An $\Omega(\sqrt{n})$ lower bound for infeasible-start long-step interior-point methods

GPT-6 Astra prompted by Wenzhi Gao and Huikang Liu

September 5, 2026

Abstract

We replace the resisting-oracle realization in the supplied pure-centering note by one explicit linear program. An elementary invariant shows that each exact Newton direction exposes at most one additional coefficient of a fixed reverse-Helmert frame, even with unequal primal and dual steps and arbitrary data-dependent or randomized scalar centering. A two-sided positivity certificate yields an exact product obstruction and a lower bound of order $\min\{n, \sqrt{n}\log(1/\varepsilon)\}$ on direction calls. For genuinely primal-and-dual infeasible starts, the upper side of the certificate combines with the standard residual-to-complementarity neighborhood condition. The construction can satisfy both primal and dual strict feasibility and an initial point that dominates an optimum. Extensions to multiple right-hand sides are stated through an explicit tail-preservation condition. These are proposed theorems with self-contained proofs and independent algebraic/numerical checks; they are not externally reviewed claims of novelty or an improvement of the known $O(n^2\log(1/\varepsilon))$ worst-case upper bound for Algorithm IIP.

1 Main theorem and precise scope

Consider the standard-form primal-dual pair

$$\min\{c^Tx : Ax = b,\ x \geq 0\}, \qquad \max\{b^Ty : A^Ty + s = c,\ s \geq 0\}. \tag{1}$$

Write $r_p = Ax - b$, $r_d = A^Ty + s - c$, $r = (r_p, r_d)$, $G = x^Ts$, and $\mu = G/n$. A scalar-centered exact Newton call solves

$$A\Delta x = -r_p, \qquad A^T\Delta y + \Delta s = -r_d, \qquad S\Delta x + X\Delta s = \theta\mathbf{1} - x \circ s, \tag{2}$$

where $X = \mathrm{Diag}(x)$, $S = \mathrm{Diag}(s)$ and $\theta \in \mathbb{R}$ is arbitrary. In particular, $\theta = \gamma\mu$ allows all $\gamma \in [0, 1]$, including repeated pure centering $\gamma = 1$. We allow updates

$$x^+ = x + \alpha_P\Delta x, \qquad (y^+, s^+) = (y, s) + \alpha_D(\Delta y, \Delta s). \tag{3}$$

The two step lengths need not agree. All coefficients may depend on the complete LP data, the computed directions, the entire history, and random bits. We impose no permutation symmetry on the algorithm. All queried points must have $x, s > 0$. Step lengths may be arbitrary real numbers when the specified positivity and neighborhood conditions hold; the usual $[0, 1]$ restriction is a special case.

For an infeasible initial point, the basic residual condition is

$$x, s > 0, \qquad \frac{\|r\|_2}{\|r^0\|_2} \leq \frac{\mu}{\mu_0}. \tag{4}$$

One may additionally require $x_i s_i \geq \delta\mu$ for any fixed $\delta \in (0,1)$; our lower bound does not use this additional restriction. A *legal connected run* means that the entire segment

$$(x + t\alpha_P \Delta x, y + t\alpha_D \Delta y, s + t\alpha_D \Delta s), \qquad 0 \leq t \leq 1,$$

satisfies (4) at every accepted update. This includes the whole-segment line search in Algorithm IIP of Monteiro–O'Neal [4]. More general continuous paths are addressed in Section 8.

Theorem 1.1 (Strictly feasible LP, dominated infeasible start)**.** *For every even $n = 2p \geq 8$ and $\varepsilon \in (0,1)$, the explicit matrix A of Section 2, with $AA^T = 2I$, and the data*

$$\lambda = \frac{\varepsilon}{8n}, \qquad b = \lambda A\mathbf{1}, \qquad c = \lambda \mathbf{1} \tag{5}$$

have the following properties. Both primal and dual problems are strictly feasible. They have an optimal triple (x^*, y^*, s^*) satisfying $0 \leq x^*, s^* \leq \mathbf{1}$. The initial point $(x^0, y^0, s^0) = (\mathbf{1}, 0, \mathbf{1})$ is perfectly centered, has $\mu_0 = 1$, and has both $r_p^0 \neq 0$ and $r_d^0 \neq 0$.

Every legal connected run of (2)–(3) reaching $G/G_0 = \mu \leq \varepsilon$ requires a number K of Newton-direction calls satisfying

$$K > \min\left\{\frac{n}{4}, \frac{\sqrt{n}}{8}\left(1 - \frac{1}{4n}\right)\log\frac{1}{\varepsilon}\right\}. \tag{6}$$

The same fixed LP works for every such run, including every realization of a randomized method. In the ordinary one-direction-per-update method, K is also the number of iterations.

The important quantifier is

$$\forall(n, \varepsilon) \quad \exists \text{ one explicit LP} \quad \forall \text{ legal methods and random realizations.}$$

It is stronger than an algorithm-dependent completion of a resisting oracle. It is *not* a lower bound on all LP algorithms: directions and updates are restricted as stated. Solving the explicit LP by another method and inserting its answer as an iterate is outside this model.

2 The explicit matrix and its prefix flag

Use p paired coordinates, grouped as $(+1, \ldots, +p, -1, \ldots, -p)$. Let e_j be the coordinate vectors in $\mathbb{R}^p$. Define the orthogonal matrix $Q = [h_0, h_1, \ldots, h_{p-1}]$ by

$$\begin{aligned} h_0 &= p^{-1/2}\mathbf{1}, \qquad\qquad (7)\\ h_j &= -\sqrt{\frac{p-j}{p-j+1}}\, e_j + \frac{1}{\sqrt{(p-j)(p-j+1)}} \sum_{i=j+1}^{p} e_i, \quad 1 \leq j \leq p-1. \end{aligned}$$

Every h_j with $j \geq 1$ has unit norm and zero sum, and two such columns are orthogonal because the later column has zero sum on the constant tail of the earlier one. Thus $Q^T Q = I$. Put

$$B = \begin{pmatrix} Q + I \\ Q - I \end{pmatrix}, \qquad C = \begin{pmatrix} Q - I \\ Q + I \end{pmatrix}, \qquad A = \frac{C^T}{\sqrt{2}}. \tag{8}$$

Direct multiplication gives

$$B^T B = C^T C = 4I, \qquad B^T C = 0, \qquad AA^T = 2I, \tag{9}$$

$$\ker A = \operatorname{col} B, \qquad \operatorname{range} A^T = \operatorname{col} C, \qquad A\mathbf{1} = \sqrt{2p}\, e_1. \tag{10}$$

In particular, $\|A\mathbf{1}\|_2 = \|\mathbf{1}\|_2 = \sqrt{n}$. The row conditioning of A is exactly one; no ill-conditioned row system is needed for this exact-arithmetic obstruction.

Let $\mathcal{E}_r = \operatorname{span}\{e_1, \ldots, e_r\}$, with $\mathcal{E}_0 = \{0\}$, and let $\mathcal{F}_r$ be the vectors in $\mathbb{R}^p$ that are constant on indices $r+1, \ldots, p$. The triangular support of (7) gives

$$Q\mathcal{E}_{r+1} = \mathcal{F}_r, \qquad 0 \le r < p. \tag{11}$$

If $u, v \in \mathcal{E}_r$, then each vector in

$$x = a\mathbf{1} + Bu, \qquad s = b_0\mathbf{1} + Cv \tag{12}$$

is individually constant on all $2(p-r)$ coordinates in pairs $i > r$. The two constants need not be equal, and no identity $s = Jx$ is imposed. We use b_0 for the slack baseline to distinguish it from the LP right-hand side b.

3 One exact Newton call exposes at most one coefficient

Lemma 3.1 (Prefix propagation)**.** *Suppose $x, s > 0$ have the form (12) with $u, v \in \mathcal{E}_r$, $r < p$. Consider an exact solve*

$$Ad_x = \zeta_P A\mathbf{1}, \qquad A^T d_y + d_s = \zeta_D \mathbf{1}, \qquad Sd_x + Xd_s = h, \tag{13}$$

where ζ_P, ζ_D are scalars and h has a single common value on all coordinates in pairs $i > r$. Then

$$d_x = \zeta_P\mathbf{1} + B\xi, \qquad d_s = \zeta_D\mathbf{1} + Cv, \qquad \xi, v \in \mathcal{E}_{r+1}.$$

Proof. The first two equations and (10) give these representations without the support conclusion. Subtract the scalar particular terms from the third equation. Its new right-hand side $\widehat{h} = h - \zeta_P s - \zeta_D x$ is still constant on the untouched tail. Let the positive tail values of x, s be $\bar{x}, \bar{s}$. On every $i > r$, the plus and minus equations are

$$\bar{s}(Q\xi + \xi)_i + \bar{x}(Qv - v)_i = \widehat{h}_i, \quad \bar{s}(Q\xi - \xi)_i + \bar{x}(Qv + v)_i = \widehat{h}_i.$$

Their difference says $w := \bar{s}\xi - \bar{x}v \in \mathcal{E}_r$. Their sum says $Qz \in \mathcal{F}_r$, where $z := \bar{s}\xi + \bar{x}v$. By (11), $z \in \mathcal{E}_{r+1}$. Since $\bar{x}, \bar{s} > 0$, $\xi = (z + w)/(2\bar{s})$ and $v = (z - w)/(2\bar{x})$ are in $\mathcal{E}_{r+1}$. The Newton system is nonsingular: its homogeneous equations imply $d_x^T d_s = 0$ and $d_s = -X^{-1}Sd_x$, hence $d_x = d_s = d_y = 0$. ∎

For $b = \lambda A\mathbf{1}, c = \lambda\mathbf{1}$, let the residual factors be

$$r_p = (1-\lambda)\eta_P A\mathbf{1}, \qquad r_d = (1-\lambda)\eta_D\mathbf{1}. \tag{14}$$

They start at $\eta_P = \eta_D = 1$ and update as $\eta_P^+ = (1-\alpha_P)\eta_P$ and $\eta_D^+ = (1-\alpha_D)\eta_D$. Consequently every state has

$$x = a\mathbf{1} + Bu, \quad s = b_0\mathbf{1} + Cv, \quad a = \lambda + (1-\lambda)\eta_P, \quad b_0 = \lambda + (1-\lambda)\eta_D. \tag{15}$$

Initially $u = v = 0$. The usual right-hand side $h = \theta\mathbf{1} - x \circ s$ is constant on the untouched tail, and (2) has the form (13). By induction, after $K < p$ calls, every point on the accepted segments has $u, v \in \mathcal{E}_K$. Rank-preserving and zero directions cause no problem: K is an upper bound on the prefix length, not an assertion that every call increases rank.

For the feasible instance $b = A\mathbf{1}, c = \mathbf{1}$, the same argument gives $a = b_0 = 1$, with no residual-neighborhood assumption.

4 An exact two-sided positivity certificate

Define

$$R_r = \prod_{j=0}^{r-1} \frac{\sqrt{p-j}-1}{\sqrt{p-j}+1}, \qquad 0 \le r < p, \qquad R_0 = 1. \tag{16}$$

The following certificate is the additional step that recovers logarithmic accuracy dependence from the static geometry of the supplied note.

Lemma 4.1 (Sharp prefix interval)**.** *If $\tau > 0$, $z \in \mathcal{E}_r$, and $\tau\mathbf{1} + Bz > 0$, then for $1 \le r < p$,*

$$R_r < g_\tau(z) := 1 + \frac{2z_1}{\tau\sqrt{p}} < R_r^{-1}. \tag{17}$$

Both endpoints are attained on the nonnegative closure. The same statement holds with C in place of B.

Proof. Let $t_0 = \tau + z_1/\sqrt{p}$, set $q_i = -z_i$ for $2 \le i \le r$, and set $q_{r+1} = 0$. For the lower endpoint set $q_1 = -z_1$; for the upper endpoint set $q_1 = z_1$. For $1 \le i \le r$ define

$$m_i = p - i + 1, \qquad \alpha_i = \sqrt{\frac{m_i - 1}{m_i}}, \qquad \beta_i = \frac{1}{\sqrt{m_i(m_i-1)}},$$
$$\epsilon_i = t_{i-1} + \alpha_i q_{i+1} - q_i, \qquad t_i = t_{i-1} - \beta_i q_{i+1}.$$

By the formula for Q, each ϵ_i is a positive coordinate of $\tau\mathbf{1} + Bz$: for the lower endpoint all are plus coordinates; for the upper endpoint the first is the minus coordinate of pair 1 and the rest are plus coordinates. Define

$$L_i^{\pm} = t_{i-1} \pm \frac{q_i}{\sqrt{m_i}}.$$

Elementary algebra gives the exact recurrences

$$L_{i+1}^{+} = \frac{\sqrt{m_i}}{\sqrt{m_i}+1} L_i^{+} + \frac{\epsilon_i}{\sqrt{m_i}+1}, \tag{18}$$

$$L_{i+1}^{-} = \frac{\sqrt{m_i}}{\sqrt{m_i}-1} L_i^{-} - \frac{\epsilon_i}{\sqrt{m_i}-1}. \tag{19}$$

Since $q_{r+1} = 0$, the two terminal values both equal t_r. Set

$$P^{\pm} = \prod_{i=1}^{r} \frac{\sqrt{m_i}}{\sqrt{m_i} \pm 1}.$$

Iteration of (18)–(19) gives the explicit positive certificate

$$P^- L_1^- - P^+ L_1^+ = \sum_{i=1}^{r} \epsilon_i \left[\frac{\prod_{j=i+1}^{r} \frac{\sqrt{m_j}}{\sqrt{m_j}+1}}{\sqrt{m_i}+1} + \frac{\prod_{j=i+1}^{r} \frac{\sqrt{m_j}}{\sqrt{m_j}-1}}{\sqrt{m_i}-1} \right] > 0. \tag{20}$$

In the lower construction $(L_1^+, L_1^-) = (\tau, \tau g_\tau(z))$; in the upper construction it is $(\tau g_\tau(z), \tau)$. Since $P^+/P^- = R_r$, both inequalities in (17) follow. No sign assumption on the intermediate q_i was used.

For sharpness, set all $\epsilon_i = 0$, choose $t_r = T > 0$ and $q_{r+1} = 0$, and reconstruct backwards using

$$t_{i-1} = t_i + \beta_i q_{i+1}, \qquad q_i = t_i + \frac{q_{i+1}}{\alpha_i}.$$

All q_i and t_i are positive. The selected coordinates are zero, their paired complements are $2q_i > 0$, and the untouched coordinates equal $T > 0$. Set $z_1 = -q_1$ for the lower case and $z_1 = q_1$ for the upper case, with $z_i = -q_i$ for $i \geq 2$. The required $\tau = t_0 - z_1/\sqrt{p}$ is positive: the homogeneous recurrences give $L_1^{\pm} = T/P^{\pm} > 0$. Scaling achieves any prescribed $\tau > 0$. Finally $C = JB$, where J swaps the plus and minus blocks, so its positivity interval is identical. □

Lemma 4.2 (Two-sided complementarity estimate)**.** *Let $x = a\mathbf{1} + Bu > 0$, $s = b_0\mathbf{1} + Cv > 0$, with $u, v \in \mathcal{E}_r$ and $1 \leq r < p$. Then $a, b_0 > 0$ and*

$$R_r < \frac{\mu}{ab_0} < R_r^{-1}. \tag{21}$$

For $r = 0$, the middle expression equals one.

Proof. Positivity of $a\mathbf{1}+Bu$ means $a\mathbf{1}+Qu > |u|$ componentwise. If $a \leq 0$, this implies $Qu > |u| \geq 0$, contradicting $\|Qu\|_2 = \|u\|_2$. Thus $a > 0$, and the same reasoning gives $b_0 > 0$. By (9) and $Q^T\mathbf{1} = \sqrt{p}\, e_1$,

$$\mu = ab_0 + \frac{b_0 u_1 + a v_1}{\sqrt{p}}, \qquad \frac{\mu}{ab_0} = \frac{g_a(u) + g_{b_0}(v)}{2}.$$

Apply Lemma 4.1 separately to x and s. □

5 Feasible starts and the accuracy-dependent bound

For the feasible LP $b = A\mathbf{1}, c = \mathbf{1}$ and start $(\mathbf{1}, 0, \mathbf{1})$, prefix propagation and Lemma 4.2 immediately imply

$$\frac{G_K}{G_0} = \mu_K > R_K, \qquad 1 \leq K < p. \tag{22}$$

Only positivity at the queried and accepted points is needed. There is no centrality, residual, equal-step, gap-monotonicity, or information restriction.

Lemma 5.1 (Converting the product to a call bound)**.** *For $n = 2p \geq 8$ and $1 \leq K \leq p/2$,*

$$-\log R_K = 2\sum_{j=0}^{K-1} \operatorname{artanh} \frac{1}{\sqrt{p-j}} \leq \frac{8K}{\sqrt{n}}. \tag{23}$$

Consequently, if $R_K < q < 1$ whenever $K < p$, then

$$K > \min\left\{\frac{n}{4}, \frac{\sqrt{n}}{8}\log\frac{1}{q}\right\}. \tag{24}$$

Proof. For $0 \leq z \leq 1/\sqrt{2}$, $\operatorname{artanh} z \leq z/(1-z^2) \leq 2z$. In the stated range, $(p-j)^{-1/2} \leq \sqrt{2/p}$, giving $-\log R_K \leq 4\sqrt{2}K/\sqrt{p} = 8K/\sqrt{n}$. If $K \leq p/2$, combine this with $\log(1/q) < -\log R_K$; if $K > p/2$, the first alternative already holds. The case $K = 0$ cannot reach a target below the initial value. □

Thus the feasible instance needs $K > \min\{n/4, (\sqrt{n}/8)\log(1/\varepsilon)\}$ calls. The exact product is stronger than the convenient constant 8. For each fixed $t = \log(1/\varepsilon) > 0$, it gives

$$\liminf_{n\to\infty} \frac{K}{\sqrt{n}} \geq \frac{t}{2\sqrt{2}},$$

because $-\log R_K = 2K/\sqrt{p} + o(1)$ when $K = O(\sqrt{p})$. The argument saturates at order n; it does not prove an unsaturated $\Omega(\sqrt{n}\log(1/\varepsilon))$ bound for arbitrarily small ε at fixed dimension.

6 A genuinely infeasible start with a dominated optimum

First take the simplest target, $b = c = 0$, so that $\lambda = 0$. The optimal triple $(0, 0, 0)$ exists and is dominated by the initial $(\mathbf{1}, 0, \mathbf{1})$. Both initial residuals are nonzero. Along any Newton run,

$$x = \eta_P \mathbf{1} + Bu, \qquad s = \eta_D \mathbf{1} + Cv.$$

By Lemma 4.2, positivity forces $\eta_P, \eta_D > 0$. The row balance in (10) yields

$$\frac{\|r\|_2}{\|r^0\|_2} = \sqrt{\frac{\eta_P^2 + \eta_D^2}{2}}. \tag{25}$$

Under (4), $\eta_P\eta_D \leq (\eta_P^2 + \eta_D^2)/2 \leq \mu^2$. The *upper* half of (21), not its lower half, therefore gives

$$\mu < \frac{\eta_P\eta_D}{R_K} \leq \frac{\mu^2}{R_K}, \qquad \text{hence} \qquad \mu > R_K. \tag{26}$$

This remains true when the actual support is zero and $K \geq 1$, since then $\mu = \eta_P\eta_D$ and $R_K < 1$. The same call bound as in the feasible case follows. Only accepted *endpoints* need satisfy the residual condition for this zero-shift instance. A common step gives $r = \eta r^0$, so the argument then works for any norm in the neighborhood, without the balancing calculation.

This example proves a direct infeasible-start obstruction under the start-dominates-an-optimum assumption. It does not have strictly feasible primal and dual target points; the next section removes that limitation.

7 Strict feasibility by a small shift: proof of Theorem 1.1

Fix $0 < \lambda < 1$ and use $b = \lambda A\mathbf{1}, c = \lambda\mathbf{1}$. The triple $(\lambda\mathbf{1}, 0, \lambda\mathbf{1})$ is strictly feasible, although it is not generally optimal. The initial residuals are

$$r_p^0 = (1-\lambda)A\mathbf{1} \neq 0, \qquad r_d^0 = (1-\lambda)\mathbf{1} \neq 0.$$

For the unscaled LP $b = A\mathbf{1}, c = \mathbf{1}$, an optimum exists: the primal feasible set is nonempty and objective sublevel sets $\{x \geq 0 : \mathbf{1}^T x \leq M\}$ are compact; LP duality gives a dual optimum. Every optimal triple obeys

$$\mathbf{1}^T x^* + \mathbf{1}^T s^* = n, \tag{27}$$

because $\mathbf{1}^T x^* = (A\mathbf{1})^T y^*$ and $s^* = \mathbf{1} - A^T y^*$. In particular $0 \leq x^*, s^* \leq n\mathbf{1}$. Scaling all three optimal variables by λ gives an optimal triple of the shifted LP. For $\lambda \leq 1/n$ it is dominated by the initial point. Also $s^0 \geq |c - A^T y^0| = \lambda\mathbf{1}$, if this additional initial condition is required by an iterative-solver analysis.

By (15) and (25), the residual condition gives

$$\begin{aligned} ab_0 &= \lambda^2 + \lambda(1-\lambda)(\eta_P + \eta_D) + (1-\lambda)^2 \eta_P \eta_D \\ &\le [\lambda + (1-\lambda)\mu]^2. \end{aligned} \tag{28}$$

Here $\eta_P + \eta_D \le 2\mu$ and $\eta_P \eta_D \le \mu^2$ follow from $\eta_P^2 + \eta_D^2 \le 2\mu^2$. Their signs are irrelevant. Every positive state in the prefix after at most $K < p$ calls therefore satisfies

$$\mu R_K < [\lambda + (1-\lambda)\mu]^2, \qquad K \ge 1. \tag{29}$$

An important distinction is that (29) alone does not give a global lower bound on μ: it may also hold on a disconnected small-μ branch. Instead, a legal connected run that starts at $\mu = 1$ and reaches $\mu \le \varepsilon$ must pass through a point with $\mu = \varepsilon$. All directions used on that path are among the K counted calls. At this crossing, (29) forces

$$R_K < \Phi_\lambda(\varepsilon) := \frac{[\lambda + (1-\lambda)\varepsilon]^2}{\varepsilon}. \tag{30}$$

No monotonicity of μ between calls or along segments is assumed.

For $\lambda = \varepsilon/(8n)$ and $t = \log(1/\varepsilon)$,

$$\begin{aligned} \log \frac{1}{\Phi_\lambda(\varepsilon)} &= t - 2\log\left(1 + \frac{\lambda(1-\varepsilon)}{\varepsilon}\right) \\ &\ge t - \frac{1-\varepsilon}{4n} \ge \left(1 - \frac{1}{4n}\right) t > 0. \end{aligned}$$

Thus $\Phi_\lambda(\varepsilon) < 1$. Apply Lemma 5.1 with $q = \Phi_\lambda(\varepsilon)$ to obtain (6). All instance data were selected before any algorithmic choices, proving the uniform and pathwise randomized assertions. This completes the proof. □

8 Extensions and boundaries of the result

8.1 Multiple right-hand sides and history-dependent corrections

The operative restriction is not a particular formula for γ. It is the following *tail-preserving Newton-span computation*:

1. Every queried base point is positive and equals the initial point plus a scalar linear combination of previously computed directions. Primal and dual/slack coefficients may differ, with the dual multiplier and slack combined consistently.

2. Each new direction solves (13). Its complementarity right-hand side is constant on all paired coordinates not yet exposed by the computation. Its linear-constraint right-hand sides lie on the scalar rays displayed there.

3. New states and any accepted continuous path remain in the initial point plus the span of already computed directions. The infeasible cases obey the residual requirement used above (at endpoints for $\lambda = 0$, along the path for $\lambda > 0$).

For the feasible-start extension, require $\zeta_P = \zeta_D = 0$, so that all generated states remain affinely feasible. Induction with Lemma 3.1 proves exactly the same K-prefix bound for this larger class.

Componentwise operations on existing vectors with uniform untouched tails preserve that property. Thus, for example, an affine predictor followed by a correction with

$$h = \gamma\mu\mathbf{1} - x \circ s - d_x^{\mathrm{aff}} \circ d_s^{\mathrm{aff}}$$

exposes at most two new coefficients: the product is constant after the predictor's expanded prefix, and the second solve adds at most one more. Homogeneous corrector constraint equations correspond to $\zeta_P = \zeta_D = 0$. Repeated product corrections and scalar combinations of computed directions are covered when the listed conditions hold. An arc assembled from such directions is also covered if it stays in the same prefix and satisfies the requisite neighborhood condition throughout.

Each right-hand-side direction counts as a call, including an auxiliary predictor that is not accepted as a step. Reusing a factorization does not make a second direction free. If there are at most q calls per reported iteration, the iteration lower bound is the direction-call bound divided by q. We do not claim that every predictor-corrector or arc-search method satisfies these hypotheses without inspecting its equations and line search.

8.2 A relaxed residual constant and arbitrary common-step norms

For $\kappa \geq 1$, replace (4) by $\|r\|_2/\|r^0\|_2 \leq \kappa\mu$. The same proof gives

$$\mu > R_K/\kappa^2 \quad (\lambda = 0), \qquad R_K < \frac{[\lambda + (1-\lambda)\kappa\varepsilon]^2}{\varepsilon} \quad (\lambda > 0 \text{ with a connected run}).$$

These imply logarithmic call bounds whenever the corresponding right-hand side is below one. This qualification is necessary. With common primal-dual steps, $r = \eta r^0$ and these conclusions hold for any norm; unequal steps use the balanced Euclidean calculation above.

8.3 Any prescribed positive, exactly centered initial pair

The normalization $x^0 = s^0 = \mathbf{1}$ is not essential. Prescribe any positive pair $(\widehat{x}^0, \widehat{s}^0)$ with $\widehat{x}_i^0 \widehat{s}_i^0 = \widehat{\mu}_0$ for every i. Let

$$D = \mathrm{Diag}(\widehat{x}^0), \qquad \rho = \frac{\|\widehat{s}^0\|_2}{\sqrt{n}}, \qquad \widehat{A} = \rho A D^{-1}, \quad \widehat{b} = \rho b, \quad \widehat{c} = \widehat{\mu}_0 D^{-1} c.$$

The change of variables $\widehat{x} = Dx$, $\widehat{s} = \widehat{\mu}_0 D^{-1} s$, $\widehat{y} = (\widehat{\mu}_0/\rho) y$ preserves the scalar-centered Newton equations, step lengths, relative complementarity, strict feasibility, and initial domination of an optimum. On the residual rays,

$$\widehat{r}_p = \rho(1-\lambda)\eta_P A\mathbf{1}, \qquad \widehat{r}_d = (1-\lambda)\eta_D \widehat{s}^0.$$

The two initial residual blocks again have equal Euclidean norm, so the normalized residual is still $\sqrt{(\eta_P^2 + \eta_D^2)/2}$. All the same bounds hold from $(\widehat{x}^0, 0, \widehat{s}^0)$, including with unequal steps. Tail preservation is understood in the normalized coordinates. The transformed matrix remains full row rank, although its row condition number need no longer be one. This does not assert a corresponding result for every noncentered initial pair.

8.4 What has not been proved

The theorem does not cover coordinate-dependent forcing on untouched coordinates, arbitrary vector-valued centering, arbitrary inexact Newton errors, unrestricted reformulations or embeddings, changed barrier metrics, or updates outside the allowed direction span. Full knowledge of A, b, c is allowed for choosing scalar coefficients, but not for injecting arbitrary new coordinate directions. No claim is made about bit complexity, rational approximations of the radical entries of Q, or stability under finite precision: the invariant is an exact-arithmetic statement.

The shift $\lambda = \varepsilon/(8n)$ depends on the target accuracy. This is legitimate for a worst-case bound over instances, but is not a uniform-in-ε claim for one strictly feasible LP with a fixed positive shift. The zero-shift and feasible instances themselves do not depend on the target. General conic and nonlinear programs are not asserted to inherit the proof merely by replacing n with a barrier parameter.

9 Relation to the supplied note and verification

The supplied report [1] motivates the reverse-Helmert geometry and the use of terminal positivity to handle pure-centering calls. Its stated bound is $(1 - e^{-t})\sqrt{n}/8$ in a deterministic symmetric information-restricted model. Here the fixed matrix and prefix proof avoid the resisting-oracle completion and zero-pivot compression entirely. The two-sided certificate sharpens the static estimate to an exact product; its upper side is what makes the direct dominated infeasible-start construction possible. The small-shift crossing argument supplies strict primal and dual feasibility without dropping the dominance assumption.

The independent verification script accompanies this note. It constructs the full KKT matrix, without projecting directions onto the claimed prefix, and checks asymmetric states, unequal steps, pure centering, other scalar parameters, and a predictor/product-corrector pair. Separate auxiliary linear programs minimize and maximize g_τ over the entire nonnegative prefix polyhedron. Both extrema agree with R_r and R_r^{-1} for tested p up to 512. Symbolic simplification verifies both scalar recurrences (18)–(19) exactly. Further tests check shifted optimal solutions, initial dominance, residual rays and infeasible trajectories. These checks support the algebra; the proofs above, not floating-point rank decisions, are the basis of the claims.

The result remains a proposed, unreviewed theorem. An independent adversarial proof audit and a broader literature comparison are required before claiming publication-level novelty. In particular, it strengthens the lower-bound side of the present investigation; it does not close the gap to the $O(n^2 \log(1/\varepsilon))$ upper bound for the specific long-step infeasible-start Algorithm IIP [4].

References

[1] *Square-Root Lower Bounds with Pure-Centering Steps: A Static Reverse-Helmert Rank Obstruction in the Todd–Ye Resisting-Oracle Model*. Revised technical report, September 4, 2026. User-supplied file `pure_centering_sqrt_lower_bound_revised.pdf`.

[2] M. J. Todd. A lower bound on the number of iterations of primal-dual interior-point methods for linear programming. In *Numerical Analysis 1993*, Pitman Research Notes in Mathematics 303, pp. 237–259. Longman, 1994.

[3] M. J. Todd and Y. Ye. A lower bound on the number of iterations of long-step primal-dual linear programming algorithms. *Annals of Operations Research* 62, 233–252, 1996. https://doi.org/10.1007/BF02206818.

[4] R. D. C. Monteiro and J. W. O'Neal. *Convergence Analysis of a Long-Step Primal-Dual Infeasible Interior-Point LP Algorithm Based on Iterative Linear Solvers*. October 31, 2003. Algorithm IIP and Theorem 2.1. https://optimization-online.org/2003/10/768/.