

Straight Line Complexity

In this lecture, we develop the concept of *straight-line complexity* (SLC) of the central path, which will give a combinatorial measure of its complexity. We will in particular relate the complexity of the central path to a geometric measure of complexity of certain (shadow vertex) *simplex paths*. We will then further connect the complexity to circuits of the constraints of the constraint matrix.

As in prior lectures, we are interested in solving the following primal-dual pair:

$$\begin{aligned} \min \langle c, x \rangle & & \max \langle b, y \rangle \\ \mathbf{A}x = b & & \mathbf{A}^\top y + s = c \\ x \geq \mathbf{0}_n, & & s \geq \mathbf{0}_n, \end{aligned} \tag{LP}$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and \mathbf{A} has rank m . $\mathcal{P}(\mathcal{P}_{++})$, $\mathcal{D}(\mathcal{D}_{++})$, denote the primal and dual (strictly) feasible regions. Let $z^* = (x^*, s^*, y^*)$ denote the optimal primal-dual pair and let $v^* = \langle c, x^* \rangle = \langle y^*, b \rangle$ denote the optimal value.

The central path $\text{CP} := \{z^{\text{CP}}(\mu) := (x^{\text{CP}}(\mu), s^{\text{CP}}(\mu), y^{\text{CP}}(\mu)) \in \mathcal{P}_{++} \times \mathcal{D}_{++} \mid \mu > 0\}$ is defined by the strictly feasible pairs satisfying the centrality equations $x^{\text{CP}}(\mu)s^{\text{CP}}(\mu) = \mu \mathbf{1}_n$. We use the notation z^{CP} in this lecture to disambiguate with the *max central path*, a combinatorial proxy for central path, that we define later.

The basic concept we will use to measure complexity is as follows:

DEFINITION 1 Let $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a function and $\eta \in (0, 1]$. The straight line complexity of f with respect to η , denoted $\text{SLC}_\eta(f)$, is the minimum number of pieces of a continuous piecewise linear function $h: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfying $\eta f \leq h \leq f$. We write $\text{SLC}_\eta(f, \mu_0)$ to indicate the same quantity when restrict the approximation of f to the interval $[0, \mu_0]$.

With this definition, one can express the performance of the interior point method we have developed as follows:

THEOREM 2 ([ADL+22]) Let $z_0 \in N_2(\beta)$, $\beta \in (0, 1/6)$ with $\mu_0 := \mu(z) > 0$. Then, for $\eta \in (0, 1]$, the TRUST-REGION IPM on initial iterate z_0 outputs an optimal solution $z^* = (x^*, s^*, y^*) \in \bar{N}_2(\beta)$ in at number of iterations bounded by

$$O \left(\min_{\eta \in (0, 1]} \frac{\sqrt{n}}{\beta} \log \left(\frac{n}{\beta \eta} \right) \sum_{i=1}^n \text{SLC}_\eta(x_i^{\text{CP}}(\cdot), \mu_0) \right).$$

At a high level, the proof of the above theorem follows by breaking up the central path into approximately linear chunks (η is the approximation factor), and showing that each linear chunk can be traversed using $O(1)$ affine scaling and trust-region steps in an amortized sense. We do not attempt to prove this here, and instead use it to prove complexity bounds.

Importantly, the straight-line complexity of the components of the central path naturally correspond to *lower bounds* on the iteration complexity of a large family of IPMs. Indeed, if one examines the trajectory of almost any IPM, it produces a piecewise linear trajectory that stays multiplicatively close to the central path. Each piece of the trajectory generally corresponds to a single iteration, and hence any lower bound on the number of pieces yields a lower bound on the number of iterations. The following result of Allamigeon, Benchimol, Gaubert and Joswig [ABGJ18] in fact proves that the straight-line complexity of the central can be exponential.

THEOREM 3 ([ABGJ18]) *Examine the following linear program induced by a parameter t :*

$$\begin{aligned}
& \min x_1 \\
& x_1 \leq t^2 \\
& x_2 \leq t \\
& x_{2j+1} \leq tx_{2j-1}, x_{2j+1} \leq tx_{2j}, 1 \leq j \leq n \\
& x_{2j+2} \leq t^{1-1/2^j} (x_{2j-1} + x_{2j}), 1 \leq j < n, \\
& x_{2n-1} \geq 0, x_{2n} \geq 0.
\end{aligned} \tag{LP(t)}$$

Then for $\eta \in (0, 1)$ and $t \geq 2^{\Omega(2^n(n \log(n/\eta)))}$, we have that $\text{SLC}_\eta(x_{2n}^{\text{cp}}(\cdot)) \geq 2^n - 1$.

The above result proves that interior point methods that follow the central path may require an exponential number of iterations, even when the linear program has $O(n)$ inequalities. Note that this does not contradict the polynomiality of IPMs, as the bit-complexity of the above linear program is exponential (i.e., $t \geq 2^{2^n}$).

Organization. In what follows, we examine the consequences of Theorem 2. In Section 1, we show that SLC of the central path can be captured by the shape of specific shadow vertex simplex paths. This will in particular, yield an exponential upper bound on SLC, since the number of edges of a simplex path on a polytope with n inequalities is at most 2^n . The material from this section is derived from [ADL⁺22]. In Section 2, we relate the SLC to the shape of the circuits of \mathbf{A} , the minimal linear dependencies, and introduce the notion of a circuit cover. In Section 3, we use this notion to show that the SLC of a linear program can be upper bounded in terms of the logarithm of the *circuit imbalance measure* of the matrix \mathbf{A} , which measures the ratios of non-zero elements in minimal linear dependencies. The material in these sections is derived from [DKN⁺24].

1 The Maximum Central Path

In this section, we introduce the *max central path*, our combinatorial proxy for the central path. This proxy will allow us to give a easy upper bound on straight-line complexity in terms of the number of vertices of the primal or dual polyhedron.

Let v^* denote the optimal value of (LP). Given $g \geq 0$, we denote the gap truncated primal and dual feasible regions by

$$\begin{aligned}
\mathcal{P}_g &:= \{x \in \mathbb{R}^n \mid \mathbf{A}x = b, x \geq \mathbf{0}, \langle c, x \rangle \leq v^* + g\}, \\
\mathcal{D}_g &:= \{s \in \mathbb{R}^n \mid \exists y \in \mathbb{R}^m \mathbf{A}^\top y + s = c, s \geq \mathbf{0}, \langle b, y \rangle \geq v^* - g\}.
\end{aligned}$$

These sets correspond to the primal and dual feasible points $(x, s, y) \in \mathcal{P} \times \mathcal{D}$ with objective value within g from the optimum v^* , respectively. Under the assumption that the primal and dual are non-empty, recall that both programs have the same value, and hence $\mathcal{P}_g, \mathcal{D}_g$ are also non-empty for all $g \geq 0$. If we further assume that the primal and dual are strictly feasible, as is usual in the IPM context, then $\mathcal{P}_g, \mathcal{D}_g$ will in fact be bounded for all $g \geq 0$ (we leave this as an exercise to the reader). We will thus assume throughout that both $\mathcal{P}_g, \mathcal{D}_g$ is are bounded.

Let (x^*, s^*, y^*) , v^* denote the optimal primal-dual pair and value as above. For any $(x, y, s) \in \mathcal{P} \times \mathcal{D}$, we derive the following identities directly from the gap formula:

$$\begin{aligned} \langle x, s^* \rangle &= \langle c, x \rangle - \langle y^*, b \rangle = \langle c, x \rangle - v^* \\ \langle x^*, s \rangle &= \langle c, x^* \rangle - \langle y, b \rangle = v^* - \langle y, b \rangle \\ \langle x, s \rangle &= \langle c, x \rangle - \langle y, b \rangle = (\langle c, x \rangle - v^*) + (v^* - \langle y, b \rangle) = \langle x, s^* \rangle + \langle x^*, s \rangle. \end{aligned} \quad (1)$$

Given the identities above, the sets \mathcal{P}_g and \mathcal{D}_g are equivalently given by

$$\mathcal{P}_g = \{x \in \mathcal{P} \mid \langle x, s^* \rangle \leq g\}, \quad \mathcal{D}_g = \{(s, y) \in \mathcal{D} \mid \langle x^*, s \rangle \leq g\}.$$

These expressions are in fact independent of the choice of optimal solutions (x^*, s^*, y^*) .

We now define *max central path* as the parametric curve

$$\text{MCP} := \{z^m(g) := (x^m(g), s^m(g)) \in \mathbb{R}_+^{2n} \mid g \geq 0\},$$

where

$$x_i^m(g) := \max\{x_i \mid x \in \mathcal{P}_g\} \quad \text{and} \quad s_i^m(g) := \max\{s_i \mid (s, y) \in \mathcal{D}_g\}, \quad \forall i \in [n]. \quad (2)$$

By our boundedness assumption on $\mathcal{P}_g, \mathcal{D}_g$, the max central path is well-defined for all $g \geq 0$. We note that $z^m(g)$ does not necessarily correspond to the slacks (x, s) of a single feasible primal dual point $(x, s, y) \in \mathcal{P} \times \mathcal{D}$. This is because each coordinate is defined as the value of a different optimization problem.

Irrespective of this, the MCP will provide a good approximation of the coordinates of the central path after an appropriate reparametrization, which we prove presently.

LEMMA 4 For every $\mu > 0$ and the central path point $z^{\text{CP}}(\mu) = (x^{\text{CP}}(\mu), s^{\text{CP}}(\mu), y^{\text{CP}}(\mu))$, then

$$\frac{1}{2n}(x^m(n\mu), s^m(n\mu)) \leq (x^{\text{CP}}(\mu), s^{\text{CP}}(\mu)) \leq (x^m(n\mu), s^m(n\mu)). \quad (3)$$

PROOF: Recall that

$$\langle x^{\text{CP}}(\mu), s^{\text{CP}}(\mu) \rangle = \underbrace{\langle x^{\text{CP}}(\mu), s^* \rangle}_{\geq 0} + \underbrace{\langle x^*, s^{\text{CP}}(\mu) \rangle}_{\geq 0} = n\mu.$$

using Equation (1). Therefore, $x^{\text{CP}}(\mu) \in \mathcal{P}_{n\mu}$ and $s^{\text{CP}}(\mu) \in \mathcal{D}_{n\mu}$. By definition of the max central path, $(x^{\text{CP}}(\mu), s^{\text{CP}}(\mu)) \leq (x^m(n\mu), s^m(n\mu))$.

For the second inequality, let us examine a single coordinate of x_i^m and s_i^m , $i \in [n]$. Now let $x^{(i)} \in \mathcal{P}_{n\mu}$ be the point satisfying $x_i^{(i)} = x_i^m(n\mu)$ and $(s^{(i)}, y^{(i)}) \in \mathcal{D}_{n\mu}$ be the point satisfying $s_i^{(i)} = s_i^m(n\mu)$. Then,

$$\frac{x_i^m(n\mu)}{x_i^{\text{CP}}(\mu)} = \frac{x_i^{(i)} s_i^{\text{CP}}(\mu)}{\mu} \leq \frac{\langle x^{(i)}, s^{\text{CP}}(\mu) \rangle}{\mu} = \frac{\langle x^{(i)}, s^* \rangle + \langle x^*, s^{\text{CP}}(\mu) \rangle}{\mu} \leq \frac{n\mu + n\mu}{\mu} = 2n.$$

By a symmetric argument on the dual, $s_i^m(n\mu) = s_i^{(i)} \leq 2n s_i^{\text{CP}}(\mu)$. The claim thus follows. \square

From the sandwiching relation between the central path and max central path, we derive the corresponding relations between the straight-line complexities as a simplex corollary.

COROLLARY 5 For $\eta \in (0, 1]$, $T > 0$, we have that

$$\begin{aligned} \text{SLC}_{\eta/(2n)}(x_i^m, nT) &\leq \text{SLC}_{\eta}(x_i^{\text{CP}}, T), & \text{SLC}_{\eta/(2n)}(x_i^{\text{CP}}, T) &\leq \text{SLC}_{\eta}(x_i^m, nT), \\ \text{SLC}_{\eta/(2n)}(s_i^m, nT) &\leq \text{SLC}_{\eta}(s_i^{\text{CP}}, T), & \text{SLC}_{\eta/(2n)}(s_i^{\text{CP}}, T) &\leq \text{SLC}_{\eta}(s_i^m, nT). \end{aligned}$$

PROOF: By symmetry, we only show the statement for the primal paths.

$\text{SLC}_{\eta/(2n)}(x_i^m, nT) \leq \text{SLC}_\eta(x_i^{\text{CP}}, T)$. Let $h : [0, T] \rightarrow \mathbb{R}_+$ be the optimal piecewise-linear approximation of x_i^{CP} satisfying $\eta x_i^{\text{CP}} \leq h \leq x_i^{\text{CP}}$ on $[0, T]$. Define $\bar{h}(g) = h(g/n)$, for $g \in [0, nT]$. Then by Lemma 4, we have that

$$\frac{\eta x_i^m(g)}{2n} \leq \eta x_i^{\text{CP}}(g/n) \leq \bar{h}(g) := h(g/n) \leq x_i^{\text{CP}}(g/n) \leq x_i^m(g), g \in [0, nT].$$

Therefore, \bar{h} is an $\eta/(2n)$ approximation of x_i^m on $[0, nT]$. Since \bar{h} has the same number of pieces as h , the claim follows.

$\text{SLC}_{\eta/2n}(x_i^{\text{CP}}, T) \leq \text{SLC}_\eta(x_i^m, nT)$. Let $h : [0, nT] \rightarrow \mathbb{R}_+$ be the optimal piecewise-linear approximation of x_i^m satisfying $\eta x_i^m \leq h \leq x_i^m$ on $[0, nT]$. Define $\bar{h}(\mu) = h(n\mu)/(2n)$, for $\mu \in [0, T]$. Then by Lemma 4, we have that

$$\frac{\eta x_i^{\text{CP}}(\mu)}{2n} \leq \frac{\eta x_i^m(n\mu)}{2n} \leq \bar{h}(\mu) := \frac{h(n\mu)}{2n} \leq \frac{x_i^m(n\mu)}{2n} \leq x_i^{\text{CP}}(\mu), g \in [0, \mu].$$

Therefore, \bar{h} is an $\eta/(2n)$ approximation of x_i^{CP} on $[0, \mu]$. Since \bar{h} has the same number of pieces as h , the claim follows. \square

Note that reducing the approximation quality from η to $\eta/(2n)$ in Theorem 2 only changes the number of iterations by a constant factor. Therefore, from the perspective of Theorem 2, one can replace x_i^{CP} with x_i^m in the SLC bound at essentially no loss.

1.1 The Max Central Path and The Simplex Method

We now give a geometric interpretation of each coordinate of the max central path coordinates in terms of shadow vertex simplex paths. Using this interpretation, we give a simple exponential upper bound on the number of iterations of the TRUST-REGION IPM.

By symmetry, we restrict ourselves to analyzing the primal central path. Let us analyze $x_i^m(g) := \{\max x_i : x \in \mathcal{P}_g\}$ as $g \geq 0$ increases. Examine the following 2 dimensional projection of the feasible region:

$$\mathcal{P}^i := \{(\langle c, x \rangle - v^*, x_i) : x \in \mathcal{P}\}.$$

Then, by definition, we have that

$$x_i^m(g) := \max\{y : (t, y) \in \mathcal{P}^i, t \leq g\}.$$

Since \mathcal{P} is a convex polyhedron, its shadow \mathcal{P}^i is a two dimensional convex polygon. Clearly $\mathcal{P}^i \subseteq \mathbb{R}_+^2$ since the optimality gap is non-negative and $x_i \geq 0$ for $x \in \mathcal{P}$. Any optimal solution x^* maps to $(\langle c, x^* \rangle - v^*, x_i^*) = (0, x_i^*)$, and similarly any point in $(g, y) \in \mathcal{P}^i$ with $g = 0$ is the image of some optimal solution in \mathcal{P}^i . Therefore, $x_i^m(0) = x_i^*$ where x^* is the optimal solution with largest i th coordinate.

If $u := \sup\{x_i : x \in \mathcal{P}\} < \infty$, we define $\bar{g} < \infty$ to be the first time $\bar{g} \in [0, \infty)$ where $x_i^m(\bar{g}) = u$. In this case, $x_i^m(\cdot)$ will increase on $[0, \bar{g}]$, then will be the constant function u on $[\bar{g}, \infty)$. If $u = \infty$, then $x_i^m(\cdot)$ will be increasing over all \mathbb{R}_+ . In this case, we set $\bar{g} = \infty$. From a geometric perspective, as g increases, $x_i^m(g)$ will first trace the increasing part of the upper convex hull of \mathcal{P}^i and will stay constant afterwards (if $\bar{g} < \infty$).

Using this viewpoint, we conclude that x_i^m is a piecewise linear function with a number of pieces equal to the number of edges on the increasing part of the upper convex hull plus one if $\bar{g} < \infty$. Furthermore, on the interval $[0, \bar{g}]$, x_i^m precisely traces a simplex path on the upper convex hull of \mathcal{P}^i . Furthermore, it can be shown that there is a simplex path in \mathcal{P} whose projection under the map $x \rightarrow (\langle c, x \rangle - v^*, x_i)$ is *precisely* the graph $\{(g, x_i^m(g)) : g \in [0, \bar{g}]\}$ of x_i^m restricted to the interval $[0, \bar{g}]$. By simplex path in this context, we mean the geometric object corresponding to the union of all the edges on the path. One can in fact always choose this path so that the projection is bijective when restricted to the path (furthermore, the path will generically be unique). A path derived in this way is known as a *shadow vertex simplex path*, as it projects to a simplex path in the two dimensional shadow.

From the above discussion, we conclude that up to an additive factor of 1, each of the functions x_1^m, \dots, x_n^m has a number of pieces that is equal to the length of of some shadow vertex simplex path. Since the number of vertices of \mathcal{P} is trivially bounded by $\binom{n}{m}$ (an upper bound on the number of bases of \mathbf{A}), we see that $\sum_{i=1}^n \text{SLC}_1(x_i^m) \leq n \binom{n}{m} \leq n2^n$. Combining Corollary 5 and Theorem 2, with $\beta = 1/6$ and $\eta = 1$, we conclude that the TRUST-REGION IPM requires at most $O(n^{1.5} \log(n)2^n)$ many iterations to solve Equation (LP).

1.2 Duality of the Max Central Path

We recall that for the regular central path we have the identity $x^{\text{CP}}(\mu)s^{\text{CP}}(\mu) = \mu 1_n$. That is, the primal central path determines the dual path and vice versa. One may wonder to which extent a similar phenomenon holds for the MCP (note that Lemma 4 establishes this approximately up to a polynomial factor in n). The following theorem shows that this relation holds approximately with an approximation factor of 2. With this relation, it can in fact be shown (the proof is not entirely trivial) that $\text{SLC}_\eta(x_i^m) = \Theta(\text{SLC}_\eta(s_i^m))$ for any $\eta \in [0, 1/2]$ (the $1/2$ relates to the factor 2 below). That is, that the SLC of the primal and dual are essentially the same.

THEOREM 6 (CENTRALITY OF THE MAX CENTRAL PATH) *For all $g \geq 0$, we have that*

$$g \leq x_i^m(g)s_i^m(g) \leq 2g \quad \forall i \in [n].$$

PROOF: We first prove the upper bound. For $i \in [n]$, let $x^{(i)} \in \text{argmax}\{x_i : x \in \mathcal{P}_g\}$ and $s^{(i)} \in \text{argmax}\{s_i : s \in \mathcal{D}_g\}$. Then,

$$x_i^m(g)s_i^m(g) = x_i^{(i)}s_i^{(i)} \leq \langle x^{(i)}, s^{(i)} \rangle = \langle x^{(i)}, s^* \rangle + \langle s^*, s^{(i)} \rangle \leq 2g,$$

where the last equality follows from Equation (1). We now prove the lower bound.

We assume $g > 0$, since the statement is trivial otherwise. Let e^1, \dots, e^n denote the standard basis vectors in \mathbb{R}^n , i.e., $e_j^i = 1$ if $i = j$ and 0 otherwise. Note that the dual program of $\max\{x_i : \mathbf{A}x = b, x \geq 0, \langle x, s^* \rangle \leq g\}$ can be expressed as

$$\min \left\{ \alpha g + \langle u, x^* \rangle : \alpha s^* + u \geq e^i, \mathbf{A}^\top y = u, \alpha \geq 0 \right\},$$

using that $\langle u, x^* \rangle = \langle y, b \rangle$ since $\mathbf{A}x^* = b$. Similarly, the dual program of $\max\{s_i : \mathbf{A}^\top y + s = c, s \geq 0, \langle s, x^* \rangle \leq g\}$ can be expressed as

$$\min \left\{ \beta g + \langle v, s^* \rangle : \beta x^* + v \geq e^i, \mathbf{A}v = 0_n, \beta \geq 0 \right\}.$$

Let us pick optimal (α, u) and (β, v) to these two programs. The product of the objective values $(\alpha g + \langle s^*, u \rangle)(\beta g + \langle x^*, v \rangle)$ is thus equal to $x_i^m(g)s_i^m(g)$. We complete the proof by showing that this product is lower bounded by g .

We first claim that

$$\langle u, x^* \rangle \geq 0 \quad \text{and} \quad \langle v, s^* \rangle \geq 0. \quad (4)$$

Recalling that $x^* \geq 0$, to show $\langle u, x^* \rangle \geq 0$ it suffices to show that $u_j \geq 0$ whenever $x_j^* > 0$, for all $j \in [n]$. If $x_j^* > 0$, then by complementary slackness $s_j^* = 0$, and thus $u_j = \alpha s_j^* + u_j \geq e_j^i \geq 0$. The proof for $\langle v, s^* \rangle \geq 0$ follows by symmetric reasoning.

Next, note that the constraints in the two programs imply

$$1 = \langle e^i, e^i \rangle \leq \langle \alpha s^* + u, \beta x^* + v \rangle = \alpha \langle v, s^* \rangle + \beta \langle u, x^* \rangle. \quad (5)$$

Now, the product of the objective values can be written as

$$\begin{aligned} x_i^m(g)s_i^m(g) &= (\alpha g + \langle u, x^* \rangle)(\beta g + \langle v, s^* \rangle) \\ &= \alpha\beta g^2 + g(\alpha \langle v, s^* \rangle + \beta \langle u, x^* \rangle) + \langle u, x^* \rangle \cdot \langle v, s^* \rangle \geq g, \end{aligned}$$

using (4) and (5). This concludes the proof. \square

2 Straight line Complexity and Circuits

We now establish an intimate connection between the SLC of an LP and its circuits.

DEFINITION 7 (ELEMENTARY VECTORS AND CIRCUITS) *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$. A vector $z \in \ker(\mathbf{A})$ is an elementary vector in $\ker(\mathbf{A})$ if z is a support-minimal nonzero vector in $\ker(\mathbf{A})$. We let $\mathcal{E}(\mathbf{A})$ denote the set of all elementary vectors. A set $C \subseteq [n]$ is a circuit of \mathbf{A} if it is the support of some elementary vector; we let $\mathcal{C}(\mathbf{A}) \subseteq 2^{[n]}$ denote the set of circuits.*

We say that a vector $y \in \mathbb{R}^n$ conforms to $x \in \mathbb{R}^n$ if $x_i y_i > 0$ whenever $y_i \neq 0$. A conformal circuit decomposition of a vector $z \in \ker(\mathbf{A})$ is a decomposition of the form

$$z = \sum_{i=1}^{\ell} h^{(i)},$$

where $h^{(1)}, \dots, h^{(\ell)} \in \mathcal{E}(\mathbf{A})$, $\ell \leq n$, and each $g^{(i)}$ conforms to z . For this definition to apply for $z = 0_n$, we use the convention that the empty decomposition equals 0_n (i.e., $\ell = 0$). The notion of conformal decomposition can be seen as a generalization of cycle decompositions of circulations of networks flows. The existence of such a decompositions was discovered by [Ful68, Roc69].

PROPOSITION 8 *For every $\mathbf{A} \in \mathbb{R}^{m \times n}$, every vector $z \in \ker(\mathbf{A})$ admits a conformal circuit decomposition.*

PROOF: By possibly multiplying the columns of \mathbf{A} and entry of z by -1 , we may us assume without loss of generality that $z \geq 0_n$. Examine the cone $C := \{x \in \mathbb{R}^n \mid \mathbf{A}x = 0, x \geq 0_n\}$.

By definition $z \in C$. By the Minkowski-Weyl theorem, there exists $h^{(1)}, \dots, h^{(k)} \in C$ such that $z = \sum_{i=1}^k h^i$ and $\mathbb{R}_+ h^i$, $i \in [k]$, are extreme rays of C (i.e., one dimensional faces of C). Note that $k = 0$ if $z = 0_n$ by our convention.

Each extreme ray $\mathbb{R}_+h^{(i)}$, $i \in [k]$, can be represented as $C_{S_i} := \{\mathbf{x} \mid \mathbf{Ax} = 0_n, x_{S_i} = 0, x \geq 0_n\}$, where $S_i \subseteq [n]$ (i.e., we force some of the inequalities to be tight). In particular, we may choose $S_i := [n] \setminus \text{support}(h^{(i)})$ (this is the inclusion-wise maximum set which is consistent with $h^{(i)}$). From here, by our assumption that C_{S_i} is one-dimensional, we must have that $\mathbb{R}h^{(i)} := \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{Ax} = 0_n, z_{S_i} = 0\}$ (otherwise, there would be a perturbation of $h^{(i)}$ that is still feasible for C_{S_i} and not collinear with $h^{(i)}$, contradicting 1-dimensionality). Given our choice $S_i := [n] \setminus \text{support}(h^{(i)})$ for any $S' \supset S_i$, we must then have that $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{Ax} = 0_n, x_{S'} = 0_{S'}\} = \{0_n\}$. Therefore, $h^{(i)} \in \mathcal{E}(\mathbf{A})$ is an elementary vector. Since $h^{(1)}, \dots, h^{(k)} \geq 0_n$, $z \geq 0_n$ and $z = \sum_{i=1}^k h^{(i)}$, we must have that each $h^{(i)}$, $i \in [k]$, conforms to z .

To ensure $k \leq n$, we apply Caratheodory's theorem. Specifically, Caratheodory allows us to ensure that the $h^{(i)}$ s in the decomposition are linearly independent, and hence there must be at most n of them. \square

DEFINITION 9 (*h*-CURVE) Let x^*, s^*, y^* be optimal primal-dual solutions to (LP). Given a vector $h \in \ker(\mathbf{A})$, the *h*-curve from x^* is the function $\bar{x}^h : \mathbb{R}_+ \rightarrow \mathbb{R}_+^n$ that maps $\bar{x}^h(g)$ to $\bar{x} + \alpha h$, for $\alpha \in \mathbb{R}_+$ chosen maximally such that $x^* + \alpha h \geq 0$ and $\langle s^*, \alpha h \rangle \leq g$.

Note that $\bar{x}^h = \bar{x}^{\gamma h}$ for all $\gamma > 0$. Furthermore, by optimality of x^* , if $x^* + \alpha h \geq 0_n$ for $\alpha > 0$ then $\langle s^*, h \rangle \geq 0$. From here, by our assumption that \mathcal{P}_g is bounded, it is easy to see that

$$\bar{x}^h(g) = x^* + \min \left(\frac{g}{\langle s^*, h \rangle}, \min_{j \in \text{supp}(h^-)} \frac{x_j^*}{|h_j|} \right), \quad (6)$$

with the convention that we ignore the first term in the minimum if $\langle s^*, h \rangle \leq 0$ and we ignore the second term if $\text{supp}(h^-) = \emptyset$ (the support of the negative coordinates). Note that the definition of the curve is invariant to the choice s^* , since $\langle s^*, h \rangle = \langle c - \mathbf{A}^\top y^*, h \rangle = \langle c, h \rangle + \langle y^*, \mathbf{A}h \rangle = \langle c, h \rangle$ as $\mathbf{A}h = 0_n$. See Figure 1 for some examples.




Figure 1: Examples of *h*-curves projected on coordinate i . Here, h 1-dominates h' on i with respect to \bar{x} .

The next lemma shows that for every $i \in [n]$ and $g \geq 0$, the i th coordinate of the max-central path at g is upper bounded by a circuit augmentation from an optimal solution, up to a factor n .

LEMMA 10 Let x^* be a primal optimal solution to (LP) and $i \in [n]$. For every $g \geq 0$ where $x_i^m(g) > x_i^*$, there exists an elementary vector $h \in \mathcal{E}(\mathbf{A})$ such that $\langle c, h \rangle \geq 0$, $h_i > 0$, $h_j \geq 0$ whenever $x_j^* = 0$, and $\bar{x}_i^h(g) \geq x_i^m(g)/n$.

PROOF: Let \hat{x} be a primal feasible solution to (LP) such that $\hat{x}_i = x_i^m(g)$. Consider a conformal circuit decomposition of $\hat{x} - x^* = \sum_{j=1}^{\ell} h^{(j)}$ as in Proposition 8. By the conformal decomposition property, $\hat{x} \geq 0$ implies that $x^* + h^{(j)} \geq 0_n$, for all $j \in [n]$, and hence these are all primal feasible. In particular, we must have $\langle c, h^{(j)} \rangle = \langle c, (x^* + h^{(j)}) - x^* \rangle \geq 0$ for all $j \in [\ell]$ because x^* is a primal optimal solution to (LP). Let $k \in \operatorname{argmax}_{j \in [\ell]} h_i^{(j)}$. Then, $h_i^{(k)} > 0$ due to $\hat{x}_i > x_i^*$. Since $\langle c, x^* + h^{(k)} \rangle \leq \langle c, x^* + \sum_{i=1}^{\ell} h^{(i)} \rangle = \langle c, \hat{x} \rangle \leq g$, we obtain

$$\bar{x}_i^{h^{(k)}}(g) \geq x_i^* + h_i^{(k)} \geq x_i^* + \frac{\sum_{j=1}^{\ell} h_i^{(j)}}{\ell} \geq \frac{\hat{x}_i}{\ell} \geq \frac{\hat{x}_i}{n} = \frac{x_i^m(g)}{n}.$$

Note also that $h_j^{(k)} \geq 0$ whenever $x_j^* = 0$ since $x^* + h^{(k)} \geq 0$. \square

DEFINITION 11 (DOMINANCE) Let x^* be a primal optimal solution to (LP). Let $i \in [n]$ and $\alpha \geq 0$. Given vectors $h, h' \in \ker(\mathbf{A})$ where $\langle c, h \rangle, \langle c, h' \rangle \geq 0$, we say that h α -dominates h' on i with respect to x^* if $\bar{x}_i^h \geq \alpha \bar{x}_i^{h'}$. More generally, given sets $S, S' \subseteq W$, we say that S α -dominates S' on i with respect to x^* if $\langle c, h \rangle \geq 0$ for all $h \in S$, and for every $h' \in S'$ with $\langle c, h' \rangle \geq 0$, there exists $h \in S$ such that h α -dominates h' on i with respect to x^* .

DEFINITION 12 (CIRCUIT COVER) Let x^* be a primal optimal solution to (LP). Let $i \in [n]$ and $\alpha \geq 0$. An α -primal circuit cover of i with respect to x^* is a set $S \subseteq \ker(\mathbf{A})$ which α -dominates $\mathcal{E}(\mathbf{A})$ on i with respect to x^* .

The utility of a circuit cover is illustrated by the following lemma. Note that $x_i^m(0)$ is the maximum of the i -th coordinate of an optimal solution.

LEMMA 13 Fix $i \in [n]$, and let x^* be a primal optimal solution to (LP) such that $x_i^* = x_i^m(0)$. If S is a α -primal circuit cover of i with respect to x^* , then $\operatorname{SLC}_{\alpha/n}(x_i^m) \leq |S| + 1$.

PROOF: We may assume that x_i^m is not constant, as otherwise $\operatorname{SLC}_{\alpha/n}(x_i^m) = 1$. Consider the function $\bar{x}_i^S : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ defined by $\bar{x}_i^S(g) := \max_{h \in S} \bar{x}_i^h(g)$. It is piecewise-linear with at most $2|S|$ pieces, and its upper convex envelope has at most $|S| + 1$ pieces (see Figure 2 for an example). So, it is left to show that $\alpha x_i^m/n \leq \bar{x}_i^S \leq x_i^m$. The upper bound is immediate by Definition 9. For the lower bound, let $g > 0$. By Lemma 10, there exists an elementary vector $h' \in \mathcal{E}(\mathbf{A})$ such that $\langle c, h' \rangle \geq 0$ and $\bar{x}_i^{h'}(g) \geq x_i^m(g)/n$. Since S is an α -primal circuit cover of i with respect to x^* , there exists a vector $h \in S$ such that $\langle c, h \rangle \geq 0$ and

$$\bar{x}_i^S(g) \geq \bar{x}_i^h(g) \geq \alpha \bar{x}_i^{h'}(g) \geq \frac{\alpha}{n} x_i^m(g).$$

\square

3 Straight line complexity in terms of the circuit imbalance measure

In this section, we show how the straight-line complexities can be bounded for (LP) in terms of the circuit imbalance $\kappa_{\mathbf{A}}$. Let us start with the definition of $\kappa_{\mathbf{A}}$.




Figure 2: An example of the curve \bar{x}_i^S , where S is a 1-primal circuit cover of i with respect to x^* .

DEFINITION 14 (CIRCUIT IMBALANCES) Let $\mathbf{A} \in \mathbb{R}^{m \times n}$. If $\ker(\mathbf{A}) = \{\mathbf{0}_n\}$ we define the circuit imbalance of \mathbf{A} as $\kappa_{\mathbf{A}} = 1$. Otherwise, we let

$$\kappa_{\mathbf{A}} := \max \left\{ \left| \frac{h_i}{h_j} \right| : h \in \mathcal{E}(\mathbf{A}), i, j \in \text{supp}(h) \right\}.$$

The circuit imbalance measure was first introduced by Vavasis [Vav94], and was first used in the context of interior point methods by [DHN20]. See the excellent recent survey [ENV22] for more details.

Importantly in the context of combinatorial optimization, matrices with bounded circuit imbalance measure provide a real analogue of integer matrices with bounded subdeterminants. In particular, $\kappa_{\mathbf{A}} = 1$ for any totally unimodular matrix \mathbf{A} . We generalize this statement below:

LEMMA 15 Let $\mathbf{A} \in \mathbb{Z}^{m \times n}$ with $\text{rank}(\mathbf{A}) \geq 1$. Let $\Delta \geq 1$ be an upper bound on the absolute value of the determinant of any square submatrix of \mathbf{A} . Then $\kappa_{\mathbf{A}} \leq \Delta$.

PROOF: Let $x \in \mathcal{E}(\mathbf{A})$. We wish to bound $|x_i/x_j|$ for $i, j \in \text{support}(x)$. If $i = j$, the statement is trivial so assume $i \neq j$. Without loss of generality, we may assume by reordering the indices $\text{support}(x) = \{1, \dots, k+1\}$ and that $i = 1$ and $j = k+1$. Furthermore, by replacing x by $-(x/x_k)$, we may assume that $x_k = -1$. From here, our goal is to upper bound $|x_1|$. Notice now that x satisfies:

$$\mathbf{A}_{\bullet, [k]} x_{[k]} = \mathbf{A}_{\bullet, \{k+1\}},$$

where we use the notation $\mathbf{A}_{\bullet, C}$ to denote the submatrix induced by the columns in $C \subseteq [n]$, and $\mathbf{A}_{R, C}$ to denote the submatrix induced by the rows $R \subseteq [m]$ and columns in C . By assumption that x is an elementary vector, we must have $\mathbf{A}_{\bullet, [k]}$ is non-singular. In particular, there exists a subset $R \subseteq [m]$, $|R| = k$, such that $\text{rank}(\mathbf{A}_{R, [k]}) = k$. In particular, $x_{[k]}$ is the unique solution to

$$\mathbf{A}_{R, [k]} x_{[k]} = \mathbf{A}_{R, \{k+1\}},$$

By Cramer's rule, we can now write

$$|x_1| = \left| \frac{\det(\mathbf{A}_{R, \{2, \dots, k+1\}})}{\det(\mathbf{A}_{R, [k]})} \right| \leq \frac{\Delta}{1} = \Delta,$$

where used that $|\det(\mathbf{A}_{R, [k]})| \geq 1$ for any integer non-singular matrix (recall that the determinant of an integer matrix is integral). \square

THEOREM 16 Assume $\mathcal{P}, \mathcal{D} \neq \emptyset$ for an instance of (LP) given by (\mathbf{A}, b, c) . For each $i \in [n]$, we have $\text{SLC}_\eta(x_i^m) \leq \min\{m, n - m\} + 1$ for $\eta = 1/(n^2\kappa_{\mathbf{A}}^2)$. Moreover, with $\beta = 1/6$, the TRUST-REGION IPM requires at most $O(n^{1.5} \min\{m, n - m\} \log(n\kappa_{\mathbf{A}}))$ iterations to solve (LP).

We recall that the TRUST-REGION IPM in Theorem 2 requires a feasible initial point. Unfortunately, in the present context, the homogeneous self-dual initialization does not preserve condition measure κ . More care is thus required to correctly initialize the system. The interested reader may consult [DKN+24] for the details on how to handle initialization in this setting.

Notwithstanding initialization, Theorem 16 provides an algorithm for solving linear programs in a number of arithmetic operations that only depends on the complexity of the constraint matrix (and not on the complexity of b or c). Such a result was first achieved by Tardos [Tar86] for integer matrices, and later by Vavasis and Ye [VY96] for real matrices using a condition measure that is roughly equivalent to $\kappa_{\mathbf{A}}$ (see [DHN20] for a discussion).

Using the tool of straight line complexity, we now give a simple proof of Theorem 16 by constructing a small circuit cover for each coordinate of the max central path. Combining Lemma 13 and Theorem 2, Theorem 16 will follow directly from the circuit cover construction given below.

LEMMA 17 (SLC CIRCUIT COVER) Fix $k \in [n]$ and let x^* be a primal optimal solution to (LP) with $x_k^* = x_k^m(0)$. Then, there exists an $1/(n\kappa_{\mathbf{A}}^2)$ -primal circuit cover S of k with respect to x^* with $|S| \leq \min\{m, n - m\} + 1$.

PROOF: Let us denote $\kappa = \kappa_{\mathbf{A}}$. We can clearly pick the primal optimal solution x^* as in the statement to be a basic solution. Let s^* be any dual basic optimal solution. Thus, $\text{supp}(x^*) \leq m$ and $\text{supp}(s^*) \leq n - m$. By the complementarity of x^* and s^* , we can reorder the index set such that $x_1^* \geq x_2^* \geq \dots \geq x_n^*$ and $s_1^* \leq s_2^* \leq \dots \leq s_n^*$. Let $p := \max\{i \in [n] \mid x_i^* > 0\}$ and $d := \min\{i \in [n] \mid s_i^* > 0\}$. Clearly, $p < d$. By the basic choice of x^* and s^* , $p \leq m$ and $d \geq n - m + 1$.

We now construct the circuit cover S of k . Note that for any circuit $C \in \mathcal{C}(\mathbf{A})$ and any elementary vector $h \in \mathcal{E}(\mathbf{A})$ such that $\langle c, h \rangle \geq 0$ and $\text{supp}(h) = C$, the h -curve \bar{x}^h is the same. Given a circuit $C \in \mathcal{C}'$, we define $h^C \in \mathcal{E}(\mathbf{A})$ as a fixed elementary vector with $\text{supp}(h^C) = C$ normalized such that $h_k^C \in \{0, 1\}$. Further, let

$$\mathcal{C}' := \{C \in \mathcal{C}(\mathbf{A}) \mid h_k^C = 1, h_j^C \geq 0 \forall p < j \leq n\}.$$

Note that if $\mathcal{C}' = \emptyset$ then Lemma 10 implies $x_k^m(g) = x_k^m(0)$ for every $g \geq 0$ and hence $\text{SLC}_1(x_k^m) = 1$. For the rest, we assume $\mathcal{C}' \neq \emptyset$.

We will define the cover S in terms of the support circuits. As the first step, let us define a ‘combinatorial signature’ of circuits. Given a circuit $C \in \mathcal{C}'$, let

$$I^p(C) := \max\{1 \leq i \leq p \mid h_i^C < 0\} \quad \text{and} \quad I^d(C) := \max\{d \leq j \leq n \mid h_j^C > 0\}.$$

We define $I^p(C) = 0$ if the first set is empty and $I^d(C) = 0$ if the second set is empty. We say that a circuit $C \in \mathcal{C}'$ is *dominated* if there exists a circuit $C' \in \mathcal{C}'$ such that

$$I^p(C) \geq I^p(C') \quad \text{and} \quad I^d(C) \geq I^d(C'),$$

and at least one of the two inequalities above is strict. Let $\mathcal{D} \subseteq \mathcal{C}'$ be a maximal collection of undominated circuits with distinct $(I^p(C), I^d(C))$, and define

$$S := \{h^C \mid C \in \mathcal{D}\}.$$

Clearly, $|S| = |\mathcal{D}| \leq \min(p+1, n-d+2) \leq \min\{m, n-m\} + 1$. We show that S is an $1/(n\kappa^2)$ -circuit cover of i with respect to x^* .

To see this, consider any $h \in \mathcal{E}(\mathbf{A})$. If $h_k \leq 0$ or if $h_j < 0$, $p < j \leq n$, then $\bar{x}_k^h(g) \leq x_k^*$, for $g \geq 0$. Hence, \bar{x}^h is 1-dominated by every $h' \in S$. Therefore, we may assume that $h_k > 0$ and $h_{[p+1, n]} \geq 0_{[p+1, n]}$. In particular, by renormalizing so that $h_k = 1$, we may assume that $h \in \mathcal{C}'$.

By definition, there is a $C' \in \mathcal{D}$ such that $I^p(C) \geq I^p(C')$ and $I^d(C) \geq I^d(C')$. For $h = h^C$ and $h' = h^{C'}$, our goal is to show that

$$\bar{x}_k^h(g) \leq n\kappa^2 \bar{x}_k^{h'}(g) \quad \forall g \geq 0. \quad (7)$$

Let $i' = I^p(C') \leq I^p(C) = i$ and $j' = I^d(C') \leq I^d(C) = j$. By the definition of $\kappa = \kappa_{\mathbf{A}}$, we have

$$\frac{1}{\kappa} \leq |h_\ell|, |h'_\ell| \leq \kappa, \quad \forall \ell \in [n]. \quad (8)$$

Using $h_k = h'_k = 1$, we can write

$$\bar{x}_k^h(g) = x_k^* + \min\left(\frac{g}{\langle s^*, h \rangle}, \min_{\ell \in \text{supp}^-(h)} \frac{x_\ell^*}{|h_\ell|}\right) \quad \text{and} \quad \bar{x}_k^{h'}(g) = x_k^* + \min\left(\frac{g}{\langle s^*, h' \rangle}, \min_{\ell \in \text{supp}^-(h')} \frac{x_\ell^*}{|h'_\ell|}\right).$$

We show that

$$\langle s^*, h' \rangle \leq n\kappa^2 \langle s^*, h \rangle \quad \text{and} \quad \min_{\ell \in \text{supp}^-(h)} \frac{x_\ell^*}{|h_\ell|} \leq n\kappa^2 \min_{\ell' \in \text{supp}^-(h')} \frac{x_{\ell'}^*}{|h'_{\ell'}|}. \quad (9)$$

Let us start by showing the first inequality. If $j' = 0$, then $\langle s^*, h' \rangle = 0$ and this trivially holds. Otherwise, by the definition of $j = I^d(C) \geq j' = I^d(C') \geq d$ and using (8), we get

$$\langle s^*, h' \rangle \leq \sum_{\ell=d}^{j'} s_\ell^* h'_\ell \leq n\kappa s_j^* \leq n\kappa^2 s_j^* h_j \leq n\kappa^2 \langle s^*, h \rangle$$

whenever $j' \geq d$; the last inequality follows since $h_\ell \geq 0$ for all $\ell > p$ since $C \in \mathcal{C}'$.

Let us now verify the second inequality in (9). Let ℓ and ℓ' denote the minimizers, respectively. Since $C \in \mathcal{C}'$, it follows that either $i = 0$, that is, $\text{supp}^-(h) = \emptyset$ or $1 \leq \ell \leq i = I^p(C)$; similarly for h' . If $i' = 0$ then the expression for h' is ∞ . Hence, we can assume $1 \leq \ell' \leq i' \leq i$. We get

$$\frac{x_{\ell'}}{|h_{\ell'}|} \geq \frac{x_{\ell'}}{\kappa} \geq \frac{x_i}{\kappa} \geq \frac{1}{\kappa^2} \cdot \frac{x_i}{|h_i|} \geq \frac{1}{\kappa^2} \cdot \frac{x_\ell^*}{|h_\ell|},$$

where the first inequality uses (8); the second inequality uses $\ell' \leq i$ and the ordering of the indices; the third inequality uses again (8); and the last inequality uses the choice of ℓ as the minimizer, noting that $h_i < 0$. \square

References

- [ABGJ18] Xavier Allamigeon, Pascal Benchimol, Stéphane Gaubert, and Michael Joswig. Log-barrier interior point methods are not strongly polynomial. *SIAM Journal on Applied Algebra and Geometry*, 2(1):140–178, 2018.

- [ADL⁺22] Xavier Allamigeon, Daniel Dadush, Georg Loho, Bento Natura, and László A Végh. Interior point methods are not worse than simplex. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 267–277. IEEE, 2022.
- [DHN^V20] Daniel Dadush, Sophie Huiberts, Bento Natura, and László A Végh. A scaling-invariant algorithm for linear programming whose running time depends only on the constraint matrix. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 761–774, 2020.
- [DKN⁺24] Daniel Dadush, Zhuan Khye Koh, Bento Natura, Neil Olver, and László A Végh. A strongly polynomial algorithm for linear programs with at most two nonzero entries per row or column. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1561–1572, 2024.
- [ENV22] Farbod Ekbatani, Bento Natura, and László A Végh. *Circuit Imbalance Measures and Linear Programming*, page 64–114. London Mathematical Society Lecture Note Series. Cambridge University Press, 2022.
- [Ful68] DR Fulkerson. Networks, frames, blocking systems. *Mathematics of the Decision Sciences, Part I, Lectures in Applied Mathematics*, 2:303–334, 1968.
- [Roc69] R. Tyrrell Rockafellar. The elementary vectors of a subspace of R^N . In *Combinatorial Mathematics and Its Applications: Proceedings North Carolina Conference, Chapel Hill, 1967*, pages 104–127. The University of North Carolina Press, 1969.
- [Tar86] Éva Tardos. A strongly polynomial algorithm to solve combinatorial linear programs. *Operations Research*, pages 250–256, 1986.
- [Vav94] Stephen A Vavasis. Stable numerical algorithms for equilibrium systems. *SIAM Journal on Matrix Analysis and Applications*, 15(4):1108–1131, 1994.
- [VY96] Stephen A. Vavasis and Yinyu Ye. A primal-dual interior point method whose running time depends only on the constraint matrix. *Mathematical Programming*, 74(1):79–120, 1996.