless bound $e^{2i_f} \leq e^2 (\log N)^{2F}$.

Round i starts at the actual base shadow $v_i \notin X$ and ends at its first subsequent visit to $S_G(v_i, r_i)$, with next root v_{i+1} . If the robber enters an X cloud she is captured on the next cop move, so we restrict attention to surviving walks avoiding them. The exit lower bound is Lemma 7.1.

To prevent an infinite round, apply Lemma 7.3 to $R = B_G(v_i, r_i) \setminus X$. By (H1) with a singleton source and (13),

$$|R| \leq 1 + 2a_2 d^{r_i} = O(N^{1/4} (\log N)^{2F}) \leq N^{1/3}.$$

If the round has not already ended, the robber is inside $B_G(v_i, r_i - 1) \setminus X$. Cleanup forces capture or an exit from R , which must first end the round. At any actual round exit the cleanup attempt is aborted and its anchor cops and reserve are reused. There are only finitely many rounds in the argument; no uniform cleanup time is required. Waiting and reversal are included in these alternatives.

8.2. Reservoir assignments and vulnerability. Let $S_0 = S_G(v_1, r_1) \setminus X$. Team i is dispatched at the start of round i , unless the finishing trigger below has already occurred. For $i \geq 1$, define

$$S_i = \{w \in S_G(v_{i+1}, r_{i+1}) \setminus X : \text{team } i \text{ has no cop assigned to } F_w\}.$$

A root v_i is *vulnerable* if

$$(14) \quad |S_{i-1}| \leq e^{-5(i-1)} |S_G(v_i, r_i)|.$$

The initial root is vulnerable. Lemma 7.2 ensures that a surviving end of round i belongs to S_{i-1} .

Fix a vulnerable root and a possible $A = S_{i-1}$ before exposing team i , and suppose

$$(15) \quad |A| > e^{-5(i_f-1)} |S_G(v_i, r_i)|.$$

Put $r = r_i$, $r' = r_{i+1}$, and $U = \bigcup_{a \in A} S_G(a, r')$. By (H1), (13), and (15), $|A| = \Omega(N^{1/4} (\log N)^{-5F}) > N^{1/4-\delta}$. Also

$$(16) \quad |U| \leq e^{-5(i-1)} a_2^2 d^{r+r'}, \quad |U| d^{r+r'} \leq a_2^2 \varepsilon_0^2 e^{5-i} N < a_3 N$$

when ε_0 is small enough. Hypothesis (H2) supplies its exceptional set Q and disjoint reservoirs of radius $t = r + r' - 4$. Their sizes satisfy

$$(17) \quad |W(w)| \geq \kappa e^{2i} \sqrt{N} \quad (w \in U \setminus Q), \quad \kappa = a_4 \min \left\{ \varepsilon_0 d^{-5}, \frac{a_5}{a_2^2 \varepsilon_0 e^5} \right\} > 0.$$

Indeed the first term follows from (12); for the second, (16) gives $N/|U| \geq e^{3i-5} \sqrt{N} / (a_2^2 \varepsilon_0) \geq e^{2i} \sqrt{N} / (a_2^2 \varepsilon_0 e^5)$.

Fix an ordering of base vertices and choose every reservoir family deterministically from the finitely many admissible choices. Assign to w the cop at the first member of $W(w) \cap Z_i$, when this

intersection is nonempty. Disjointness makes the assignment injective. The empty-hit events E_w are mutually independent and satisfy

$$(18) \quad \Pr(E_w) \leq \exp(-C\kappa e^i) \leq \frac{1}{2}e^{-5i}$$

for all $i \geq 1$ once $C \geq \max\{1, 6/(e\kappa)\}$. Targets in X need not be activated, because their prephase guards remain; including them in this assignment and count only wastes available cops.

For fixed $u \in A$, set $H_u = S_G(u, r') \setminus Q$. By (H1), (H2), and (13), $|H_u| \geq \frac{1}{2}a_1 d^{r'} = \Omega(N^{1/4})$ for large N . The number of empty-hit targets in H_u is stochastically dominated by a binomial variable of mean $\mu = \frac{1}{2}e^{-5i}|H_u| = \Omega(N^{1/4}(\log N)^{-5F})$. Chernoff's inequality gives at most $\frac{2}{3}e^{-5i}|H_u|$ such targets except with probability $\exp(-\mu/27)$. Moreover,

$$|Q \cap S_G(u, r')| < N^{1/4-2\delta} \leq \frac{1}{3}e^{-5i}|S_G(u, r')|$$

for large N . A union bound over $u \in A$ therefore shows that, with failure probability $O(N^{-\log N})$, every possible actual next sphere has at most an e^{-5i} fraction of unassigned targets outside X . No restriction on the size of the fixed positive δ is needed for this estimate.

Send the assigned cops to their fixed anchors immediately at the start of round i . When its actual next root $u = v_{i+1} \in A$ is observed, Lemma 7.2 activates the assigned cops on $S_G(u, r') \setminus X$ and establishes their guards before any entry to that sphere. Therefore the next root is vulnerable, or capture occurs. The four-move and two-move margins in (10) hold for every possible duration of the first round.

We must make this assertion simultaneous despite the robber seeing all sampled teams. Expose teams successively for the analysis. Conditional on the earlier samples, deterministic assignment choices specify a unique set S_{i-1} for each base exit history $(v_1, \dots, v_i)$; there are at most N^i histories. A union bound over those histories and $i < i_f = O(\log \log N)$ gives failure probability $o(1)$. The assignments depend only on these histories and samples. Their subsequent motions may depend on the entire robber walk, but the deterministic synchronization lemma succeeds uniformly over every such walk. It introduces no extra probability event. In particular v_1 may depend on all the samples through the preprocessing: the event already holds for every possible initial root. This is the same all-history exposure argument as [3, Claim 4.2].

8.3. The full remaining-ball matching and capture. Dispatch the still-unmoved final team i_f at the start of the first round i^* for which

$$|S_{i^*-1}| \leq e^{-5(i_f-1)}|S_G(v_{i^*}, r_{i^*})|.$$

The preceding argument ensures $i^* \leq i_f$, unless capture occurred earlier. Teams $i^*, \dots, i_f - 1$ are then unnecessary. Write $r = r_{i^*}$ and put

$$r'' = r_{i_f} + r_{i_f+1} - r, \quad \hat{r} = r_{i_f} + r_{i_f+1} + 1, \quad \hat{t} = \hat{r} - 5 = r + r'' - 4,$$

$$U_{\text{fin}} = \bigcup_{u \in S_{i^*-1}} B_G(u, r'') \setminus X.$$

Monotonicity of the parity subsequences implies $r'' \geq \min\{r_{i_f}, r_{i_f+1}\} \geq 2$ for large N . Also (12) and (13) give $d^{r''} = O(N^{1/4}(\log N)^{2F})$, so all singleton sphere bounds through radius r'' used next are admissible in (H1). The target consists of balls, not just spheres; arbitrary waiting requires the interior targets in the finishing argument. Summing the singleton sphere upper bounds in (H1) gives

$$(19) \quad |U_{\text{fin}}| \leq e^{-5(i_f-1)}a_2 d^r 3a_2 d^{r''} \leq 3a_2^2 e^5 \varepsilon_0 e^{-3i_f} \sqrt{N}.$$

For Hall's theorem, take any nonempty $V'' \subseteq U_{\text{fin}}$ of order m . Apply (H1) with center v_{i^*} , containing radius $\hat{r} - 1$, and neighborhood radius $\hat{t}$. Its source condition holds because

$$V'' \subseteq B_G(v_{i^*}, r + r'') \setminus X = B_G(v_{i^*}, \hat{r} - 1) \setminus X.$$

Both radius powers are below $N^{1/2+\delta}$ by (12). The size condition follows from (19):

$$md^{\hat{t}} \leq |U_{\text{fin}}|d^{\hat{r}} = O(Ne^{-i_f}) \leq N/(\log N)^J.$$

Consequently

$$(20) \quad |S_G(V'', \hat{t})| \geq a_1 md^{\hat{t}} \geq a_1 \varepsilon_0 d^{-5} m \sqrt{N} (\log N)^{2F}.$$

The final team samples each base vertex with probability $N^{-1/2}$. Its number of starting vertices in this sphere is binomial with mean at least $a_1 \varepsilon_0 d^{-5} m (\log N)^{2F}$, which exceeds $m (\log N)^2$ for large N . The probability of fewer than m cops is at most $\exp(-m (\log N)^2 / 8)$. Summing over the at most N^m source sets of each size gives failure probability $O(\exp(-(\log N)^2 / 9))$. Conditional on earlier samples, the possible finishing sets are indexed by $N^{O(i_f)}$ base exit histories, so the additional union bound still gives failure probability $o(1)$.

Thus, simultaneously for every feasible finishing history, Hall's theorem assigns a distinct final-team cop to every target $w \in U_{\text{fin}}$, from a starting base vertex within distance $\hat{t}$ of w . Choose such a matching deterministically. Each assigned cop starts at its original port anchor and arrives at the target anchor by

$$B = \lambda \hat{t} + L = \lambda(r + r'' - 3) - 1$$

cop moves after dispatch.

If round i^* never ends, cleanup captures the robber. Otherwise let q be its actual end, and let $u = v_{i^*+1} \in S_{i^*-1}$. At time $\sigma = \max\{B, q\}$, activate the anchor cops for the actual ball $B_G(u, r'') \setminus X$ simultaneously. The calculation in Lemma 7.2, with r'' in place of r' , finishes setup strictly before the first possible entry to $S_G(u, r'')$. At most two clouds are visited during setup, by Lemma 2.5; both are interior to the ball. Every other target is guarded, and every boundary cloud is guarded either by a successful target cop or by its retained X guard.

Keep these guards active and use the $K + 1$ reserve as in Lemma 7.3 to pursue within the one or two exceptional clouds. An exit from their union inside the ball enters a guarded cloud; an exit from the ball must cross its guarded boundary. Hence the exit alternative is unavailable to a surviving robber and capture follows. The reserve may begin this final deployment anywhere: the boundary guards remain in place during all travel. This proves the first inequality in (8) using (11). Proposition 4.4 gives its second inequality.

8.4. Absorbing the internal setup budget.

Lemma 8.1 (Degree control from the exact sphere bounds). *For sufficiently large N , a base in Definition 6.1 satisfies*

$$\Delta(G) \leq \max\{|X| - 1, 1 + a_2 d + a_2 d^2\}.$$

Proof. For $v \notin X$, apply (H1) to $V' = \{v\}$ with containing radius 1 and neighborhood radii 1 and 2. All fixed-radius admissibility conditions hold for large N , giving $|S_G(v, 1)| \leq a_2 d$ and $|S_G(v, 2)| \leq a_2 d^2$. If $x \in X$ has a neighbor $v \notin X$, then

$$N_G(x) \subseteq \{v\} \cup S_G(v, 1) \cup S_G(v, 2),$$

so $\deg_G(x) \leq 1 + a_2 d + a_2 d^2$. This uses the untrimmed target spheres in (H1). If x has no neighbor outside X , its degree is at most $|X| - 1$. Vertices outside X already have degree at most $a_2 d$. $\square$

Since $|X| \leq a_0 \sqrt{N}$, Lemma 8.1 absorbs the additive $2D + 1$ in (8), proving (9). If the tower has zero rounds, $H = G$ and its clouds are singletons. The same construction then uses $\lambda = 1$ and $L = 0$: an anchor cop occupies its entire target cloud, cleanup is ordinary vertex occupation, and the finishing targets are all occupied at activation. Thus the same bounds hold without a shadow setup. Equivalently, [3, Theorem 4.1] applies after replacing its accessibility constant a_4 by $a_4 d^{-5}$. This completes the proof of Theorem 6.2.

Corollary 8.2 (Fixed-degree random regular bases). *Fix an integer $d \geq 2$. With probability tending to one, a uniformly random $(d+1)$ -regular base on N vertices is connected and its actual stopping-rule HMGHM tower has cop number $O_d(\sqrt{N})$, with the cloud order absent from the transfer constant.*

Proof. The random-regular expansion and reservoir estimates of Prałat–Wormald [4, Theorem 3.1, Lemma 3.2, and Sublemma 3.4] give the base hypotheses, with the five-layer reading of their reservoir trees and repaired exceptional set established in [5, Section 9, early-accessibility proposition and random-regular corollary]. Apply Theorem 6.2. This use is for fixed degree and a fixed level shift. $\square$

9. THE REMAINING HALL-DEMAND OBSTRUCTION

Timed guards reduce the cost of protecting an assigned cloud. They do not supply the distinct-cop assignment demanded by the finishing step. The following elementary inequality isolates a limitation of full-region demand. It uses actual graph neighborhoods, without a thin-sphere hypothesis.

For nonempty $U \subseteq V(G)$ and an integer $t \geq 0$, define

$$\mathcal{R}_t(U) = \min_{\emptyset \neq D \subseteq U} \frac{|B_G(D, t)|}{|D|}, \quad V_t(U) = \max_{w \in U} |B_G(w, t)|.$$

If every target in U has an assigned disjoint reservoir of common integer size m inside its radius- t ball, then $m \leq \lfloor \mathcal{R}_t(U) \rfloor$. Conversely Hall's theorem, applied to m copies of each target, gives such reservoirs whenever $m \leq \mathcal{R}_t(U)$. This observation concerns a uniform reservoir certificate, rather than every possible sampled matching.

Proposition 9.1 (Capacity of a region containing a ball). *Suppose $B_G(a, u) \subseteq U$, where $u \geq 1$ is an integer, and put $M = |B_G(a, u)|$. Then*

$$\mathcal{R}_t(U) \leq 1 + V_t(U)(1 - M^{-1/u}) \leq 1 + V_t(U) \frac{\log M}{u}.$$

Proof. For $D \subseteq V(G)$ let $\partial_{\text{in}} D$ consist of the vertices of D with a neighbor outside D . Any path from D to a vertex outside D passes through this inner boundary, so

$$B_G(D, t) \subseteq D \cup B_G(\partial_{\text{in}} D, t).$$

If $D \subseteq U$, it follows that

$$|B_G(D, t)| \leq |D| + V_t(U) |\partial_{\text{in}} D|.$$

Write $b_j = |B_G(a, j)|$ for $0 \leq j \leq u$. Since $b_0 = 1$, $\prod_{j=1}^u b_{j-1}/b_j = M^{-1}$, and some j has $b_{j-1}/b_j \geq M^{-1/u}$. For $D = B_G(a, j) \subseteq U$, the inner boundary is contained in $S_G(a, j)$, whence

$$\frac{|\partial_{\text{in}} D|}{|D|} \leq 1 - \frac{b_{j-1}}{b_j} \leq 1 - M^{-1/u}.$$

The first bound follows by using this D in the capacity minimum. The second is $1 - e^{-x} \leq x$ with $x = (\log M)/u$. $\square$

Corollary 9.2 (An incompatible slow-growth certificate). *Fix $a_*, C_b > 0$. There is no sequence of such demand regions containing $B_G(a, u)$ with $q \geq 1$, $q \rightarrow 1$, $u \rightarrow \infty$, and $Z \rightarrow \infty$ that simultaneously satisfies*

$$|B_G(a, u)| \leq 1 + C_b q^{u+1}, \quad V_t(U) \leq 1 + C_b q Z, \quad \mathcal{R}_t(U) \geq a_* Z.$$

Proof. The ball upper bound implies $\log M/u \leq \log q + u^{-1} \log(1 + C_b q)$. By Theorem 9.1, the three inequalities require

$$a_* \leq Z^{-1} + (Z^{-1} + C_b q) \left(\log q + \frac{\log(1 + C_b q)}{u} \right).$$

The right side tends to zero. $\square$

In the exception-free proposed shrinking-base model, the finishing union contains a full remaining ball. The corollary therefore obstructs the uniform Hall demand even after optional sphere bounds at the remaining and finishing radii are removed. It does not assume any intermediate-radius growth law: the inner balls used in the proof are merely subsets of the same demand region, and the Hall radius stays t . With a trimmed region $U = (\bigcup_a B_G(a, u)) \setminus X$, the conclusion applies only to ambient balls actually contained in U . A path avoiding X alone does not supply such a ball. Neither statement lower-bounds the cop number of the graph.

10. CONSEQUENCES AND FURTHER QUESTIONS

The lower comparison $c(G) \leq c(H)$ is unconditional for these towers by Theorem 2.7. Thus a universal bound in the base order could feed back to the original graph without a cop-number loss. The conditional input in our upper theorem is the multistage base class, not a conjectural lower transfer. This distinction does not remove any of the capacity or retained-resource obstacles below.

The transfer separates three tasks: locating distinct cop resources, establishing local protection under a deadline, and capturing the robber inside the residual domain. Actual cloud ports solve the second task, and synchronized setup makes the third inexpensive when the residual domain contains at most two clouds. The present proof still assigns one cop to every target of a full remaining-ball union. The capacity obstruction therefore survives the removal of the cloud-order factor.

A concrete next target is a strategy that assigns guards to a smaller exit set and proves an affordable decrease in the residual domain. Such a lemma must count retained guards and the mobile reserve simultaneously; finite setup alone does not provide a deadline. It must also supply an internal pursuit bound for the remaining component, since a thin boundary does not make arbitrary interior pursuit inexpensive. Weakening H1 would additionally remove the automatic estimate $\Delta(G) = O(\sqrt{N})$ used to absorb the reserve here. That resource would then need a new bound.

The late-handoff example does not exclude earlier deployment or exterior prepositioning. Conversely the positive theorem does not simulate every base strategy. Even though its conclusion is on the square-root scale in the tower's order, its base hypotheses remain substantially stronger than polynomially shrinking expansion. It establishes no universal progress in the exponent of Meyniel's conjecture.

VERIFICATION

The September 9 follow-up makes the lower comparison self-contained, generalizes the short-walk count, records isometric geodesic lifts and the exact adjacent-cloud retraction criterion, and replaces the internal subdivided-clique citation with a direct strategy. The alleged three-move interception is checked in both turn orders; only the cop-first order captures, and that order does not satisfy the simulation's checkpoint. The lower comparison and the main timed upper theorem keep their conclusions.

The proofs use the pinned source versions in the bibliography and exact combinatorial arguments. Finite checks independently rebuild the HMGHM gadget, its stopped towers, retractions, and the 104-vertex handoff example. They exhaust lazy exterior shadow games on small clouds, check the first possible third-cloud visit, and verify both orders of anchor arrival and root observation in the synchronized clock. Separate checks validate the Hall-capacity inequality on finite graphs and the generic serial-setup arithmetic. These computations support identities and explicit examples; they do not certify the asymptotic theorem or replace its universal quantifiers. The accompanying research source contains the checks `hmghm-cloud-retractions.py`, `hmghm-exterior-tree-setup.py`, `hmghm-timed-handoff.py`, `synchronized-cloud-schedule.py`, `timed-cloud-deadlines.py`, and `remaining-ball-hall-capacity.py` under `checks/`. The additional `hmghm-lower-transfer.py` checks every cop-reply distance certificate and prescribed-endpoint geodesic lift in its finite sample. The check `hmghm-adjacent-cloud-retract.py` verifies the new retractions and their

exterior tree games, together with triangle and girth-five neighborhood counterexamples. Neither is a cop-number census.

REFERENCES

- [1] S. A. Hosseini, B. Mohar, and S. Gonzalez Hermosillo de la Maza, *Meyniel's conjecture on graphs of bounded degree*, Journal of Graph Theory 97 (2021), 401–407; inspected preprint: [arXiv:1912.06957v2](https://arxiv.org/abs/1912.06957v2) (2020).
- [2] A. Bonato, N. E. Clarke, S. Finbow, S. Fitzpatrick, and M.-E. Messinger, *A note on bounds for the cop number using tree decompositions*, [arXiv:1308.2839v1](https://arxiv.org/abs/1308.2839v1) (2013).
- [3] P. Prałat and N. Wormald, *Meyniel's conjecture holds for random graphs*, Random Structures & Algorithms 48 (2016), 396–421; inspected preprint: [arXiv:1301.2841v2](https://arxiv.org/abs/1301.2841v2) (2014).
- [4] P. Prałat and N. Wormald, *Meyniel's conjecture holds for random d -regular graphs*, Random Structures & Algorithms 55 (2019), 719–741; inspected preprint: [arXiv:1510.03003v2](https://arxiv.org/abs/1510.03003v2) (2018).
- [5] L. Neuwirth, *Ball-Occupation Certificates under Coarse Graph Projections: Degree Reduction, Square-Root Hard Families, and Toroidal Barriers*, preprint, September 6, 2026, Section 9, <https://levineuwirth.org/papers/ball-occupation-paper.pdf>.

BROWN UNIVERSITY