



On circuit diameter bounds via circuit imbalances

Daniel Dadush¹ · Zhuan Khye Koh¹ · Bento Natura² · László A. Végh³

Received: 12 August 2022 / Accepted: 3 June 2024 / Published online: 8 July 2024
© The Author(s) 2024

Abstract

We study the circuit diameter of polyhedra, introduced by Borgwardt, Finhold, and Hemmecke (SIAM J. Discrete Math. **29**(1), 113–121 (2015)) as a relaxation of the combinatorial diameter. We show that the circuit diameter of a system $\{x \in \mathbb{R}^n : Ax = b, 0 \leq x \leq u\}$ for $A \in \mathbb{R}^{m \times n}$ is bounded by $O(m \min\{m, n - m\} \log(m + \kappa_A) + n \log n)$, where κ_A is the circuit imbalance measure of the constraint matrix. This yields a strongly polynomial circuit diameter bound if e.g., all entries of A have polynomially bounded encoding length in n . Further, we present circuit augmentation algorithms for LPs using the minimum-ratio circuit cancelling rule. Even though the standard minimum-ratio circuit cancelling algorithm is not finite in general, our variant can solve an LP in $O(mn^2 \log(n + \kappa_A))$ augmentation steps.

Mathematics Subject Classification 52B05 · 90C05

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreements ScaleOpt–757481 and QIP–805241). This work was done while the authors participated in the Discrete Optimization Trimester Program at the Hausdorff Institute for Mathematics in Bonn in 2021. An extended abstract of this paper has appeared in Proceedings of the 23rd Conference on Integer Programming and Combinatorial Optimization, IPCO 2022.

✉ Zhuan Khye Koh
zhuan.koh@cwi.nl

Daniel Dadush
dadush@cwi.nl

Bento Natura
bnatura3@gatech.edu

László A. Végh
l.vegh@lse.ac.uk

¹ Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

² Georgia Institute of Technology, Atlanta, USA

³ London School of Economics and Political Science, London, UK

1 Introduction

The *combinatorial diameter* of a polyhedron P is the diameter of the vertex-edge graph associated with P . Hirsch's famous conjecture from 1957 asserted that the combinatorial diameter of a d -dimensional polytope (bounded polyhedron) with f facets is at most $f - d$. This was disproved by Santos in 2012 [30]. The *polynomial Hirsch conjecture*, i.e., finding a poly(f) bound on the combinatorial diameter remains a central question in the theory of linear programming.

The first quasipolynomial bound was given by Kalai and Kleitman [24, 25], see [32] for the best current bound and an overview of the literature. Dyer and Frieze [11] proved the polynomial Hirsch conjecture for totally unimodular (TU) matrices. For a system $\{x \in \mathbb{R}^d : Mx \leq b\}$ with integer constraint matrix M , polynomial diameter bounds were given in terms of the maximum subdeterminant Δ_M [4, 7, 12, 20]. These arguments can be strengthened to using a parametrization by a 'discrete curvature measure' $\delta_M \geq 1/(d\Delta_M^2)$. The best such bound was given by Dadush and Hähnle [12] as $O(d^3 \log(d/\delta_M)/\delta_M)$, using a shadow vertex simplex algorithm.

As a natural relaxation of the combinatorial diameter, Borgwardt, Finhold, and Hemmecke [5] initiated the study of *circuit diameters*. Consider a polyhedron in standard equality form

$$P = \{x \in \mathbb{R}^n : Ax = b, x \geq 0\} \quad (\text{P})$$

for $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$; we assume $\text{rk}(A) = m$. For the linear space $W = \ker(A) \subseteq \mathbb{R}^n$, $g \in W$ is an *elementary vector* if g is a support-minimal nonzero vector in W , that is, no $h \in W \setminus \{0\}$ exists such that $\text{supp}(h) \subsetneq \text{supp}(g)$. A *circuit* in W is the support of some elementary vector; these are precisely the circuits of the associated linear matroid $\mathcal{M}(A)$. We remark that many papers on circuit diameter, e.g., [2, 3, 5, 8, 26], refer to elementary vectors as circuits; we follow the traditional convention of [21, 27, 29]. We let $\mathcal{E}(W) = \mathcal{E}(A) \subseteq W$ and $\mathcal{C}(W) = \mathcal{C}(A) \subseteq 2^n$ denote the set of elementary vectors and circuits in the space $W = \ker(A)$, respectively. All edge directions of P are elementary vectors, and the set of elementary vectors $\mathcal{E}(A)$ equals the set of all possible edge directions of P in the form (P) for varying $b \in \mathbb{R}^m$ [31].

A *circuit walk* is a sequence of points $x^{(0)}, x^{(1)}, \dots, x^{(k)}$ in P such that for each $i = 0, \dots, k-1$, $x^{(i+1)} = x^{(i)} + \alpha^{(i)}g^{(i)}$ for some $g^{(i)} \in \mathcal{E}(A)$ and $\alpha^{(i)} > 0$, and further, $x^{(i)} + \alpha g^{(i)} \notin P$ for any $\alpha > \alpha^{(i)}$, i.e., each consecutive circuit step is *maximal*. The *circuit diameter* of P is the maximum length (number of steps) of a shortest circuit walk between any two vertices $x, y \in P$. Note that, in contrast to walks in the vertex-edge graph, circuit walks are non-reversible and the minimum length from x to y may be different from the one from y to x ; this is due to the maximal step requirement. The circuit-analogue of Hirsch conjecture, formulated in [5], asserts that the circuit diameter of d -dimensional polyhedron with f facets is at most $f - d$; this may be true even for unbounded polyhedra, see [8]. For P in the form (P), $d = n - m$ and the number of facets is at most n ; hence, the conjectured bound is m .

Circuit diameter bounds have been shown for some combinatorial polytopes such as dual transportation polyhedra [5], matching, travelling salesman, and fractional stable set polytopes [26]. The paper [2] introduced several other variants of circuit diameter, and explored the relation between them. We note that [2, 16, 26] considers circuits for

LPs given in the general form $\{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$. In Sect. 8, we show that this setting can be reduced to the form (P).

Circuit augmentation algorithms Circuit diameter bounds are inherently related to *circuit augmentation algorithms*. This is a general algorithmic scheme to solve an LP

$$\min \langle c, x \rangle \quad \text{s.t.} \quad Ax = b, \quad x \geq 0. \tag{LP}$$

The algorithm proceeds through a sequence of feasible solutions $x^{(t)}$. An initial feasible $x^{(0)}$ is required in the input. For $t = 0, 1, \dots$, the current $x^{(t)}$ is updated to $x^{(t+1)} = x^{(t)} + \alpha g$ for some $g \in \mathcal{E}(A)$ such that $\langle c, g \rangle \leq 0$, and $\alpha > 0$ such that $x^{(t)} + \alpha g$ is feasible. The elementary vector g is an *augmenting direction* if $\langle c, g \rangle < 0$ and such an $\alpha > 0$ exists; by LP duality, $x^{(t)}$ is optimal if and only if no augmenting direction exists. The augmentation is *maximal* if $x^{(t)} + \alpha' g$ is infeasible for any $\alpha' > \alpha$; α is called the maximal stepsize for $x^{(t)}$ and g . Clearly, an upper bound on the number of steps of a circuit augmentation algorithm with maximal augmentations for arbitrary cost c and starting point $x^{(0)}$ yields an upper bound on the circuit diameter.

Simplex is a circuit augmentation algorithm that is restricted to using special elementary vectors corresponding to edges of the polyhedron. Many network optimization algorithms can be seen as special circuit augmentation algorithms. Bland [6] introduced a circuit augmentation algorithm for LP, that generalizes the Edmonds–Karp–Dinic maximum flow algorithm and its analysis, see also [27, Proposition 3.1]. Circuit augmentation algorithms were revisited by De Loera, Hemmecke, and Lee in 2015 [15], analyzing different augmentation rules and also extending them to integer programming. De Loera, Kafer, and Sanità [16] studied the convergence of these rules on 0/1-polytopes, as well as the computational complexity of performing them. We refer the reader to [15] and [16] for a more detailed overview of the background and history of circuit augmentations.

The circuit imbalance measure For a linear space $W = \ker(A) \subseteq \mathbb{R}^n$, the *circuit imbalance* $\kappa_W = \kappa_A$ is defined as the maximum of $|g_j/g_i|$ over all elementary vectors $g \in \mathcal{E}(W)$, $i, j \in \text{supp}(g)$. It can be shown that $\kappa_W = 1$ if and only if W is a unimodular space, i.e., the kernel of a totally unimodular matrix. This parameter and related variants have been used implicitly or explicitly in many areas of linear programming and discrete optimization, see [19] for a recent survey. It is closely related to the Dikin–Stewart–Todd condition number $\bar{\chi}_W$ that plays a key role in layered-least-squares interior point methods introduced by Vavasis and Ye [38]. An LP of the form (LP) for $A \in \mathbb{R}^{m \times n}$ can be solved in time $\text{poly}(n, m, \log \kappa_A)$, which is strongly polynomial if $\kappa_A \leq 2^{\text{poly}(n)}$; see [13, 17] for recent developments and references.

Imbalance and diameter The combinatorial diameter bound $O(d^3 \log(d/\delta_M)/\delta_M)$ from [12] mentioned above translates to a bound $O((n - m)^3 m \kappa_A \log(\kappa_A + n))$ for the system in the form (P), see [19]. For circuit diameters, the Goldberg-Tarjan minimum-mean cycle cancelling algorithm for minimum-cost flows [23] naturally extends to a circuit augmentation algorithm for general LPs using the *steepest-descent* rule. This yields a circuit diameter bound $O(n^2 m \kappa_A \log(\kappa_A + n))$ [19], see also [22]. However, note that these bounds may be exponential in the bit-complexity of the input.

1.1 Our contributions

Our first main contribution improves the κ_A dependence to a $\log \kappa_A$ dependence for circuit diameter bounds.

Theorem 1.1 *The circuit diameter of a system in the form (P) with constraint matrix $A \in \mathbb{R}^{m \times n}$ is $O(m \min\{m, n - m\} \log(m + \kappa_A))$.*

The proof in Sect. 3 is via a simple ‘shoot towards the optimum’ scheme. We need the well-known concept of *conformal circuit decompositions*. We say that $x, y \in \mathbb{R}^n$ are *sign-compatible* if $x_i y_i \geq 0$ for all $i \in [n]$. We write $x \sqsubseteq y$ if they are sign-compatible and further $|x_i| \leq |y_i|$ for all $i \in [n]$. It follows from Carathéodory’s theorem and Minkowski–Weyl theorem that for any linear space $W \subseteq \mathbb{R}^n$ and $x \in W$, there exists a decomposition $x = \sum_{j=1}^k h^{(j)}$ such that $h^{(j)} \in \mathcal{E}(W)$, $h^{(j)} \sqsubseteq x$ for all $j \in [k]$ and $k \leq \dim(W)$. This is called a *conformal circuit decomposition* of x (see also Definition 2.2 and Lemma 2.3 below).

Let $B \subseteq [n]$ be a feasible basis and $N = [n] \setminus B$, i.e., $x^* = (A_B^{-1}b, 0_N) \geq 0_n$ is a basic feasible solution. This is the unique optimal solution to (LP) for the cost function $c = (0_B, \mathbb{1}_N)$. Let $x^{(0)} \in P$ be an arbitrary vertex. We may assume that $n \leq 2m$, by restricting to the union of the support of x^* and $x^{(0)}$, and setting all other variables to 0. For the current iterate $x^{(t)}$, let us consider a conformal circuit decomposition $x^* - x^{(t)} = \sum_{j=1}^k h^{(j)}$. Note that the existence of such a decomposition *does not* yield a circuit diameter bound of n , due to the maximality requirement in the definition of circuit walks. For each $j \in [k]$, $x^{(t)} + h^{(j)} \in P$, but there might be a larger augmentation $x^{(t)} + \alpha h^{(j)} \in P$ for some $\alpha > 1$.

Still, one can use this decomposition to construct a circuit walk. Let us pick the most improving circuit from the decomposition, i.e., the one maximizing $-\langle c, h^{(j)} \rangle = \|h_N^{(j)}\|_1$, and obtain $x^{(t+1)} = x^{(t)} + \alpha^{(t)} h^{(j)}$ for the maximum stepsize $\alpha^{(t)} \geq 1$. The proof of Theorem 1.1 is based on analyzing this procedure. The first key observation is that $\langle c, x^{(t)} \rangle = \|x_N^{(t)}\|_1$ decreases geometrically. Then, we look at the set of indices $L_t = \{i \in [n] : x_i^* > n\kappa_A \|x_N^{(t)}\|_1\}$ and $R_t = \{i \in [n] : x_i^{(t)} \leq (n - m)x_i^*\}$, and show that indices may never leave these sets once they enter. Moreover, a new index is added to either set every $O(m \log(m + \kappa_A))$ iterations. In Sect. 4, we extend this bound to the setting with upper bounds on the variables.

Theorem 1.2 *The circuit diameter of a system in the form $Ax = b$, $0 \leq x \leq u$ with constraint matrix $A \in \mathbb{R}^{m \times n}$ is $O(m \min\{m, n - m\} \log(m + \kappa_A) + (n - m) \log n)$.*

There is a straightforward reduction from the capacitated form to (P) by adding n slack variables; however, this would give an $O(n^2 \log(n + \kappa_A))$ bound. For the stronger bound, we use a preprocessing that involves cancelling circuits in the support of the current solution; this eliminates all but $O(m)$ of the capacity bounds in $O(n \log n)$ iterations, independently of κ_A .

For rational input, $\log(\kappa_A) = O(\text{size}(A))$ where $\text{size}(A)$ denotes the total encoding length of A [13]. Hence, our result yields an $O(m \min\{m, n - m\} \text{size}(A) + n \log n)$ diameter bound on $Ax = b$, $0 \leq x \leq u$. This can be compared with the bounds $O(n \text{size}(A, b))$ using deepest descent augmentation steps in [15, 16], where $\text{size}(A, b)$

is the encoding length of (A, b) . (Such a bound holds for every augmentation rule that decreases the optimality gap geometrically, including the minimum-ratio circuit rule discussed below.) Note that our bound is independent of b . Furthermore, it is also applicable to systems given by irrational inputs, in which case arguments based on subdeterminants and bit-complexity cannot be used.

In light of these results, the next important step towards the polynomial Hirsch conjecture might be to show a $\text{poly}(n, \log \kappa_A)$ bound on the combinatorial diameter of (P) . Note that—in contrast with the circuit diameter—not even a $\text{poly}(n, \text{size}(A, b))$ bound is known. In this context, the best known general bound is $O((n - m)^3 m \kappa_A \log(\kappa_A + n))$ implied by [12].

Circuit augmentation algorithms The diameter bounds in Theorems 1.1 and 1.2 rely on knowing the optimal solution x^* ; thus, they do not provide efficient LP algorithms. We next present circuit augmentation algorithms with $\text{poly}(n, m, \log \kappa_A)$ bounds on the number of iterations. Such algorithms require subroutines for finding augmenting circuits. In many cases, such subroutines are LPs themselves. However, they may be of a simpler form, and might be easier to solve in practice. Borgwardt and Viss [9] exhibit an implementation of a steepest-descent circuit augmentation algorithm with encouraging computational results.

We assume that a subroutine $\text{RATIO-CIRCUIT}(A, c, w)$ is available; this implements the well-known minimum-ratio circuit rule. It takes as input a matrix $A \in \mathbb{R}^{m \times n}$, $c \in \mathbb{R}^n$, $w \in (\mathbb{R}_{++} \cup \{\infty\})^n$, and returns a basic optimal solution to the system

$$\min \langle c, z \rangle \quad \text{s.t.} \quad Az = 0, \langle w, z^- \rangle \leq 1, \tag{1}$$

where $(z^-)_i := \max\{0, -z_i\}$ for $i \in [n]$. Here, we use the convention $w_i z_i = 0$ if $w_i = \infty$ and $z_i = 0$. This system can be equivalently written as an LP using auxiliary variables. If bounded, a basic optimal solution is either 0 or an elementary vector $z \in \mathcal{E}(A)$ that minimizes $\langle c, z \rangle / \langle w, z^- \rangle$.

Given $x \in P$, we use weights $w_i = 1/x_i$ (with $w_i = \infty$ if $x_i = 0$). For minimum-cost flow problems, this rule was proposed by Wallacher [39]; such a cycle can be found in strongly polynomial time for flows. The main advantage of this rule is that the optimality gap decreases by a factor $1 - 1/n$ in every iteration. This rule, along with the same convergence property, can be naturally extended to linear programming [28], and has found several combinatorial applications, e.g., [40, 41], and has also been used in the context of integer programming [33].

On the negative side, Wallacher’s algorithm is *not* strongly polynomial: it does not terminate finitely for minimum-cost flows, as shown in [28]. In contrast, our algorithms achieve a strongly polynomial running time whenever $\kappa_A \leq 2^{\text{poly}(n)}$. An important modification is the occasional use of a second type of circuit augmentation step SUPPORT-CIRCUIT that removes circuits in the support of the current (non-basic) iterate $x^{(t)}$ (see Subroutine 2.1); this can be implemented using simple linear algebra. Our first result addresses the feasibility setting:

Theorem 1.3 *Consider an LP of the form (LP) with cost function $c = (\mathbb{0}_{[n] \setminus N}, \mathbb{1}_N)$ for some $N \subseteq [n]$. There exists a circuit augmentation algorithm that either finds a solution x such that $x_N = 0$ or a dual certificate that no such solution exists, using*

$O(mn \log(n + \kappa_A))$ RATIO- CIRCUIT and $(m + 1)n$ SUPPORT- CIRCUIT augmentation steps.

Such problems typically arise in Phase I of the Simplex method when we add auxiliary variables in order to find a feasible solution. The algorithm is presented in Sect. 6. The analysis extends that of Theorem 1.1, tracking large coordinates $x_i^{(t)}$. Our second result considers general optimization:

Theorem 1.4 *Consider an LP of the form (LP). There exists a circuit augmentation algorithm that finds an optimal solution or concludes unboundedness using $O(mn^2 \log(n + \kappa_A))$ RATIO- CIRCUIT and $(m + 1)n^2$ SUPPORT- CIRCUIT augmentation steps.*

The proof is given in Sect. 7. The main subroutine identifies a new index $i \in [n]$ such that $x_i^{(t)} = 0$ in the current iteration and $x_i^* = 0$ in an optimal solution; we henceforth fix this variable to 0. To derive this conclusion, at the end of each phase the current iterate $x^{(t)}$ will be optimal to (LP) with a slightly modified cost function \tilde{c} ; the conclusion follows using a proximity argument (Theorem 5.4). The overall algorithm repeats this subroutine n times. The subroutine is reminiscent of the feasibility algorithm (Theorem 1.3) with the following main difference: whenever we identify a new ‘large’ coordinate, we slightly perturb the cost function.

Comparison to black-box LP approaches An important milestone towards strongly polynomial linear programming was Tardos’s 1986 paper [35] on solving (LP) in time $\text{poly}(n, m, \log \Delta_A)$, where Δ_A is the maximum subdeterminant of A . Her algorithm makes $O(nm)$ calls to a weakly polynomial LP solver for instances with small integer capacities and costs, and uses proximity arguments to gradually learn the support of an optimal solution. This approach was extended to the real model of computation for a $\text{poly}(n, m, \log \kappa_A)$ bound [17]. The latter result uses proximity arguments with circuit imbalances κ_A , and eliminates all dependence on bit-complexity.

The proximity tool Theorem 5.4 derives from [17], and our circuit augmentation algorithms are inspired by the feasibility and optimization algorithms in this paper. However, using circuit augmentation oracles instead of an approximate LP oracle changes the setup. Our arguments become simpler since we proceed through a sequence of feasible solutions, whereas much effort in [17] is needed to deal with infeasibility of the solutions returned by the approximate solver. On the other hand, we need to be more careful as all steps must be implemented using circuit augmentations in the original system, in contrast to the higher degree of freedom in [17] where we can make approximate solver calls to arbitrary modified versions of the input LP.

Organization of the paper The rest of the paper is organized as follows. We first provide the necessary preliminaries in Sect. 2. In Sect. 3, we upper bound the circuit diameter of (P). In Sect. 4, this bound is extended to the setting with upper bounds on the variables. Then, we develop circuit-augmentation algorithms for solving (LP). We first present the required proximity results in Sect. 5, Sect. 6 contains the algorithm for finding a feasible solution, whereas Sect. 7 contains the algorithm for solving (LP) given an initial feasible solution. Section 8 shows how circuits in LPs of more general forms can be reduced to the notion used in this paper.

2 Preliminaries

Let $[n] = \{1, 2, \dots, n\}$. Let \mathbb{R}_+ and \mathbb{R}_{++} be the set of nonnegative and positive real numbers respectively. For $\alpha \in \mathbb{R}$, we denote $\alpha^+ = \max\{0, \alpha\}$ and $\alpha^- = \max\{0, -\alpha\}$. For a vector $z \in \mathbb{R}^n$, we define $z^+, z^- \in \mathbb{R}^n$ as $(z^+)_i = (z_i)^+, (z^-)_i = (z_i)^-$ for $i \in [n]$. For $z \in \mathbb{R}^n$, we let $\text{supp}(z) = \{i \in [n] : z_i \neq 0\}$ denote its support, and $1/z \in (\mathbb{R} \cup \{\infty\})^n$ denote the vector $(1/z_i)_{i \in [n]}$. We use $\|\cdot\|_p$ to denote the ℓ_p -norm; we simply write $\|\cdot\|$ for $\|\cdot\|_2$. For $A \in \mathbb{R}^{m \times n}$ and $S \subseteq [n]$, we let $A_S \in \mathbb{R}^{m \times |S|}$ denote the submatrix corresponding to columns S . We denote $\text{rk}(S) := \text{rk}(A_S)$, i.e., the rank of the set S in the linear matroid associated with A . A *spanning subset of S* is a subset $T \subseteq S$ such that $\text{rk}(T) = \text{rk}(S)$. The *closure* of S is defined as $\text{cl}(S) := \{i \in [n] : \text{rk}(S \cup \{i\}) = \text{rk}(S)\}$. The dual linear program of (LP) is

$$\max \langle b, y \rangle \quad \text{s.t.} \quad A^T y + s = c, \quad s \geq 0. \tag{DLP}$$

Note that y uniquely determines s , and due to the assumption $\text{rk}(A) = m$, s also uniquely determines y . For this reason, given a dual feasible solution (y, s) , we can just focus on y or s .

For $A \in \mathbb{R}^{m \times n}$, let $W = \ker(A)$. Recall that $\mathcal{C}(W) = \mathcal{C}(A)$ and $\mathcal{E}(W) = \mathcal{E}(A)$ are the set of circuits and elementary vectors in W respectively. Note that every circuit has size at most $m + 1$ because we assumed that $\text{rk}(A) = m$. The *circuit imbalance measure* of W is defined as

$$\kappa_W := \kappa_A := \max_{g \in \mathcal{E}(W)} \left\{ \frac{|g_i|}{|g_j|} : i, j \in \text{supp}(g) \right\}$$

if $W \neq \{0\}$. Otherwise, it is defined as $\kappa_W := \kappa_A := 1$. For a linear space $W \subseteq \mathbb{R}^n$, let W^\perp denote the orthogonal complement. Thus, for $W = \ker(A)$, $W^\perp = \text{Im}(A^T)$. According to the next lemma, circuit imbalances are self-dual.

Lemma 2.1 ([13]) *For a linear space $W \subseteq \mathbb{R}^n$, we have $\kappa_W = \kappa_{W^\perp}$.*

For P as in (P), $x \in P$ and an elementary vector $g \in \mathcal{E}(A) \setminus \mathbb{R}_+^n$, we let $\text{aug}_P(x, g) := x + \alpha g$ where $\alpha = \max\{\bar{\alpha} : x + \bar{\alpha} g \in P\}$.

Definition 2.2 [14] We say that $x, y \in \mathbb{R}^n$ are *sign-compatible* if $x_i y_i \geq 0$ for all $i \in [n]$. We write $x \sqsubseteq y$ if they are sign-compatible and further $|x_i| \leq |y_i|$ for all $i \in [n]$. For a linear space $W \subseteq \mathbb{R}^n$ and $x \in W$, a *conformal circuit decomposition* of x is a set of elementary vectors $h^{(1)}, h^{(2)}, \dots, h^{(k)}$ in W such that $x = \sum_{j=1}^k h^{(j)}$, $k \leq \dim(W)$, and $h^{(j)} \sqsubseteq x$ for all $j \in [k]$.

The following lemma shows that every vector in a linear space has a conformal circuit decomposition. It is a simple corollary of the Minkowski–Weyl and Carathéodory theorems.

Lemma 2.3 *For a linear space $W \subseteq \mathbb{R}^n$, every $x \in W$ has a conformal circuit decomposition $x = \sum_{j=1}^k h^{(j)}$ such that $k \leq \min\{\dim(W), |\text{supp}(x)|\}$.*

2.1 Circuit oracles

In Sects. 4, 6, and 7, we use a simple circuit finding subroutine $\text{SUPPORT-CIRCUIT}(A, c, x, S)$ that will be used to identify circuits in the support of a solution x . This can be implemented easily using Gaussian elimination. Note that the constraint $\langle c, z \rangle \leq 0$ is superficial as $-z$ is also an elementary vector for every elementary vector z .

Subroutine 2.1. $\text{SUPPORT-CIRCUIT}(A, c, x, S)$

For a matrix $A \in \mathbb{R}^{m \times n}$, vectors $c, x \in \mathbb{R}^n$ and $S \subseteq [n]$, the output is an elementary vector $z \in \mathcal{E}(A)$ with $\text{supp}(z) \subseteq \text{supp}(x)$, $\text{supp}(z) \cap S \neq \emptyset$ with $\langle c, z \rangle \leq 0$, or concludes that no such elementary vector exists.

The circuit augmentation algorithms in Sects. 6 and 7 will use the subroutine $\text{RATIO-CIRCUIT}(A, c, w)$.

Subroutine 2.2. $\text{RATIO-CIRCUIT}(A, c, w)$

For a matrix $A \in \mathbb{R}^{m \times n}$ and vectors $c \in \mathbb{R}^n$, $w \in (\mathbb{R}_{++} \cup \{\infty\})^n$, the output is a basic optimal solution to the system:

$$\min \langle c, z \rangle \quad \text{s.t.} \quad Az = 0, \langle w, z^- \rangle \leq 1, \tag{2}$$

and an optimal solution to the following dual program:

$$\max -\lambda \quad \text{s.t.} \quad A^\top y + s = c \quad 0 \leq s \leq \lambda w \tag{3}$$

Note that (2) can be reformulated as an LP using additional variables, and its dual LP can be equivalently written as (3). Recall that we use the convention $w_i z_i = 0$ if $w_i = \infty$ and $z_i = 0$ in (2). The opposite convention is used in (3), i.e., $\lambda_i w_i = \infty$ if $\lambda = 0$ and $w_i = \infty$. If (2) is bounded, then a basic optimal solution is either 0 or an elementary vector $z \in \mathcal{E}(A)$ that minimizes $\langle c, z \rangle / \langle w, z^- \rangle$. Moreover, observe that every feasible solution to (3) is also feasible to (DLP).

We will use the following lemma, a direct consequence of [18, Lemma 4.3].

Lemma 2.4 *Given $A \in \mathbb{R}^{m \times n}$, $W = \ker(A)$, $\ell \in (\mathbb{R} \cup \{-\infty\})^n$ and $u \in (\mathbb{R} \cup \{\infty\})^n$, let $r \in W$ such that $\ell \leq r \leq u$. In $\text{poly}(m, n)$ time, we can find a vector $r' \in W$ such that $\ell \leq r' \leq u$ and $\|r'\|_\infty \leq \kappa_A \|\ell^+ + u^-\|_1$.*

This lemma, together with Lemma 2.1, allows us to assume that the optimal dual solution s returned by RATIO-CIRCUIT satisfies

$$\|s\|_\infty \leq 2\kappa_A \|c\|_1. \tag{4}$$

To see this, let (y, s, λ) be an optimal solution to (3). We know that $-c \leq s - c \leq \lambda w - c$. Let $\ell := -c, r := s - c$ and $u := \lambda w - c$. By Lemma 2.4, we can in $\text{poly}(m, n)$ time compute $r' \in W^\perp$ such that $\ell \leq r' \leq u$ and

$$\|r'\|_\infty \leq \kappa_{W^\perp} \|\ell^+ + u^-\|_1 \leq \kappa_{W^\perp} \|c^- + c^+\|_1 = \kappa_{W^\perp} \|c\|_1.$$

Then, $s' := r' + c$ is an optimal solution to (3) which satisfies

$$\|s'\|_\infty \leq \|r'\|_\infty + \|c\|_\infty \leq (\kappa_{W^\perp} + 1)\|c\|_1 \leq 2\kappa_{W^\perp} \|c\|_1.$$

Thus, (4) follows using Lemma 2.1, since $\kappa_{W^\perp} = \kappa_W = \kappa_A$.

The following lemma is well-known, see e.g., [28, Lemma 2.2].

Lemma 2.5 *Let OPT be the optimal value of (LP), and assume that it is finite. Given a feasible solution x to (LP), let g be the optimal solution to (2) returned by RATIO-CIRCUIT($A, c, 1/x$).*

- (i) *If $\langle c, g \rangle = 0$, then x is optimal to (LP).*
- (ii) *If $\langle c, g \rangle < 0$, then letting $x' = \text{aug}_P(x, g)$, we have $\alpha \geq 1$ for the augmentation stepsize and*

$$\langle c, x' \rangle - \text{OPT} \leq \left(1 - \frac{1}{|\text{supp}(x)|}\right) (\langle c, x \rangle - \text{OPT}).$$

Proof We only prove (ii) because (i) is trivial. The stepsize bound $\alpha \geq 1$ follows since $\langle 1/x, g^- \rangle \leq 1$; thus, $x + g \in P$. Let x^* be an optimal solution to (LP), and let $z = (x^* - x)/|\text{supp}(x)|$. Note that $g \not\geq 0$, as otherwise (2) is unbounded. So, $x \neq 0$. Then, z is feasible to (2) for $w = 1/x$. Therefore,

$$\alpha \langle c, g \rangle \leq \langle c, g \rangle \leq \langle c, z \rangle = \frac{\text{OPT} - \langle c, x \rangle}{|\text{supp}(x)|},$$

implying the second claim. □

Remark 2.6 It is worth noting that Lemma 2.5 shows that applying RATIO-CIRCUIT to vectors x with small support gives better convergence guarantees. Algorithms 3 and 4 for feasibility and optimization in Sects. 6 and 7 apply RATIO-CIRCUIT to vectors x which have large support $|\text{supp}(x)| = \Theta(n)$ in general. These algorithms could be reformulated in that one first runs SUPPORT-CIRCUIT to reduce the size of the support to size $O(m)$ and only then runs RATIO-CIRCUIT. The guarantees of Lemma 2.5 now imply that to reduce the optimality gap by a constant factor we would replace $O(n)$ calls to RATIO-CIRCUIT with only $O(m)$ calls. On the other hand, this comes at the cost of n additional calls to SUPPORT-CIRCUIT for every call to RATIO-CIRCUIT.

2.2 A norm bound

We now formulate a proximity bound asserting that if the columns of A outside N are linearly independent, then we can bound the ℓ_∞ -norm of any vector in $\ker(A)$ by the norm of its coordinates in N . This can be seen as a special case of Hoffman-proximity results; see Sect. 5 for more such results and references.

Lemma 2.7 For $A \in \mathbb{R}^{m \times n}$, let $N \subseteq [n]$ such that $A_{[n] \setminus N}$ has full column rank. Then, for any $z \in \ker(A)$, we have $\|z\|_\infty \leq \kappa_A \|z_N\|_1$.

Proof Let $h^{(1)}, \dots, h^{(k)}$ be a conformal circuit decomposition of z . Then, $\|z\|_\infty \leq \sum_{t=1}^k \|h^{(t)}\|_\infty$. For each $h^{(t)}$, we have $\text{supp}(h^{(t)}) \cap N \neq \emptyset$ because $A_{[n] \setminus N}$ has full column rank. Hence, $\|h^{(t)}\|_\infty \leq \kappa_A |h_{j(t)}^{(t)}|$ for some $j(t) \in N$. Conformality implies that

$$\sum_{t=1}^k |h_{j(t)}^{(t)}| = \sum_{s \in N} \sum_{j(t)=s} |h_{j(t)}^{(t)}| \leq \sum_{s \in N} |z_s| = \|z_N\|_1.$$

The lemma follows by combining all the previous inequalities. □

2.3 Estimating circuit imbalances

The circuit augmentation algorithms in Sects. 6 and 7 explicitly use the circuit imbalance measure κ_A . However, this is NP-hard to approximate within a factor $2^{O(n)}$, see [13, 36]. We circumvent this problem using a standard guessing procedure, see e.g., [13, 38]. Instead of κ_A , we use an estimate $\hat{\kappa}$, initialized as $\hat{\kappa} = n$. Running the algorithm with this estimate either finds the desired feasible or optimal solution (which one can verify), or fails. In case of failure, we conclude that $\hat{\kappa} < \kappa_A$, and replace $\hat{\kappa}$ by $\hat{\kappa}^2$. Since the running time of the algorithms is linear in $\log(n + \hat{\kappa})$, the running time of all runs will be dominated by the last run, giving the desired bound. For simplicity, the algorithm descriptions use the explicit value κ_A .

3 The circuit diameter bound

In this section, we show Theorem 1.1, namely the bound $O(m \min\{m, n - m\} \log(m + \kappa_A))$ on the circuit diameter of a polyhedron in standard form (P). As outlined in the Introduction, let $B \subseteq [n]$ be a feasible basis and $N = [n] \setminus B$ such that $x^* = (A_B^{-1}b, 0_N)$ is a basic solution to (LP). We can assume $n \leq 2m$: the union of the supports of the starting vertex $x^{(0)}$ and the target vertex x^* is at most $2m$; we can fix all other variables to 0. Defining $\tilde{n} := |\text{supp}(x^*) \cup \text{supp}(x^{(0)})| \leq 2m$ and restricting A to these columns, we show a circuit diameter bound $O(\tilde{n}(\tilde{n} - m) \log(m + \kappa_A))$. This implies Theorem 1.1 for general n . In the rest of this section, we use n instead of \tilde{n} , but assume $n \leq 2m$. The simple ‘shoot towards the optimum’ procedure is shown in Algorithm 1.

A priori, even finite termination is not clear. First, we show that the ‘cost’ $\|x_N^{(t)}\|_1$ decreases geometrically. It is a consequence of choosing the most improving circuit $g^{(t)}$ in each iteration.

Lemma 3.1 For every iteration $t \geq 0$, we have $\|x_N^{(t+1)}\|_1 \leq (1 - \frac{1}{n-m}) \|x_N^{(t)}\|_1$. Furthermore, $|x_i^{(t+1)} - x_i^{(t)}| \leq (n - m) |x_i^* - x_i^{(t)}|$ for all $i \in [n]$.

Algorithm 1: DIAMETER- BOUND

Input : Polyhedron in standard form (P) , basis $B \subseteq [n]$ with its corresponding vertex $x^* = (A_B^{-1}b, 0_N)$, and initial vertex $x^{(0)}$.

Output: Length of a circuit walk from $x^{(0)}$ to x^* .

```

1  $t \leftarrow 0$ ;
2 while  $x^{(t)} \neq x^*$  do
3   Let  $h^{(1)}, h^{(2)}, \dots, h^{(k)}$  be a conformal circuit decomposition of  $x^* - x^{(t)}$ ;
4    $g^{(t)} \leftarrow h^{(j)}$  for any  $j \in \arg \max_{i \in [k]} \|h_N^{(i)}\|_1$ ;
5    $x^{(t+1)} \leftarrow \text{aug}_P(x^{(t)}, g^{(t)})$ ;  $t \leftarrow t + 1$ ;
6 return  $t$ ;
    
```

Proof Let $h^{(1)}, \dots, h^{(k)}$ with $k \leq n - m$ be the conformal circuit decomposition of $x^* - x^{(t)}$ used in iteration t of Algorithm 1. Note that $h_N^{(i)} \leq 0_N$ for all $i \in [k]$ because $x_N^* = 0_N$ and $x^{(t)} \geq 0$. By our choice of $g^{(t)}$,

$$\|g_N^{(t)}\|_1 = \max_{i \in [k]} \|h_N^{(i)}\|_1 \geq \frac{1}{k} \sum_{i \in [k]} \|h_N^{(i)}\|_1 = \frac{1}{k} \|x_N^{(t)}\|_1$$

where the last equality uses the conformality of the decomposition. Let $\alpha^{(t)}$ be such that $x^{(t+1)} = x^{(t)} + \alpha^{(t)}g^{(t)}$. Clearly, $\alpha^{(t)} \geq 1$ because $x^{(t)} + g^{(t)} \in P$. Hence,

$$\begin{aligned} \|x_N^{(t+1)}\|_1 &= \|x_N^{(t)} + \alpha^{(t)}g_N^{(t)}\|_1 \leq \|x_N^{(t)} + g_N^{(t)}\|_1 \\ &= \|x_N^{(t)}\|_1 - \|g_N^{(t)}\|_1 \leq \left(1 - \frac{1}{k}\right) \|x_N^{(t)}\|_1. \end{aligned}$$

Further, using $0 \leq x_N^{(t+1)} \leq x_N^{(t)}$, we see that

$$\alpha^{(t)} = \frac{\|x_N^{(t+1)} - x_N^{(t)}\|_1}{\|g_N^{(t)}\|_1} \leq \frac{\|x_N^{(t)}\|_1}{\|g_N^{(t)}\|_1} \leq k,$$

and so for all $i \in [n]$ we have $|x_i^{(t+1)} - x_i^{(t)}| = \alpha^{(t)}|g_i^{(t)}| \leq k|g_i^{(t)}| \leq k|x_i^* - x_i^{(t)}|$. \square

Our convergence proof is based on analyzing the following sets

$$\begin{aligned} L_t &:= \{i \in [n] : x_i^* > n\kappa_A \|x_N^{(t)}\|_1\}, & T_t &:= [n] \setminus L_t, \\ R_t &:= \{i \in [n] : x_i^{(t)} \leq (n - m)x_i^*\}. \end{aligned}$$

The set L_t consists of indices i where x_i^* is much larger than the current ‘cost’ $\|x_N^{(t)}\|_1$. On the other hand, the set R_t consists of indices i where $x_i^{(t)}$ is not much above x_i^* . The next lemma shows that the sets L_t and R_t are monotonically growing.

Lemma 3.2 *For every iteration $t \geq 0$, we have $L_t \subseteq L_{t+1} \subseteq B$ and $R_t \subseteq R_{t+1}$.*

Proof Clearly, $L_t \subseteq L_{t+1}$ as $\|x_N^{(t)}\|_1$ is monotonically decreasing by Lemma 3.1, and $L_t \subseteq B$ as $x_N^* = \mathbb{0}_N$. Next, let $j \in R_t$. If $x_j^{(t)} \geq x_j^*$, then $x_j^{(t+1)} \leq x_j^{(t)}$ by conformality. If $x_j^{(t)} < x_j^*$, then $x_j^{(t+1)} \leq x_j^{(t)} + (n - m)(x_j^* - x_j^{(t)}) \leq (n - m)x_j^*$ by Lemma 3.1. In both cases, we conclude that $j \in R_{t+1}$. \square

Our goal is to show that R_t or L_t is extended within $O((n - m) \log(n + \kappa_A))$ iterations. By the maximality of the augmentation, we know that at least one variable is set to zero in every iteration t . The following lemma shows that these variables do not lie in L_t .

Lemma 3.3 *For every iteration $t \geq 0$, we have $\emptyset \neq \text{supp}(x^{(t)}) \setminus \text{supp}(x^{(t+1)}) \subseteq T_t$.*

Proof Let $i \in \text{supp}(x^{(t)}) \setminus \text{supp}(x^{(t+1)})$. Such a variable exists by the maximality of the augmentation. Clearly, $x_i^{(t+1)} = 0$. Applying Lemma 2.7 to $x^{(t+1)} - x^* \in \ker(A)$ yields

$$x_i^* \leq \|x^{(t+1)} - x^*\|_\infty \leq \kappa_A \|x_N^{(t+1)} - x_N^*\|_1 = \kappa_A \|x_N^{(t+1)}\|_1 \leq \kappa_A \|x_N^{(t)}\|_1.$$

The equality is due to $x_N^* = \mathbb{0}$, while the last inequality follows from Lemma 3.1. So, $i \in T_t$. \square

Clearly, any variable i that is set to zero in iteration t belongs to R_{t+1} . So, if $i \notin R_t$, then we make progress as $R_t \subsetneq R_{t+1}$. Note that this is always the case if $i \in N$. We show that if $\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty$ is sufficiently large, then $i \notin R_t$.

Lemma 3.4 *If $\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty > 2mn^2\kappa_A^2 \left\|x_{T_t}^*\right\|_\infty$ for some iteration t , then $R_t \subsetneq R_{t+1}$.*

Proof Let $i \in \text{supp}(x^{(t)}) \setminus \text{supp}(x^{(t+1)})$. Clearly, $i \in R_{t+1}$ because $x_i^{(t+1)} = 0$. So, it suffices to show that $i \notin R_t$. Since $x^{(t+1)} - x^{(t)}$ is an elementary vector, we have $\|x^{(t+1)} - x^{(t)}\|_\infty \leq \kappa_A |x_i^{(t+1)} - x_i^{(t)}| = \kappa_A x_i^{(t)}$. As $|\text{supp}(x^{(t+1)} - x^{(t)})| \leq m + 1$, we obtain

$$\|x_N^{(t)} - x_N^{(t+1)}\|_1 \leq (m + 1) \|x_i^{(t)} - x_i^{(t+1)}\|_\infty \leq (m + 1) \kappa_A x_i^{(t)} \leq 2m\kappa_A x_i^{(t)}. \tag{5}$$

Let $h^{(1)}, \dots, h^{(k)}$ with $k \leq n - m$ be the conformal circuit decomposition of $x^* - x^{(t)}$ used in iteration t of Algorithm 1. Let $j \in T_t$ such that $|x_j^{(t)} - x_j^*| = \|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty$. There exists $\tilde{h} = h^{(\ell)}$ for some $\ell \in [k]$ in this decomposition such that $|\tilde{h}_j| \geq \frac{1}{k} |x_j^{(t)} - x_j^*|$. Since A_B has full column rank, we have $\text{supp}(\tilde{h}) \cap N \neq \emptyset$ and so

$$\|\tilde{h}_N\|_1 \geq \frac{|\tilde{h}_j|}{\kappa_A} \geq \frac{|x_j^{(t)} - x_j^*|}{k\kappa_A}. \tag{6}$$

From (5), (6) and noting that $\|\tilde{h}_N\|_1 \leq \|g_N^{(t)}\|_1 \leq \|x_N^{(t)} - x_N^{(t+1)}\|_1$ by our choice of $g^{(t)}$, we get

$$x_i^{(t)} \geq \frac{\|x_N^{(t)} - x_N^{(t+1)}\|_1}{2m\kappa_A} \geq \frac{\|\tilde{h}_N\|_1}{2m\kappa_A} \geq \frac{\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty}{2mk\kappa_A^2}.$$

Thus, if $\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty > 2mn^2\kappa_A^2\|x_{T_t}^*\|_\infty$ as in the assumption of the lemma, then $x_i^{(t)} > n\|x_{T_t}^*\|_\infty \geq nx_i^*$, where the last inequality is due to $i \in T_t$ by Lemma 3.3. It follows that $i \notin R_t$ as desired. \square

We are ready to give the convergence bound. We have just proved that a large $\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty$ guarantees the extension of R_t . Using the geometric decay of $\|x_N^{(t)}\|$ (Lemma 3.1), we now show that if $\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty$ is small, then $\|x_N^{(t)}\|_1$ drops sufficiently such that a new variable enters L_t .

Proof of Theorem 1.1 Recall that we assumed $n \leq 2m$ without loss of generality. In light of Lemma 3.2, it suffices to show that either L_t or R_t is extended in every $O((n - m) \log(n + \kappa_A))$ iterations. If $\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty > 2mn^2\kappa_A^2\|x_{T_t}^*\|_\infty$, then $R_t \subsetneq R_{t+1}$ by Lemma 3.4.

So, let us assume that $\|x_{T_t}^{(t)} - x_{T_t}^*\|_\infty \leq 2mn^2\kappa_A^2\|x_{T_t}^*\|_\infty$, that is, $\|x_{T_t}^{(t)}\|_\infty \leq (2mn^2\kappa_A^2 + 1)\|x_{T_t}^*\|_\infty$. We may also assume that $\|x_N^{(t)}\|_1 > 0$, as otherwise $x^{(t)} = x^*$. By Lemma 3.1, there is an iteration $r = t + O((n - m) \log(n + \kappa_A))$ such that $n^2\kappa_A(2mn^2\kappa_A^2 + 1)\|x_N^{(r)}\|_1 < \|x_N^{(t)}\|_1$. Hence,

$$\begin{aligned} (2mn^2\kappa_A^2 + 1)\|x_{T_t}^*\|_\infty &\geq \|x_{T_t}^{(t)}\|_\infty \geq \|x_N^{(t)}\|_\infty \\ &\geq \frac{1}{n}\|x_N^{(t)}\|_1 > n\kappa_A(2mn^2\kappa_A^2 + 1)\|x_N^{(r)}\|_1, \end{aligned}$$

where the second inequality is due to $N \subseteq T_t$ by Lemma 3.2. Thus, $\|x_{T_t}^*\|_\infty > n\kappa_A\|x_N^{(r)}\|_1$ and so $L_t \subsetneq L_r$. \square

4 Circuit diameter bound for the capacitated case

In this section we consider diameter bounds for systems of the form

$$P_u = \{x \in \mathbb{R}^n : Ax = b, 0 \leq x \leq u\}. \tag{Cap-P}$$

The theory in Sect. 3 carries over to P_u at the cost of turning m into n via the standard reformulation

$$\tilde{P}_u = \left\{ (x, y) \in \mathbb{R}^{n+n} : \begin{bmatrix} A & 0 \\ I & I \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b \\ u \end{bmatrix}, x, y \geq 0 \right\}, \quad P_u = \{x : (x, y) \in \tilde{P}_u\}. \tag{7}$$

Corollary 4.1 *The circuit diameter of a system in the form (Cap-P) with constraint matrix $A \in \mathbb{R}^{m \times n}$ is $O(n^2 \log(n + \kappa_A))$.*

Proof Follows straightforward from Theorem 1.1 together with the reformulation (7). Let \tilde{A} denote the constraint matrix of (7). It is easy to check that $\kappa_A = \kappa_{\tilde{A}}$, and that there is a one-to-one mapping between the circuits and maximal circuit augmentations of the two systems. □

Intuitively, the polyhedron should not become more complex; related theory in [37] also shows how two-sided bounds can be incorporated in a linear program without significantly changing the complexity of solving the program.

Theorem 1.2 is proved using a new procedure, which we outline below. A basic feasible point $x^* \in P_u$ is characterised by a partition $B \cup L \cup H = [n]$ where A_B is a basis (has full column rank), $x_L^* = 0_L$ and $x_H^* = u_H$. In $O(n \log n)$ iterations, we fix all but $2m$ variables to the same bound as in x^* ; for the remaining system with $2m$ variables, we can use the standard reformulation.

Algorithm 2 starts with a preprocessing. We let $S_t \subseteq L \cup H$ denote the set of indices where $x_i^{(t)} \neq x_i^*$, i.e., we are not yet at the required lower and upper bound. If $|S_t| \leq m$, then we remove the indices in $(L \cup H) \setminus S_t$, and use the diameter bound resulting from the standard embedding as in Corollary 4.1.

As long as $|S_t| > m$, we proceed as follows. We define the cost function $c \in \mathbb{R}^n$ by $c_i = 0$ for $i \in B$, $c_i = 1/u_i$ for $i \in L$, and $c_i = -1/u_i$ for $i \in H$. For this choice, we see that the optimal solution of the LP $\min_{x \in P_u} \langle c, x \rangle$ is x^* with optimal value $\langle c, x^* \rangle = -|H|$.

Depending on the value of $\langle c, x^{(t)} \rangle$, we perform one of two updates. As long as $\langle c, x^{(t)} \rangle \geq -|H| + 1$, we take a conformal decomposition of $x^* - x^{(t)}$, and pick the most improving augmenting direction from the decomposition. If $\langle c, x^{(t)} \rangle < -|H| + 1$, then we use a support circuit augmentation obtained from SUPPORT-CIRCUIT($A, c, x^{(t)}, S_t$).

Let us show that whenever SUPPORT-CIRCUIT is called, $g^{(t)}$ is guaranteed to exist. This is because $|S_t| > m$ and $x_i^{(t)} > 0$ for all $i \in S_t$. Indeed, if $x_j^{(t)} = 0$ for some $j \in S_t$, then $j \in H$ from the definition of S_t . However, this implies that

$$\langle c, x^{(t)} \rangle \geq \sum_{i \in H \setminus \{j\}} c_i x_i^{(t)} \geq -|H| + 1,$$

which is a contradiction.

The cost $\langle c, x^{(t)} \rangle$ is monotone decreasing, and it is easy to see that $\langle c, x^{(0)} \rangle \leq n$ for any initial solution $x^{(0)}$. Hence, within $O((n - m) \log n)$ iterations we must reach $\langle c, x^{(t)} \rangle < -|H| + 1$. Each support circuit augmentation sets $x_i^{(t+1)} = 0$ for $i \in L$ or $x_i^{(t+1)} = u_i$ for $i \in H$; hence, we perform at most $n - m$ such augmentations. The formal proof is given below.

Proof of Theorem 1.2 We show that Algorithm 2 has the claimed number of iterations. As previously mentioned, $\langle c, x^* \rangle = -|H|$ is the optimal value of the LP $\min_{x \in P_u} \langle c, x \rangle$. Initially, $\langle c, x^{(0)} \rangle = -\sum_{i \in H} \frac{x^{(0)}}{u_i} + \sum_{i \in L} \frac{x^{(0)}}{u_i} \leq n$. Similar to

Algorithm 2: CAPACITATED- DIAMETER- BOUND

Input : Polyhedron in the form (Cap-P), partition $B \cup L \cup H = [n]$ with its corresponding vertex $x^* = (A_B^{-1}b, 0_L, u_H)$, and initial vertex $x^{(0)}$.

Output: Length of a circuit walk from $x^{(0)}$ to x^* .

- 1 Set the cost $c \in \mathbb{R}^n$ as $c_i = 0$ if $i \in B$, $c_i = 1/u_i$ if $i \in L$, and $c_i = -1/u_i$ if $i \in H$;
- 2 $t \leftarrow 0$;
- 3 $S_0 \leftarrow \{i \in L \cup H : x_i^{(0)} \neq x_i^*\}$;
- 4 **while** $|S_t| > m$ **do**
- 5 **if** $\langle c, x^{(t)} \rangle \geq -|H| + 1$ **then**
- 6 Let $h^{(1)}, h^{(2)}, \dots, h^{(k)}$ be a conformal circuit decomposition of $x^* - x^{(t)}$;
- 7 $g^{(t)} \leftarrow h^{(j)}$ for any $j \in \arg \min_{i \in [k]} \langle c, h^{(i)} \rangle$;
- 8 **else**
- 9 $g^{(t)} \leftarrow \text{SUPPORT- CIRCUIT}(A, c, x^{(t)}, S_t)$
- 10 $x^{(t+1)} \leftarrow \text{aug}_P(x^{(t)}, g^{(t)})$;
- 11 $S_{t+1} \leftarrow \{i \in L \cup H : x_i^{(t+1)} \neq x_i^*\}$; $t \leftarrow t + 1$;
- 12 Run Algorithm 1 on $\tilde{A} := \begin{bmatrix} A_{B \cup S_t} & 0 \\ I & I \end{bmatrix}$ and $\tilde{b} = \begin{bmatrix} b \\ u \end{bmatrix}$ to get $t' \in \mathbb{Z}_+$;
- 13 **return** $t + t'$;

Lemma 3.1, due to our choice of $g^{(t)}$ from the conformal circuit decomposition, we have $\langle c, x^{(t+1)} \rangle + |H| \leq (1 - \frac{1}{n-m})(\langle c, x^{(t)} \rangle + |H|)$. In particular, $O((n - m) \log n)$ iterations suffice to find an iterate t such that $\langle c, x^{(t)} \rangle < -|H| + 1$.

Note that the calls to SUPPORT- CIRCUIT do not increase $\langle c, x^{(t)} \rangle$, so from now we will never make use of the conformal circuit decomposition again. An augmentation resulting from a call to SUPPORT- CIRCUIT will set at least one variable $i \in \text{supp}(g^{(t)})$ to either 0 or u_i . We claim that either $x_i^{(t+1)} = 0$ for some $i \in L$, or $x_i^{(t+1)} = u_i$ for some $i \in H$, that is, we set a variable to the ‘correct’ boundary. To see this, note that if $x_i^{(t+1)}$ hits the wrong boundary, then the gap between $\langle c, x^{(t+1)} \rangle$ and $-|H|$ must be at least 1, a clear contradiction to $\langle c, x^{(t+1)} \rangle < -|H| + 1$.

Thus, after at most $n - m$ calls to SUPPORT- CIRCUIT, we get $|S_t| \leq m$, at which point we call Algorithm 1 with at most $2m$ variables, so the diameter bound of Theorem 1.1 applies. □

5 Proximity results

We now present Hoffman-proximity bounds in terms of the circuit imbalance measure κ_A . A simple such bound was Lemma 2.7; we now present additional norm bounds. These can be derived from more general results in [17]; see also [19]. The references also explain the background and similar results in previous literature, in particular, to proximity bounds via Δ_A in e.g., [35] and [10]. For completeness, we include the proofs.

The next technical lemma will be key in our arguments. See Corollary 5.2 below for a simple implication.

Lemma 5.1 *Let $A \in \mathbb{R}^{m \times n}$ and $x \in \mathbb{R}^n$. Let $L \subseteq \text{supp}(x)$ and $S \subseteq [n] \setminus L$. If there is no circuit $C \subseteq \text{supp}(x)$ such that $C \cap S \neq \emptyset$, then*

$$\|x_S\|_\infty \leq \kappa_A \min_{z \in \ker(A)+x} \|z_{[n] \setminus \text{cl}(L)}\|_1.$$

Before the proof, it is worth stating a useful special case $L = \emptyset$ and $S = [n]$.

Corollary 5.2 *Let x be a basic (but not necessarily feasible) solution to (LP). Then, for any z where $Az = b$, we have $\|x\|_\infty \leq \kappa_A \|z\|_1$.*

Proof of Lemma 5.1 First, we show that $x_{S \cap \text{cl}(L)} = 0$ due to our assumption. Indeed, any $i \in S \cap \text{cl}(L)$ with $x_i \neq 0$ gives rise to a circuit in $L \cup \{i\} \subseteq \text{supp}(x)$, contradicting the assumption in the lemma. It follows that $\|x_S\|_\infty = \|x_{S \setminus \text{cl}(L)}\|_\infty$; let $j \in S \setminus \text{cl}(L)$ such that $|x_j| = \|x_S\|_\infty$. Let $z \in \ker(A)+x$ be a minimizer of the RHS in the statement. We may assume that $|x_j| > |z_j|$, as otherwise we are done because $\kappa_A \geq 1$.

Let $h^{(1)}, \dots, h^{(k)}$ be a conformal circuit decomposition of $z - x \in \ker(A)$. Among these elementary vectors, consider the set $R := \{t \in [k] : h_j^{(t)} \neq 0\}$.

Claim 5.3 *For each $t \in R$, there exists an index $i(t) \in \text{supp}(h^{(t)}) \setminus \text{cl}(L)$ such that $x_{i(t)} = 0$ and $z_{i(t)} \neq 0$.*

Proof For the purpose of contradiction, suppose that $\text{supp}(h^{(t)}) \setminus \text{cl}(L) \subseteq \text{supp}(x)$. For every $i \in \text{cl}(L) \setminus L$, we can write $A_i = Ay^{(i)}$ where $\text{supp}(y^{(i)}) \subseteq L$. Consider the vector

$$h := h^{(t)} + \sum_{i \in \text{cl}(L) \setminus L} h_i^{(t)}(y^{(i)} - e_i).$$

Clearly, $h_{\text{cl}(L) \setminus L} = 0$ and $h_{[n] \setminus \text{cl}(L)} = h_{[n] \setminus \text{cl}(L)}^{(t)}$. Since $L \subseteq \text{supp}(x)$ and we assumed $\text{supp}(h^{(t)}) \setminus \text{cl}(L) \subseteq \text{supp}(x)$, it follows that $\text{supp}(h) \subseteq \text{supp}(x)$. Moreover, $j \in \text{supp}(h)$ because $j \in S \setminus \text{cl}(L)$. Hence, applying Lemma 2.3 to $h \in \ker(A)$ yields an elementary vector $g \in \mathcal{E}(A)$ such that $\text{supp}(g) \subseteq \text{supp}(x)$ and $\text{supp}(g) \cap S \neq \emptyset$. This contradicts the assumption of the lemma. \square

By conformality of the decomposition, $|x_j - z_j| = \sum_{t \in R} |h_j^{(t)}|$. According to Claim 5.3, for every $t \in R$, we have $|h_j^{(t)}| \leq \kappa_A |h_{i(t)}^{(t)}|$ where $i(t) \in [n] \setminus (\text{cl}(L) \cup \{j\})$; notice that $i(t) \neq j$ for all $t \in R$ due to our assumption $|x_j| > 0$. Applying conformality again yields

$$\sum_{i \in R} |h_{i(t)}^{(t)}| = \sum_{s \in [n] \setminus (\text{cl}(L) \cup \{j\})} \sum_{i(t)=s} |h_{i(t)}^{(t)}| \leq \sum_{s \in [n] \setminus (\text{cl}(L) \cup \{j\})} |z_s| = \|z_{[n] \setminus (\text{cl}(L) \cup \{j\})}\|_1.$$

Therefore,

$$\|x_S\|_\infty = |x_j| \leq |z_j| + |x_j - z_j| \leq \kappa_A \|z_{[n] \setminus \text{cl}(L)}\|_1$$

where the last inequality is obtained by combining the previous equation and inequalities. □

The following proximity theorem will be key to derive $x_i^* = 0$ for certain variables in our optimization algorithm; see [17] and [19, Theorem 6.5]. For $\tilde{c} \in \mathbb{R}^n$, we use $LP(\tilde{c})$ to denote (LP) with cost vector \tilde{c} , and $OPT(\tilde{c})$ as the optimal value of $LP(\tilde{c})$.

Theorem 5.4 *Let $c, c' \in \mathbb{R}^n$ be two cost vectors, such that both $LP(c)$ and $LP(c')$ have finite optimal values. Let s' be a dual optimal solution to $LP(c')$. For all indices $j \in [n]$ such that*

$$s'_j > (m + 1)\kappa_A \|c - c'\|_\infty,$$

it follows that $x_j^ = 0$ for every optimal solution x^* to $LP(c)$.*

Proof We may assume that $c \neq c'$, as otherwise we are done by complementary slackness. Let x' be an optimal solution to $LP(c')$. By complementary slackness, $s'_j x'_j = 0$, and therefore $x'_j = 0$. For the purpose of contradiction, suppose that there exists an optimal solution x^* to $LP(c)$ such that $x_j^* > 0$. Let $h^{(1)}, \dots, h^{(k)}$ be a conformal circuit decomposition of $x^* - x'$. Then, $h_j^{(t)} > 0$ for some $t \in [k]$. Since $h^{(t)}$ is an elementary vector, $|\text{supp}(h^{(t)})| \leq m + 1$ and so $\|h^{(t)}\|_1 \leq (m + 1)\|h^{(t)}\|_\infty \leq (m + 1)\kappa_A h_j^{(t)}$. Observe that for any $i \in [n]$ where $h_i^{(t)} < 0$, we have $s'_i = 0$ because $x'_i > x_i^* \geq 0$. Hence,

$$\begin{aligned} \langle c, h^{(t)} \rangle &= \langle c - c', h^{(t)} \rangle + \langle c', h^{(t)} \rangle \geq -\|c - c'\|_\infty \|h^{(t)}\|_1 + \langle s', h^{(t)} \rangle \\ &\geq -(m + 1)\kappa_A \|c - c'\|_\infty h_j^{(t)} + s'_j h_j^{(t)} > 0. \end{aligned}$$

The first inequality here used Hölder’s inequality and that $\langle c', h^{(t)} \rangle = \langle s', h^{(t)} \rangle$ since $c' - s' \in \text{Im}(A^\top)$ and $h^{(t)} \in \text{ker}(A)$. Since $x^* - h^{(t)}$ is feasible to $LP(c)$, this contradicts the optimality of x^* . □

The following lemma provides an upper bound on the norm of the perturbation $c - c'$ for which the existence of an index j as in Theorem 5.4 is guaranteed.

Lemma 5.5 *Let $c, c' \in \mathbb{R}^n$ be two cost vectors, and let s' be an optimal dual solution to $LP(c')$. If $c \in \text{ker}(A)$, $\|c\|_2 = 1$ and $\|c - c'\|_\infty < 1/(\sqrt{n}(m + 2)\kappa_A)$, then there exists an index $j \in [n]$ such that*

$$s'_j > \frac{m + 1}{\sqrt{n}(m + 2)}.$$

Proof Let $r = c - c'$. Note that $s' + r \in \text{Im}(A^\top) + c$. Then,

$$\|s'\|_\infty + \|r\|_\infty \geq \|s' + r\|_\infty \geq \frac{1}{\sqrt{n}} \|s' + r\|_2 \geq \frac{1}{\sqrt{n}} \|c\|_2 = \frac{1}{\sqrt{n}},$$

where the last inequality is due to $s' + r - c$ and c being orthogonal. This gives us

$$\|s'\|_\infty \geq \frac{1}{\sqrt{n}} - \|r\|_\infty > \frac{(m + 2)\kappa_A - 1}{\sqrt{n}(m + 2)\kappa_A} \geq \frac{m + 1}{\sqrt{n}(m + 2)}$$

as desired because $\kappa_A \geq 1$. □

6 A circuit augmentation algorithm for feasibility

In this section we prove Theorem 1.3: given a linear program (LP) with cost $c = (\mathbb{0}_{[n] \setminus N}, \mathbb{1}_N)$ for some $N \subseteq [n]$, find a solution x with $x_N = \mathbb{0}$ (showing that the optimum value is 0), or certify that no such solution exists. A dual certificate in the latter case is a vector $y \in \mathbb{R}^m$ such that $A^\top y \leq c$ and $\langle b, y \rangle > 0$.

Theorem 1.3 can be used to solve the feasibility problem for linear programs. Given a polyhedron in standard form (P), we construct an auxiliary linear program whose feasibility problem is trivial, and whose optimal solutions correspond to feasible solutions to (P). This is in the same tune as Phase I of the Simplex method:

$$\min \langle \mathbb{1}, z \rangle \quad \text{s.t.} \quad Ay - Az = b, \quad y, z \geq \mathbb{0}. \tag{Aux-LP}$$

For the constraint matrix $\tilde{A} = [A \ -A]$, it is easy to see that $\kappa_{\tilde{A}} = \kappa_A$ and that any solution $Ax = b$ can be converted into a feasible solution to (Aux-LP) via $(y, z) = (x^+, x^-)$. Hence, if the subroutines SUPPORT- CIRCUIT and RATIO- CIRCUIT are available for (Aux-LP), then we can invoke Theorem 1.3 with $N = \{n + 1, n + 2, \dots, 2n\}$ on (Aux-LP) to solve the feasibility problem of (P) in $O(mn \log(n + \kappa_A))$ augmentation steps.

Our algorithm is presented in Algorithm 3. We maintain a set $\mathcal{L}_t \subseteq [n] \setminus N$, initialized as \emptyset . Whenever $x_i^{(t)} \geq 4mn\kappa_A \|x_N^{(t)}\|_1$ for the current iterate $x^{(t)}$, we add i to \mathcal{L}_t . Note that once an index i enters \mathcal{L}_t , it is never removed, even though x_i might drop below this threshold in the future. Still, we will show that $\mathcal{L}_t \subseteq \text{supp}(x^{(t)})$ in every iteration.

Whenever $\text{rk}(\mathcal{L}_t)$ increases, we run SUPPORT- CIRCUIT($A, c, x^{(t)}, N$) iterations as long as there exists a circuit in $\text{supp}(x^{(t)})$ intersecting N . Afterwards, we run a sequence of RATIO- CIRCUIT iterations until $\text{rk}(\mathcal{L}_t)$ increases again. The key part of the analysis is to show that $\text{rk}(\mathcal{L}_t)$ increases in every $O(n \log(n + \kappa_A))$ iterations.

Let us first analyze what happens during RATIO- CIRCUIT iterations.

Lemma 6.1 *If RATIO- CIRCUIT is called in iteration t , then either $\|x_N^{(t+1)}\|_1 \leq (1 - \frac{1}{n}) \|x_N^{(t)}\|_1$, or the algorithm terminates with a dual certificate.*

Proof The oracle returns $g^{(t)}$ that is optimal to (2) and $(y^{(t)}, s^{(t)})$ that is optimal to (3) with optimum value $-\lambda$. Thus, $A^\top y + s = c$ and $\mathbb{0} \leq s \leq \lambda w$. Recall that we use weights $w_i = 1/x_i^{(t)}$. If $\langle b, y^{(t)} \rangle > 0$, the algorithm terminates. Otherwise, note that

Algorithm 3: FEASIBILITY- ALGORITHM

Input : Linear program in standard form (LP) with cost $c = (0_{[n] \setminus N}, \mathbb{1}_N)$ for some $N \subseteq [n]$, and initial feasible solution $x^{(0)}$.

Output: A solution x with $x_N = 0$, or a dual solution $y \in \mathbb{R}^m, A^\top y \leq c, \langle b, y \rangle > 0$.

```

1  $t \leftarrow 0; \mathcal{L}_{t-1} \leftarrow \emptyset;$ 
2 while  $x_N^{(t)} \neq 0$  do
3    $\mathcal{L}_t \leftarrow \mathcal{L}_{t-1} \cup \{i \in [n] : x_i^{(t)} \geq 4mn\kappa_A \|x_N^{(t)}\|_1\};$ 
4   if  $t = 0$  or  $\text{rk}(\mathcal{L}_t) > \text{rk}(\mathcal{L}_{t-1})$  then
5     while  $\exists$  a circuit in  $\text{supp}(x^{(t)})$  intersecting  $N$  do
6        $g^{(t)} \leftarrow \text{SUPPORT-CIRCUIT}(A, c, x^{(t)}, N);$ 
7        $x^{(t+1)} \leftarrow \text{aug}_P(x^{(t)}, g^{(t)});$ 
8        $\mathcal{L}_{t+1} \leftarrow \mathcal{L}_t; t \leftarrow t + 1;$ 
9     if  $x_N^{(t)} = 0$  then
10      return  $x^{(t)};$ 
11    $(g^{(t)}, y^{(t)}, s^{(t)}) \leftarrow \text{RATIO-CIRCUIT}(A, c, 1/x^{(t)});$ 
12   if  $\langle b, y^{(t)} \rangle > 0$  then
13     Terminate with infeasibility certificate ;
14    $x^{(t+1)} \leftarrow \text{aug}_P(x^{(t)}, g^{(t)}); t \leftarrow t + 1;$ 
15 return  $x^{(t)};$ 

```

$$\langle c, x^{(t)} \rangle = \langle b, y^{(t)} \rangle + \langle s^{(t)}, x^{(t)} \rangle \leq \lambda \langle w_{\text{supp}(x^{(t)})}, x_{\text{supp}(x^{(t)})}^{(t)} \rangle \leq n\lambda,$$

implying $\lambda \geq \langle c, x^{(t)} \rangle / n$, and therefore $\langle c, g^{(t)} \rangle = -\lambda \leq -\langle c, x^{(t)} \rangle / n$. This implies the lemma, noting that

$$\|x_N^{(t+1)}\|_1 = \langle c, x^{(t+1)} \rangle \leq \langle c, x^{(t)} \rangle + \langle c, g^{(t)} \rangle \leq \left(1 - \frac{1}{n}\right) \|x_N^{(t)}\|_1.$$

□

Next, we analyze what happens during SUPPORT- CIRCUIT iterations.

Lemma 6.2 *If SUPPORT- CIRCUIT is called in iteration t , then $\|x^{(t+1)} - x^{(t)}\|_\infty \leq \kappa_A \|x_N^{(t)}\|_1$.*

Proof We have $g_i^{(t)} < 0$ for some $i \in N$ because $\text{supp}(g^{(t)}) \cap N \neq \emptyset$ and $\langle c, g^{(t)} \rangle \leq 0$. Hence,

$$\|x^{(t+1)} - x^{(t)}\|_\infty \leq \kappa_A |x_i^{(t+1)} - x_i^{(t)}| \leq \kappa_A x_i^{(t)} \leq \kappa_A \|x_N^{(t)}\|_1.$$

□

The following lemma shows that once a coordinate enters \mathcal{L}_t , its value stays above a certain threshold.

Lemma 6.3 For every iteration $t \geq 0$, we have $x_j^{(t)} \geq 2mn\kappa_A \|x_N^{(t)}\|_1$ for all $j \in \mathcal{L}_t$.

Proof Fix an iteration $t \geq 0$ and a coordinate $j \in \mathcal{L}_t$. We may assume that $\|x_N^{(t)}\|_1 > 0$, as otherwise the lemma trivially holds because $x^{(t)} \geq 0$. Let $r \leq t$ be the iteration in which j was added to \mathcal{L}_r ; the lemma clearly holds at iteration r .

We analyze the ratio $x_j^{(t')}/\|x_N^{(t')}\|_1$ for iterations $t' = r, \dots, t$. At an iteration $r \leq t' < t$ that performs RATIO-CIRCUIT, observe that if $x_j^{(t')}/\|x_N^{(t')}\|_1 \geq 2n\kappa_A$, then

$$\begin{aligned} \frac{x_j^{(t'+1)}}{\|x_N^{(t'+1)}\|_1} &\geq \frac{x_j^{(t')} - \kappa_A \|x_N^{(t'+1)} - x_N^{(t')}\|_1}{(1 - \frac{1}{n}) \|x_N^{(t')}\|_1} \geq \frac{x_j^{(t')} - 2\kappa_A \|x_N^{(t')}\|_1}{(1 - \frac{1}{n}) \|x_N^{(t')}\|_1} \\ &\geq \frac{(1 - \frac{1}{n})x_j^{(t')}}{(1 - \frac{1}{n}) \|x_N^{(t')}\|_1} = \frac{x_j^{(t')}}{\|x_N^{(t')}\|_1}. \end{aligned}$$

The first inequality is due to Lemma 6.1 and the fact that $x^{(t'+1)} - x^{(t')}$ is an elementary vector whose support intersects N . This fact follows from $\langle c, g^{(t')} \rangle < 0$ because $\|x_N^{(t')}\|_1 \geq \|x_N^{(t)}\|_1 > 0$ and $\langle b, y^{(t')} \rangle \leq 0$. The second inequality uses the monotonicity $\|x_N^{(t'+1)}\|_1 \leq \|x_N^{(t')}\|_1$ and the triangle inequality. The third inequality uses the assumption $x_j^{(t')}/\|x_N^{(t')}\|_1 \geq 2n\kappa_A$.

Hence, it suffices to show that SUPPORT-CIRCUIT maintains the invariant $x_j^{(t')}/\|x_N^{(t')}\|_1 \geq 2n\kappa_A$. At an iteration $r \leq t' < t$ which performs SUPPORT-CIRCUIT, we have

$$\frac{x_j^{(t'+1)}}{\|x_N^{(t'+1)}\|_1} \geq \frac{x_j^{(t')} - \kappa_A \|x_N^{(t')}\|_1}{\|x_N^{(t')}\|_1} = \frac{x_j^{(t')}}{\|x_N^{(t')}\|_1} - \kappa_A$$

by Lemma 6.2. Since Algorithm 3 performs at most $(m + 1)n$ SUPPORT-CIRCUIT iterations, the total decrease of this ratio is at most $(m + 1)n\kappa_A \leq 2mn\kappa_A$. As the starting value is at least $4mn\kappa_A$, it follows that this ratio does not drop below $2mn\kappa_A$. \square

Proof of Theorem 1.3 The correctness of Algorithm 3 is obvious. If the algorithm terminates due to $x_N^{(t)} = 0$, then $x^{(t)}$ is the desired solution to (LP). Otherwise, if the algorithm terminates due to $\langle b, y^{(t)} \rangle > 0$, then $y^{(t)}$ is the desired dual certificate as it is feasible to (DLP).

Next, we show that if $\text{rk}(\mathcal{L}_t) = m$, then the algorithm will terminate in iteration $r \leq t + n$ with $x_N^{(r)} = 0$. As long as $x_N^{(t)} \neq 0$, we have $\mathcal{L}_t \subseteq [n] \setminus N$ by Lemma 6.3. Moreover, any $i \in \text{supp}(x_N^{(t)})$ induces a circuit in $\mathcal{L}_t \cup \{i\}$, so SUPPORT-CIRCUIT will be invoked. Since every call to SUPPORT-CIRCUIT reduces $\text{supp}(x^{(t)})$, all the coordinates in N will be zeroed-out in at most n calls.

It is left to bound the number of iterations of Algorithm 3. In the first iteration and whenever $\text{rk}(\mathcal{L}_t)$ increases, we perform a sequence of at most n SUPPORT CIRCUIT cancellations. Let us consider an iteration t right after we are done with the SUPPORT CIRCUIT cancellations. Then, there is no circuit in $\text{supp}(x^{(t)})$ intersecting N . We show that $\text{rk}(\mathcal{L}_t)$ increases within $O(n \log(n + \kappa_A))$ consecutive calls to RATIO-CIRCUIT; this completes the proof.

By Lemma 6.1, within $O(n \log(n\kappa_A)) = O(n \log(n + \kappa_A))$ consecutive RATIO-CIRCUIT augmentations, we reach an iterate $r = t + O(n \log(n + \kappa_A))$ such that $\|x_N^{(r)}\|_1 \leq (4mn^3\kappa_A^2)^{-1} \|x_N^{(t)}\|_1$. Since $\mathcal{L}_t \subseteq \text{supp}(x^{(t)})$ and $N \subseteq [n] \setminus \mathcal{L}_t$ by Lemma 6.3, and there is no circuit in $\text{supp}(x^{(t)})$ intersecting N , applying Lemma 5.1 with $x = x^{(t)}$ and $z = x^{(r)}$ yields

$$\|x_{[n] \setminus \text{cl}(\mathcal{L}_t)}^{(r)}\|_\infty \geq \frac{\|x_{[n] \setminus \text{cl}(\mathcal{L}_t)}^{(r)}\|_1}{n} \geq \frac{\|x_N^{(t)}\|_\infty}{n\kappa_A} \geq \frac{\|x_N^{(t)}\|_1}{n^2\kappa_A} \geq 4mn\kappa_A \|x_N^{(r)}\|_1,$$

showing that some $j \in [n] \setminus \text{cl}(\mathcal{L}_t)$ must be included in \mathcal{L}_r . □

7 A circuit augmentation algorithm for optimization

In this section, we give a circuit-augmentation algorithm for solving (LP), given by $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$. We also assume that an initial feasible solution $x^{(0)}$ is provided. In every iteration t , the algorithm maintains a feasible solution $x^{(t)}$ to (LP), initialized with $x^{(0)}$. The goal is to augment $x^{(t)}$ using the subroutines SUPPORT-CIRCUIT and RATIO-CIRCUIT until the emergence of a nonempty set $N \subseteq [n]$ which satisfies $x_N^{(t)} = x_N^* = 0$ for every optimal solution x^* to (LP). When this happens, we have reached a lower dimensional face of the polyhedron that contains the optimal face. Hence, we can fix $x_N^{(t')} = 0$ in all subsequent iterations $t' \geq t$. In particular, we repeat the same procedure on a smaller LP with constraint matrix $A_{[n] \setminus N}$, RHS vector b , and cost $c_{[n] \setminus N}$, initialized with the feasible solution $x_{[n] \setminus N}^{(t)}$. Note that a circuit walk of this smaller LP corresponds to a circuit walk of the original LP. This gives the overall circuit-augmentation algorithm.

In what follows, we focus on the aforementioned variable fixing procedure (Algorithm 4), since the main algorithm just calls it at most n times.

We fix parameters

$$\delta := \frac{1}{2n^{3/2}(m+2)\kappa_A}, \quad T := \Theta(n \log(n + \kappa_A)), \quad \Gamma := \frac{6(m+2)\sqrt{n}\kappa_A^2 T}{\delta}.$$

Throughout the procedure, A and b will be fixed, but we will sometimes modify the cost function c . Recall that for any $\tilde{c} \in \mathbb{R}^n$, we use $\text{LP}(\tilde{c})$ to denote the problem with cost vector \tilde{c} , and the optimal value is $\text{OPT}(\tilde{c})$. We will often use the fact that if $\tilde{s} \in \text{Im}(A^\top) + \tilde{c}$, then the linear programs $\text{LP}(\tilde{s})$ and $\text{LP}(\tilde{c})$ are equivalent.

Let us start with a high level overview before presenting the algorithm. The inference that $x_N^{(t)} = x_N^* = 0$ for every optimal x^* will be made using Theorem 5.4. To apply

this, our goal is to find a cost function c' and an optimal dual solution s' to $\text{LP}(c')$ such that the set of indices $N := \{j : s'_j > (m+1)\kappa_A \|c - c'\|_\infty\}$ is nonempty.

If $c = 0$, then we can return $x^{(0)}$ as an optimal solution. Otherwise, we can normalize to $\|c\| = 1$.¹ Let us start from any primal and dual feasible solutions $(x^{(0)}, s^{(0)})$ to $\text{LP}(c)$; we can obtain $s^{(0)}$ from a call to **RATIO-CIRCUIT**. Within $O(n \log(n + \kappa_A))$ **RATIO-CIRCUIT** augmentations, we arrive at a pair of primal and dual feasible solutions $(x, s) = (x^{(t)}, s^{(t)})$ such that $\langle x, s \rangle \leq \varepsilon := \langle x^{(0)}, s^{(0)} \rangle / \text{poly}(n, \kappa_A)$.

We now describe the high level motivation for the algorithm. Suppose that for every $i \in \text{supp}(x)$, s_i is small, say $s_i < \delta$. Let $\tilde{c}_i := s_i$ if $i \notin \text{supp}(x)$ and $\tilde{c}_i := 0$ if $i \in \text{supp}(x)$. Then, $\|\tilde{c} - s\|_\infty < \delta$ and x and \tilde{c} are primal and dual optimal solutions to $\text{LP}(\tilde{c})$. This follows because they are primal and dual feasible and satisfy complementary slackness. Consider the vector $c' := c - s + \tilde{c}$, which satisfies $\|c - c'\|_\infty < \delta$. Since $c - s \in \text{Im}(A^\top)$, $\text{LP}(c')$ and $\text{LP}(\tilde{c})$ are equivalent. Thus, x and \tilde{c} are primal and dual optimal solutions to $\text{LP}(c')$. Then, Theorem 5.4 is applicable for the costs c, c' and the dual optimal solution \tilde{c} . However, to be able to make progress by fixing variables, we also need to guarantee that $N \neq \emptyset$. Following Tardos [34, 35], this can be ensured if we pre-process by projecting the cost vector c onto $\ker(A)$; this guarantees that $\|s\|$ —and thus $\|\tilde{c}\|$ —must be sufficiently large.

Let us now turn to the case when the above property does not hold for (x, s) : for certain coordinates we could have $x_i > 0$ and $s_i \geq \delta$. We enter the second *phase* of the algorithm. Let $S = \{i \in [n] : s_i \geq \delta\}$ be the coordinates with large dual slack. Since $x_i s_i \leq \langle x, s \rangle \leq \varepsilon$, this implies $x_i \leq \varepsilon/\delta$ for all $i \in S$. Therefore, $\|x_S\|$ is sufficiently small, and one can show that the set of ‘large’ indices $\mathcal{L} = \{i \in [n] : x_i \geq \Gamma \|x_S\|_1\}$ is nonempty. We proceed by defining a new cost function $\tilde{c}_i := s_i$ if $i \in S$ and $\tilde{c}_i := 0$ if $i \notin S$. We perform **SUPPORT-CIRCUIT** iterations as long as there exist circuits in $\text{supp}(x)$ intersecting $\text{supp}(\tilde{c})$, and then perform further $O(n \log(n + \kappa_A))$ **RATIO-CIRCUIT** iterations for the cost function \tilde{c} . If we now arrive at an iterate $(x, s) = (x^{(t')}, s^{(t')})$ such that $s_i < \delta$ for every $i \in \text{supp}(x)$, then we truncate s as before to an optimal dual solution to $\text{LP}(c'')$ for some vector c'' where $\|c - c''\|_\infty < 2\delta$. After that, Theorem 5.4 is applicable for the costs c, c'' and said optimal dual solution. Otherwise, we continue with additional phases.

The algorithm formalizes the above idea, with some technical modifications. The algorithm comprises at most $m+1$ phases; the main potential is that the rank of the large index set \mathcal{L} increases in every phase. We show that if an index $i \notin \text{cl}(\mathcal{L})$ was added to \mathcal{L} , then it must have $s_i < \delta$ at the beginning of every later phase. Thus, these indices cannot be violating anymore.

We now turn to a more formal description of Algorithm 4. We start by orthogonally projecting the input cost vector c to $\ker(A)$. This does not change the optimal face of **(LP)**. If $c = 0$, then we terminate and return the current feasible solution $x^{(0)}$ as it is optimal. Otherwise, we scale the cost to $\|c\|_2 = 1$, and use **RATIO-CIRCUIT** to obtain a feasible solution $\tilde{s}^{(-1)}$ to the dual of **(LP)**.

¹ Taking square roots can be avoided by normalizing with $\|c\|_1 = 1$ or $\|c\|_\infty = 1$ instead, and changing the parameters of the algorithm accordingly.

Algorithm 4: VARIABLE- FIXING

```

Input : Linear program in standard form (LP), and initial feasible solution  $x^{(0)}$ .
Output: Either an optimal solution to (LP), or a feasible solution  $x$  and  $\emptyset \neq N \subseteq [n]$  such that
 $x_N = x_N^* = 0$  for every optimal solution  $x^*$  to (LP).
1  $t \leftarrow 0; k \leftarrow 0; \mathcal{L}_{t-1} \leftarrow \emptyset;$ 
2  $c \leftarrow \Pi_{\ker(A)}(c);$ 
3 if  $c = 0$  then
4   return  $x^{(0)}$ ;
5  $c \leftarrow c/\|c\|_2;$ 
6  $(\cdot, \cdot, \tilde{s}^{(-1)}) \leftarrow \text{RATIO-CIRCUIT}(A, c, \mathbb{1});$   $\triangleright$  Any dual feasible solution to LP(c)
7 while  $\langle \tilde{s}^{(t-1)}, x^{(t)} \rangle > 0$  do
8    $S_t \leftarrow \{i \in [n] : \tilde{s}_i^{(t-1)} \geq \delta\};$ 
9    $\mathcal{L}_t \leftarrow \mathcal{L}_{t-1} \cup \{i \in [n] : x_i^{(t)} \geq \Gamma \|x_{S_t}^{(t)}\|_1\};$ 
10  if  $t = 0$  or  $\text{rk}(\mathcal{L}_t) > \text{rk}(\mathcal{L}_{t-1})$  then
11     $k \leftarrow k + 1;$   $\triangleright$  New phase
12    Set modified cost  $\tilde{c}^{(k)} \in \mathbb{R}_+^n$  as  $\tilde{c}_i^{(k)} \leftarrow \tilde{s}_i^{(t-1)}$  if  $i \in S_t$ , and  $\tilde{c}_i^{(k)} \leftarrow 0$  otherwise;
13    while  $\exists$  a circuit in  $\text{supp}(x^{(t)})$  intersecting  $\text{supp}(\tilde{c}^{(k)})$  do
14       $g^{(t)} \leftarrow \text{SUPPORT-CIRCUIT}(A, \tilde{c}^{(k)}, x^{(t)}, \text{supp}(\tilde{c}^{(k)}));$ 
15       $x^{(t+1)} \leftarrow \text{aug}_P(x^{(t)}, g^{(t)});$ 
16       $\mathcal{L}_{t+1} \leftarrow \mathcal{L}_t; t \leftarrow t + 1;$ 
17     $(g^{(t)}, y^{(t)}, s^{(t)}) \leftarrow \text{RATIO-CIRCUIT}(A, \tilde{c}^{(k)}, 1/x^{(t)});$ 
18    if  $\langle \tilde{c}^{(k)}, g^{(t)} \rangle = 0$  then
19       $x^{(t+1)} \leftarrow x^{(t)};$   $\triangleright$  Terminating in the next iteration by Claim 7.2
20    else
21       $x^{(t+1)} \leftarrow \text{aug}_P(x^{(t)}, g^{(t)});$ 
22       $\tilde{s}^{(t)} \leftarrow \arg \min_{s \in \langle \tilde{c}^{(k)}, s^{(t)} \rangle} \langle s, x^{(t+1)} \rangle; t \leftarrow t + 1;$ 
23   $N \leftarrow \{i \in [n] : \tilde{s}_i^{(t-1)} > \kappa_A(m + 1)n\delta\};$ 
24 return  $(x^{(t)}, N);$ 

```

The rest of Algorithm 4 consists of repeated *phases*, ending when $\langle \tilde{s}^{(t-1)}, x^{(t)} \rangle = 0$. In an iteration t , let $S_t = \{i \in [n] : \tilde{s}_i^{(t-1)} \geq \delta\}$ be the set of coordinates with large dual slack. The algorithm keeps track of the following set

$$\mathcal{L}_t := \mathcal{L}_{t-1} \cup \left\{ i \in [n] : x_i^{(t)} \geq \Gamma \|x_{S_t}^{(t)}\|_1 \right\}.$$

These are the variables that were once large with respect to $\|x_{S_t}^{(t)}\|_1$ in iteration $t' \leq t$. Note that $|\mathcal{L}_t|$ is monotone nondecreasing.

The first phase starts at $t = 0$, and we enter a new phase k whenever $\text{rk}(\mathcal{L}_t) > \text{rk}(\mathcal{L}_{t-1})$. Such an iteration t is called the *first iteration* in phase k . At the start of the phase, we define a new *modified cost* $\tilde{c}^{(k)}$ from the dual slack $\tilde{s}^{(t-1)}$ by truncating entries less than δ to 0. This cost vector will be used until the end of the phase. Then, we call $\text{SUPPORT-CIRCUIT}(A, \tilde{c}^{(k)}, x^{(t)}, \text{supp}(\tilde{c}^{(k)}))$ to eliminate circuits in $\text{supp}(x^{(t)})$ intersecting $\text{supp}(\tilde{c}^{(k)})$. Note that there are at most n such calls because each call sets a primal variable $x_i^{(t)}$ to zero.

In the remaining part of the phase, we augment $x^{(t)}$ using $\text{RATIO-CIRCUIT}(A, \tilde{c}^{(k)}, 1/x^{(t)})$ until $\text{rk}(\mathcal{L}_t)$ increases, triggering a new phase. In every iteration, $\text{RATIO-CIRCUIT}(A, \tilde{c}^{(k)}, 1/x^{(t)})$ returns a minimum cost-to-weight ratio circuit $g^{(t)}$, where the choice of weights $1/x^{(t)}$ follows Wallacher [39]. It also returns a feasible solution $(y^{(t)}, s^{(t)})$ to the dual of $\text{LP}(\tilde{c}^{(k)})$. After augmenting $x^{(t)}$ to $x^{(t+1)}$ using $g^{(t)}$, we update the dual slack as

$$\tilde{s}^{(t)} := \arg \min_{s \in \{\tilde{c}^{(k)}, s^{(t)}\}} \langle s, x^{(t+1)} \rangle.$$

This finishes the description of a phase.

Since $\text{rk}(A) = m$, clearly there are at most $m + 1$ phases. Let k and t be the final phase and iteration of Algorithm 4 respectively. As $\langle \tilde{s}^{(t-1)}, x^{(t)} \rangle = 0$, and $x^{(t)}, \tilde{s}^{(t-1)}$ are primal-dual feasible solutions to $\text{LP}(\tilde{c}^{(k)})$, they are also optimal. Now, it is not hard to see that $\tilde{c}^{(k)} \in \text{Im}(A^\top) + c - r$ for some $0 \leq r \leq (m + 1)\delta \mathbb{1}$ (Claim 7.3). Hence, $\tilde{s}^{(t-1)}$ is also an optimal solution to the dual of $\text{LP}(c - r)$. The last step of the algorithm consists of identifying the set N of coordinates with large dual slack $\tilde{s}_i^{(t-1)}$. Then, applying Theorem 5.4 for $c' = c - r$ allows us to conclude that they can be fixed to zero.

In order to prove Theorem 1.4, we need to show that $N \neq \emptyset$. Moreover, we need to show that there are at most T iterations of RATIO-CIRCUIT per phase. First, we show that the objective value is monotone nonincreasing.

Lemma 7.1 *For any two iterations $r \geq t$ in phases $\ell \geq k \geq 1$ respectively,*

$$\langle \tilde{c}^{(\ell)}, x^{(r)} \rangle \leq \langle \tilde{c}^{(k)}, x^{(t)} \rangle.$$

Proof We proceed by induction on $\ell - k \geq 0$. For the base case $\ell - k = 0$, iterations r and t occur in the same phase. So, the objective value is nonincreasing from the definition of SUPPORT-CIRCUIT and RATIO-CIRCUIT . Next, suppose that the statement holds for $\ell - k = d$, and consider the inductive step $\ell - k = d + 1$. Let q be the first iteration in phase $k + 1$; note that $r \geq q > t$. Then, we have

$$\langle \tilde{c}^{(\ell)}, x^{(r)} \rangle \leq \langle \tilde{c}^{(k+1)}, x^{(q)} \rangle \leq \langle \tilde{s}^{(q-1)}, x^{(q)} \rangle \leq \langle \tilde{c}^{(k)}, x^{(q)} \rangle \leq \langle \tilde{c}^{(k)}, x^{(t)} \rangle.$$

The first inequality uses the inductive hypothesis. In the second inequality, we use that $\tilde{c}^{(k+1)}$ is obtained from $\tilde{s}^{(q-1)}$ by setting some nonnegative coordinates to 0. The third inequality is by the definition of $\tilde{s}^{(q-1)}$. The final inequality is by monotonicity within the same phase. □

The following claim gives a sufficient condition for Algorithm 4 to terminate.

Claim 7.2 *Let t be an iteration in phase $k \geq 1$. If RATIO-CIRCUIT returns an elementary vector $g^{(t)}$ such that $\langle \tilde{c}^{(k)}, g^{(t)} \rangle = 0$, then Algorithm 4 terminates in iteration $t + 1$.*

Proof Recall that the weights w in RATIO- CIRCUIT are chosen as $w = 1/x^{(t)}$. Recall also the constraint $s^{(t)} \leq \lambda w$ in the dual program (3). Hence, for every $i \in \text{supp}(x^{(t)})$, $s_i^{(t)} x_i^{(t)} \leq \lambda = -\langle \tilde{c}^{(k)}, g^{(t)} \rangle$, where the equality is due to strong duality. It follows that $\langle s^{(t)}, x^{(t)} \rangle \leq -n \langle \tilde{c}^{(k)}, g^{(t)} \rangle = 0$. Since $\tilde{s}^{(t)}, x^{(t+1)} \geq 0$, we have

$$0 \leq \langle \tilde{s}^{(t)}, x^{(t+1)} \rangle \leq \langle s^{(t)}, x^{(t+1)} \rangle = \langle s^{(t)}, x^{(t)} \rangle \leq 0.$$

Thus, the algorithm terminates in the next iteration. □

The next two claims provide some basic properties of the modified cost $\tilde{c}^{(k)}$. For convenience, we define $\tilde{c}^{(0)} := c$.

Claim 7.3 *For every phase $k \geq 0$, we have $\tilde{c}^{(k)} \in \text{Im}(A^\top) + c - r$ for some $0 \leq r \leq k\delta\mathbb{1}$.*

Proof We proceed by induction on k . The base case $k = 0$ is trivial. Next, suppose that the statement holds for k , and consider the inductive step $k + 1$. Let t be the first iteration of phase $k + 1$, i.e., $\tilde{c}_i^{(k+1)} = \tilde{s}_i^{(t-1)}$ if $i \in S_t$, and $\tilde{c}_i^{(k+1)} = 0$ otherwise. Note that $\tilde{s}^{(t-1)} \in \{\tilde{c}^{(k)}, s^{(t-1)}\}$. Since both of them are feasible to the dual of LP($\tilde{c}^{(k)}$), we have $\tilde{s}^{(t-1)} \in \text{Im}(A^\top) + \tilde{c}^{(k)}$. By the inductive hypothesis, $\tilde{c}^{(k)} \in \text{Im}(A^\top) + c - r$ for some $0 \leq r \leq k\delta\mathbb{1}$. Hence, from the definition of $\tilde{c}^{(k+1)}$, we have $\tilde{c}^{(k+1)} \in \text{Im}(A^\top) + c - r - q$ for some $0 \leq q \leq \delta\mathbb{1}$ as required. □

Claim 7.4 *For every phase $k \geq 0$, we have $\|\tilde{c}^{(k)}\|_\infty \leq 3\sqrt{n}\kappa_A$.*

Proof We proceed by induction on k . The base case $k = 0$ is easy because $\|c\|_\infty \leq \|c\|_2 = 1$. Next, suppose that the statement holds for k , and consider the inductive step $k + 1$. Let t be the first iteration of phase $k + 1$. If $\tilde{s}^{(t-1)} = \tilde{c}^{(k)}$, then $\tilde{c}^{(k+1)}$ is obtained from $\tilde{c}^{(k)}$ by setting some coordinates to 0, so we are done by the inductive hypothesis. Otherwise, $\tilde{s}^{(t-1)} = s^{(t-1)}$. We know that $s^{(t-1)}$ is an optimal solution to (3) for RATIO- CIRCUIT($A, \tilde{c}^{(k)}, 1/x^{(t-1)}$). Since $c - r \in \text{Im}(A^\top) + \tilde{c}^{(k)}$ for some $0 \leq r \leq k\delta\mathbb{1}$ by Claim 7.3, $s^{(t-1)}$ is also an optimal solution to (3) for RATIO- CIRCUIT($A, c - r, 1/x^{(t-1)}$). By (4), we obtain

$$\begin{aligned} \|s^{(t-1)}\|_\infty &\leq 2\kappa_A \|c - r\|_1 \leq 2\kappa_A (\|c\|_1 + \|r\|_1) \\ &\leq 2\kappa_A (\sqrt{n} + nk\delta) \leq 2\kappa_A (\sqrt{n} + n(m + 1)\delta) \leq 3\sqrt{n}\kappa_A. \end{aligned}$$

The third inequality is due to $\|c\|_2 = 1$, the fourth inequality follows from the fact that there are at most $m + 1$ phases, and the last inequality follows from the definition of δ . □

We next show a primal proximity lemma that holds for iterates throughout the algorithm.

Lemma 7.5 *Let t be the first iteration of a phase $k \geq 1$. For any iteration $r \geq t$,*

$$\left\| x^{(r+1)} - x^{(r)} \right\|_\infty \leq \frac{3\sqrt{n}\kappa_A^2}{\delta} \left\| x_{S_t}^{(t)} \right\|_1. \tag{8}$$

Proof Fix an iteration $r \geq t$ and let $\ell \geq k$ be the phase in which iteration r occurred. Consider the elementary vector $g^{(r)}$. If it is returned by SUPPORT-CIRCUIT, then $g_i^{(r)} < 0$ for some $i \in \text{supp}(\tilde{c}^{(\ell)})$ by definition. If it is returned by RATIO-CIRCUIT, we also have $g_i^{(r)} < 0$ for some $i \in \text{supp}(\tilde{c}^{(\ell)})$ unless $\langle \tilde{c}^{(\ell)}, g^{(r)} \rangle = 0$. Note that if $\langle \tilde{c}^{(\ell)}, g^{(r)} \rangle = 0$, then the algorithm sets $x^{(r+1)} = x^{(r)}$, which makes the lemma trivially true. Hence, we may assume that such an iteration does not occur.

By construction, we have $x^{(r+1)} - x^{(r)} = \alpha g^{(r)}$ for some $\alpha > 0$, and $\alpha |g_i^{(r)}| \leq x_i^{(r)}$. Applying the definition of κ_A yields

$$\|x^{(r+1)} - x^{(r)}\|_\infty \leq \kappa_A x_i^{(r)} \leq \frac{\kappa_A}{\delta} \langle \tilde{c}^{(\ell)}, x^{(r)} \rangle \leq \frac{\kappa_A}{\delta} \langle \tilde{c}^{(k)}, x^{(t)} \rangle \leq \frac{3\sqrt{n}\kappa_A^2}{\delta} \|x_{S_t}^{(t)}\|_1.$$

The second inequality uses that all nonzero coordinates of $\tilde{c}^{(\ell)}$ are at least δ . The third inequality is by Lemma 7.1, whereas the fourth inequality is by Claim 7.4 and $\text{supp}(\tilde{c}^{(k)}) = S_t$. \square

With the above lemma, we show that any variable which enters \mathcal{L}_t at the start of a phase, is lower bounded by $\text{poly}(n, \kappa_A) \|x_{S_t}^{(t)}\|_1$ in the next $\Theta(mT)$ iterations.

Lemma 7.6 *Let t be the first iteration of a phase $k \geq 1$ and let $i \in \mathcal{L}_t \setminus \mathcal{L}_{t-1}$. For any iteration $t \leq t' \leq t + 2(m + 1)T$,*

$$x_i^{(t')} \geq \frac{6\sqrt{n}\kappa_A^2}{\delta} \|x_{S_t}^{(t)}\|_1.$$

Proof By definition, we have that $x_i^{(t)} \geq \Gamma \|x_{S_t}^{(t)}\|_1$. With Lemma 7.5 we get

$$\begin{aligned} x_i^{(t')} &\geq x_i^{(t)} - \|x^{(t')} - x^{(t)}\|_\infty \geq x_i^{(t)} - \sum_{r=t}^{t'-1} \|x^{(r+1)} - x^{(r)}\|_\infty \\ &\geq \left(\Gamma - \frac{6(m+1)\sqrt{n}\kappa_A^2 T}{\delta} \right) \|x_{S_t}^{(t)}\|_1 \\ &\geq \frac{6\sqrt{n}\kappa_A^2 T}{\delta} \|x_{S_t}^{(t)}\|_1. \end{aligned}$$

The lower bound follows from $T \geq 1$, as long as the constant in the definition of T is chosen large enough. \square

For any iteration t in phase $k \geq 1$, let us define

$$D_t := \bigcup \{ \mathcal{L}_{t'} \setminus \mathcal{L}_{t'-1} : t' \text{ is the first iteration of phase } k' = 1, 2, \dots, k \}. \tag{9}$$

These are the variables which entered $\mathcal{L}_{t'}$ at the start of a phase for all $t' \leq t$. Note that $\text{rk}(D_t) = \text{rk}(\mathcal{L}_t)$ holds. As a consequence of Lemma 7.6, D_t remains disjoint from the support of the modified cost $\tilde{c}^{(k)}$.

Lemma 7.7 *Let $0 \leq t \leq 2(m + 1)T$ be an iteration and let $k \geq 1$ be the phase in which iteration t occurred. Let $D_t \subseteq \mathcal{L}_t$ be defined as in (9). If $\langle \tilde{c}^{(k)}, x^{(t)} \rangle > 0$, then*

$$D_t \cap \text{supp}(\tilde{c}^{(k)}) = \emptyset.$$

Proof For the purpose of contradiction, suppose that there exists an index $i \in D_t \cap \text{supp}(\tilde{c}^{(k)})$. Let $r \leq t$ be the iteration in which i was added to \mathcal{L}_r . By our choice of D_t , r is the first iteration of phase j for some $j \leq k$, which implies that $S_r = \text{supp}(\tilde{c}^{(j)})$. Since $\langle \tilde{c}^{(j)}, x^{(r)} \rangle \geq \langle \tilde{c}^{(k)}, x^{(t)} \rangle > 0$ by Lemma 7.1, we have $\|x_{S_r}^{(r)}\|_1 > 0$. However, we get the following contradiction

$$6\sqrt{n}\kappa_A^2 \|x_{S_r}^{(r)}\|_1 \leq \delta x_i^{(t)} \leq \langle \tilde{c}^{(k)}, x^{(t)} \rangle \leq \langle \tilde{c}^{(j)}, x^{(r)} \rangle \leq 3\sqrt{n}\kappa_A \|x_{S_r}^{(r)}\|_1.$$

The first inequality is by Lemma 7.6, the third inequality is by Lemma 7.1, while the fourth inequality is by Claim 7.4. □

The following lemma shows that RATIO- CIRCUIT geometrically decreases the norm $\|x_{S_t}^{(t)}\|_1$.

Lemma 7.8 *Let t be the first RATIO- CIRCUIT iteration in phase $k \geq 1$. After $p \in \mathbb{N}$ consecutive RATIO- CIRCUIT iterations in phase k ,*

$$\|x_{S_{t+p}}^{(t+p)}\|_1 \leq \frac{3n^{1.5}\kappa_A}{\delta} \left(1 - \frac{1}{n}\right)^{p-1} \|x_{\text{supp}(\tilde{c}^{(k)})}^{(t)}\|_1,$$

Proof

$$\begin{aligned} \|x_{S_{t+p}}^{(t+p)}\|_1 &\leq \frac{1}{\delta} \langle \tilde{s}^{(t+p-1)}, x^{(t+p)} \rangle && \text{(as } \tilde{s}_i^{(t+p-1)} \geq \delta \text{ for all } i \in S_{t+p}) \\ &\leq \frac{1}{\delta} \langle s^{(t+p-1)}, x^{(t+p)} \rangle && \text{(from the definition of } \tilde{s}^{(t+p-1)}) \\ &= \frac{1}{\delta} \langle s^{(t+p-1)}, x^{(t+p-1)} + \alpha g^{(t+p-1)} \rangle && \text{(for some augmentation step size } \alpha) \\ &= \frac{1}{\delta} \left(\langle s^{(t+p-1)}, x^{(t+p-1)} \rangle + \alpha \langle \tilde{c}^{(k)}, g^{(t+p-1)} \rangle \right) && \text{(as } s^{(t+p-1)} \in \text{Im}(A^\top) + \tilde{c}^{(k)}) \\ &\leq \frac{1}{\delta} \langle s^{(t+p-1)}, x^{(t+p-1)} \rangle && \text{(because } \langle \tilde{c}^{(k)}, g^{(t+p-1)} \rangle \leq 0) \\ &\leq -\frac{n}{\delta} \langle \tilde{c}^{(k)}, g^{(t+p-1)} \rangle && (s_i^{(t+p-1)} \leq -\langle \tilde{c}^{(k)}, g^{(t+p-1)} \rangle / x_i^{(t+p-1)} \text{ (by (3))} \\ &\leq \frac{n}{\delta} \left(\langle \tilde{c}^{(k)}, x^{(t+p-1)} \rangle - \text{OPT}(\tilde{c}^{(k)}) \right) && \text{(by step size } \alpha \geq 1 \text{ in Lemma 2.5)} \\ &\leq \frac{n}{\delta} \left(1 - \frac{1}{n}\right)^{p-1} \left(\langle \tilde{c}^{(k)}, x^{(t)} \rangle - \text{OPT}(\tilde{c}^{(k)}) \right) && \text{(by geometric decay in Lemma 2.5)} \\ &\leq \frac{n}{\delta} \left(1 - \frac{1}{n}\right)^{p-1} \langle \tilde{c}^{(k)}, x^{(t)} \rangle && \text{(because } \tilde{c}^{(k)} \geq 0) \\ &\leq \frac{3n^{1.5}\kappa_A}{\delta} \left(1 - \frac{1}{n}\right)^{p-1} \|x_{\text{supp}(\tilde{c}^{(k)})}^{(t)}\|_1 && \text{(by Claim 7.4).} \end{aligned}$$

□

Recall Lemma 5.5 which guarantees the existence of a coordinate with large dual slack. It explains why we chose to work with a projected and normalized cost vector in Algorithm 4. We are now ready to prove the main result of this section.

Proof of Theorem 1.4 We first prove the correctness of Algorithm 4. Suppose that the algorithm terminates in iteration t . We may assume that there is at least 1 phase, as otherwise $x^{(0)}$ is an optimal solution to (LP). Let $k \geq 1$ be the phase in which iteration t occurred. Since $\langle \tilde{s}^{(t-1)}, x^{(t)} \rangle = 0$ and $x^{(t)}, \tilde{s}^{(t-1)}$ are primal-dual feasible solutions to $\text{LP}(\tilde{c}^{(k)})$, they are also optimal. By Claim 7.3, we know that $\tilde{c}^{(k)} \in \text{Im}(A^\top) + c - r$ for some $\|r\|_\infty \leq (m + 1)\delta$. Hence, $\tilde{s}^{(t-1)}$ is also an optimal dual solution to $\text{LP}(c')$ where $c' := c - r$. Since $c \in \ker(A)$, $\|c\|_2 = 1$, and

$$\|c - c'\|_\infty \leq (m + 1)\delta = \frac{m + 1}{2n^{3/2}(m + 2)\kappa_A} < \frac{1}{\sqrt{n}(m + 2)\kappa_A},$$

where the strict inequality is due to $n \geq m$ and $n > 1$, Lemma 5.5 guarantees the existence of an index $j \in [n]$ such that

$$\tilde{s}_j^{(t)} > \frac{(m + 1)}{\sqrt{n}(m + 2)} > (m + 1)\kappa_A \|c - c'\|_\infty.$$

Thus, the algorithm returns $N \neq \emptyset$. Moreover, for all $j \in N$, Theorem 5.4 allows us to conclude that $x_j^{(t)} = x_j^* = 0$ for every optimal solution x^* to $\text{LP}(c)$.

Next, we show that if $\text{rk}(\mathcal{L}_t) = m$ in some phase k , then the algorithm will terminate in iteration $r \leq t + n + 1$. As long as $\langle \tilde{c}^{(k)}, x^{(t)} \rangle > 0$, we have $D_t \subseteq [n] \setminus \text{supp}(\tilde{c}^{(k)})$ by Lemma 7.7. Moreover, any $i \in \text{supp}(\tilde{c}^{(k)}) \cap \text{supp}(x^{(t)})$ induces a circuit in $D_t \cup \{i\}$, so SUPPORT- CIRCUIT will be invoked. Since every call to SUPPORT- CIRCUIT reduces $\text{supp}(x^{(t)})$, all the coordinates in $\text{supp}(\tilde{c}^{(k)})$ will be zeroed-out in at most n calls. Let $t \leq t' \leq t + n$ be the first iteration when $\langle \tilde{c}^{(k)}, x^{(t')} \rangle = 0$. Since RATIO- CIRCUIT returns $g^{(t')}$ with $\langle \tilde{c}^{(k)}, g^{(t')} \rangle = 0$, the algorithm terminates in the next iteration by Claim 7.2.

It is left to bound the number of iterations of Algorithm 4. Clearly, there are at most $m + 1$ phases. In every phase, there are at most n SUPPORT- CIRCUIT iterations because each call sets a primal variable to 0. It is left to show that there are at most T RATIO- CIRCUIT iterations in every phase.

Fix a phase $k \geq 1$ and assume that every phase $\ell < k$ consists of at most T many RATIO- CIRCUIT iterations. Let t be the first iteration in phase k . We may assume that $\text{rk}(\mathcal{L}_t) < m$, as otherwise there is only one RATIO- CIRCUIT iteration in this phase by the previous argument. Note that this implies $\|x_{S'}^{(t')}\|_1 > 0$ for all $t' \leq t$. Otherwise, $\mathcal{L}_{t'} = [n]$ and $\text{rk}(\mathcal{L}_{t'}) = m$, which contradicts $\text{rk}(\mathcal{L}_{t'}) \leq \text{rk}(\mathcal{L}_t)$.

Let $r \geq t$ be the first RATIO- CIRCUIT iteration in phase k . Let $D_r \subseteq \mathcal{L}_r$ be as defined in (9). By Lemma 7.6 and our assumption, we have $x_{D_r}^{(r)} > 0$. We claim that $D_r \cap \text{supp}(\tilde{c}^{(k)}) = \emptyset$. This is clearly the case if $\langle \tilde{c}^{(k)}, x^{(r)} \rangle = 0$. Otherwise, it is given by Lemma 7.7. We also know that there is no circuit in $\text{supp}(x^{(r)})$ which intersects $\text{supp}(\tilde{c}^{(k)})$. Hence, applying Lemma 5.1 with $L = D_r$, $S = \text{supp}(\tilde{c}^{(k)})$, $x = x^{(r)}$, $z = x^{(r+T)}$ yields

$$\|x_{[n] \setminus \text{cl}(D_r)}^{(r+T)}\|_\infty \geq \frac{\|x_{[n] \setminus \text{cl}(D_r)}^{(r+T)}\|_1}{n} \geq \frac{\|x_{\text{supp}(\tilde{z}^{(k)})}^{(r)}\|_\infty}{n\kappa_A} \geq \frac{\|x_{\text{supp}(\tilde{z}^{(k)})}^{(r)}\|_1}{n^2\kappa_A} \geq \Gamma \|x_{S_{r+T}}^{(r+T)}\|_1,$$

where the last inequality follows from Lemma 7.8 by choosing a sufficiently large constant in the definition of T . Note that $\text{cl}(D_r) = \text{cl}(\mathcal{L}_r)$ because D_r is a spanning subset of \mathcal{L}_r . Thus, there exists an index $i \in [n] \setminus \text{cl}(\mathcal{L}_r)$ which is added to \mathcal{L}_{r+T} , showing that $\text{rk}(\mathcal{L}_{r+T}) > \text{rk}(\mathcal{L}_r)$ as required.

Since the main circuit-augmentation algorithm consists of applying Algorithm 4 at most n times, we obtain the desired bound on the number of iterations. \square

8 Circuits in general form

There are many instances in the literature where circuits are considered outside standard equality form. For example, [2, 16, 26] defined circuits for polyhedra in the general form

$$P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}, \tag{10}$$

where $A \in \mathbb{R}^{m_A \times n}$, $B \in \mathbb{R}^{m_B \times n}$, $b \in \mathbb{R}^{m_A}$, $c \in \mathbb{R}^{m_B}$. It implicitly includes polyhedra in inequality form, which were considered by e.g., [5, 8]. For this setup, they define $g \in \mathbb{R}^n$ to be an elementary vector if

- (i) $g \in \ker(A)$, and
- (ii) Bg is support minimal in the collection $\{By : y \in \ker(A), y \neq 0\}$.

In the aforementioned works, the authors use the term ‘circuit’ also for elementary vectors.

Let us assume that

$$\text{rk} \begin{pmatrix} A \\ B \end{pmatrix} = n. \tag{11}$$

This assumption is needed to ensure that P is pointed; otherwise, there exists a vector $z \in \mathbb{R}^n$, $z \neq 0$ such that $Az = 0$, $Bz = 0$. Thus, the lineality space of P is nontrivial. Note that the circuit diameter is defined as the maximum length of a circuit walk between two vertices; this implicitly assumes that vertices exists and therefore the lineality space is trivial.

Under this assumption, we show that circuits in the above definition are a special case of our definition in the Introduction, and explain how our results in the standard form are applicable. Consider the matrix and vector

$$M := \begin{pmatrix} A & 0 \\ B & I_{m_B} \end{pmatrix}, \quad q := \begin{pmatrix} b \\ d \end{pmatrix},$$

and let $\bar{W} := \ker(M) \subseteq \mathbb{R}^{n+m_B}$. Let J denote the set of the last m_B indices, and $W := \pi_J(\bar{W})$ denote the coordinate projection to J . The assumption (11) guarantees that for each $s \in W$, there is a unique $(x, s) \in \bar{W}$; further, $x \neq 0$ if and only if $s \neq 0$.

Consider the polyhedron

$$\bar{P} = \{(x, s) \in \mathbb{R}^n \times \mathbb{R}^{m_B} : M(x, s) = q, s \geq 0\}.$$

Note that P is the projection of \bar{P} onto the x variables. Let $Q := \pi_J(\bar{P}) \subseteq \mathbb{R}^{m_B}$ be the projection of \bar{P} onto the s variables. It is easy to verify the following statements.

Lemma 8.1 *If (11) holds, then there is an invertible affine one-to-one mapping ψ between Q and P , defined by*

$$M(\psi(s), s) = q.$$

Further, $g \in \mathbb{R}^n$ is an elementary vector as in (i),(ii) above if and only if there exists $h \in \mathbb{R}^{m_B}$ such that $(g, h) \in \bar{W}$, $h \neq 0$ and h is support minimal.

Given such a pair $(g, h) \in \bar{W}$ of elementary vectors, let $s \in Q$ and let $s' := \text{aug}_Q(s, h)$ denote the result of the circuit augmentation starting from s . Then, $\psi(s') = \text{aug}_P(\psi(s), g)$.

Consequently, the elementary vectors of (10) are in one-to-one mapping to elementary vectors in the subspace W as used in this paper. This was also independently shown by Borgwardt and Brugger [1, Corollary 3]. By the last part of the statement, analyzing circuit walks on P reduces to analyzing circuit walks of Q that is given in the subspace form $Q = \{s \in \mathbb{R}^{m_B} : s \in W + r, s \geq 0\}$.

Finally, we can represent Q in standard equality form as follows. Using row operations on M , we can create an $n \times n$ identity matrix in the first n columns. Thus, we can construct a representation $Q = \{s \in \mathbb{R}^{m_B} : Hs = f, s \geq 0\}$, where $H \in \mathbb{R}^{(m_A+m_B-n) \times m_B}$, $f \in \mathbb{R}^{m_A+m_B-n}$. By Lemma 8.1,

$$\kappa_H = \max \left\{ \frac{|(Bg)_i|}{|(Bg)_j|} : i, j \in \text{supp}(Bg), g \text{ is an elementary vector of (10)} \right\}.$$

Acknowledgements The authors would like to thank the anonymous reviewers for their valuable comments, which have helped improve the presentation of this paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Borgwardt, S., Brugger, M.: Circuits in extended formulations. *Discrete Optim.* **52**, 100825 (2024). <https://doi.org/10.1016/j.disopt.2024.100825>
2. Borgwardt, S., De Loera, J.A., Finhold, E.: Edges versus circuits: a hierarchy of diameters in polyhedra. *Adv. Geom.* **16**(4), 511–530 (2016)

3. Borgwardt, S., De Loera, J.A., Finhold, E., Miller, J.: The hierarchy of circuit diameters and transportation polytopes. *Discrete Appl. Math.* **240**, 8–24 (2018)
4. Bonifas, N., Di Summa, M., Eisenbrand, F., Hähnle, N., Niemeier, M.: On sub-determinants and the diameter of polyhedra. *Discrete Comput. Geom.* **52**(1), 102–115 (2014)
5. Borgwardt, S., Finhold, E., Hemmecke, R.: On the circuit diameter of dual transportation polyhedra. *SIAM J. Discrete Math.* **29**(1), 113–121 (2015)
6. Bland, R.G.: On the generality of network flow theory. In: Presented at the ORSA/TIMS Joint National Meeting, Miami, FL (1976)
7. Brunsch, T., Röglin, H.: Finding short paths on polytopes by the shadow vertex algorithm. In: Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP), pp. 279–290. Springer (2013)
8. Borgwardt, S., Stephen, T., Yusun, T.: On the circuit diameter conjecture. *Discrete Comput. Geom.* **60**(3), 558–587 (2018)
9. Borgwardt, S., Viss, C.: An implementation of steepest-descent augmentation for linear programs. *Oper. Res. Lett.* **48**(3), 323–328 (2020)
10. Cook, W., Gerards, A.M., Schrijver, A., Tardos, É.: Sensitivity theorems in integer linear programming. *Math. Program.* **34**(3), 251–264 (1986)
11. Dyer, M., Frieze, A.: Random walks, totally unimodular matrices, and a randomised dual simplex algorithm. *Math. Program.* **64**(1), 1–16 (1994)
12. Dadush, D., Hähnle, N.: On the shadow simplex method for curved polyhedra. *Discrete Comput. Geom.* **56**(4), 882–909 (2016)
13. Dadush, D., Huijberts, S., Natura, B., Végh, L.A.: A scaling-invariant algorithm for linear programming whose running time depends only on the constraint matrix. In: *Mathematical Programming (2023)*
14. De Loera, J.A., Hemmecke, R., Köppe, M.: *Algebraic and Geometric Ideas in the Theory of Discrete Optimization*. Society for Industrial and Applied Mathematics, Philadelphia (2012)
15. De Loera, J.A., Hemmecke, R., Lee, J.: On augmentation algorithms for linear and integer-linear programming: from Edmonds-Karp to Bland and beyond. *SIAM J. Optim.* **25**(4), 2494–2511 (2015)
16. De Loera, J.A., Kafer, S., Sanità, L.: Pivot rules for circuit-augmentation algorithms in linear optimization. *SIAM J. Optim.* **32**(3), 2156–2179 (2022)
17. Dadush, D., Natura, B., Végh, L.A.: Revisiting Tardos’s framework for linear programming: Faster exact solutions using approximate solvers. In: Proc. 61st IEEE Symposium on Foundations of Computer Science (FOCS), pp. 931–942 (2020)
18. Dadush, D., Natura, B., Végh, L.A.: Revisiting Tardos’s framework for linear programming: faster exact solutions using approximate solvers. *arXiv Preprint arXiv:2009.04942* (2020)
19. Ekbatani, F., Natura, B., Végh, L.A.: Circuit imbalance measures and linear programming. In: *Surveys in Combinatorics 2022*. London Mathematical Society Lecture Note Series, pp. 64–114. Cambridge University Press, Cambridge (2022)
20. Eisenbrand, F., Vempala, S.: Geometric random edge. *Math. Program.* **164**(1–2), 325–339 (2017)
21. Fulkerson, D.R.: *Networks, Frames, Blocking Systems*. RAND Corporation, Santa Monica, CA (1967)
22. Gauthier, J.B., Desrosiers, J.: The minimum mean cycle-canceling algorithm for linear programs. *Eur. J. Oper. Res.* **298**, 36–44 (2021)
23. Goldberg, A.V., Tarjan, R.E.: Finding minimum-cost circulations by canceling negative cycles. *J. ACM (JACM)* **36**(4), 873–886 (1989)
24. Kalai, G.: A subexponential randomized simplex algorithm. In: *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pp. 475–482 (1992)
25. Kalai, G., Kleitman, D.J.: A quasi-polynomial bound for the diameter of graphs of polyhedra. *Bull. Am. Math. Soc.* **26**(2), 315–316 (1992)
26. Kafer, S., Pashkovich, K., Sanità, L.: On the circuit diameter of some combinatorial polytopes. *SIAM J. Discrete Math.* **33**(1), 1–25 (2019)
27. Lee, J.: Subspaces with well-scaled frames. *Linear Algebra Appl.* **114**, 21–56 (1989)
28. McCormick, S.T., Shioura, A.: Minimum ratio canceling is oracle polynomial for linear programming, but not strongly polynomial, even for networks. *Oper. Res. Lett.* **27**(5), 199–207 (2000)
29. Rockafellar, R.T.: The elementary vectors of a subspace of \mathbb{R}^n . In: Bose, R.C., Dowling, T.A. (eds.) *Combinatorial Mathematics and Its Applications*, Proceedings of the Chapel Hill Conference, pp. 104–127 (1967)
30. Santos, F.: A counterexample to the Hirsch conjecture. *Ann. Math.* **176**, 383–412 (2012)

31. Sturmfels, B., Thomas, R.R.: Variation of cost functions in integer programming. *Math. Program.* **77**(2), 357–387 (1997)
32. Sukegawa, N.: Improving bounds on the diameter of a polyhedron in high dimensions. *Discrete Math.* **340**(9), 2134–2142 (2017)
33. Schulz, A.S., Weismantel, R.: An oracle-polynomial time augmentation algorithm for integer programming. In: *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 967–968 (1999)
34. Tardos, É.: A strongly polynomial minimum cost circulation algorithm. *Combinatorica* **5**(3), 247–255 (1985)
35. Tardos, É.: A strongly polynomial algorithm to solve combinatorial linear programs. *Oper. Res.* **34**, 250–256 (1986)
36. Tunçel, L.: Approximating the complexity measure of Vavasis–Ye algorithm is NP-hard. *Math. Program.* **86**(1), 219–223 (1999)
37. van den Brand, J., Liu, Y.P., Lee, Y.-T., Saranurak, T., Sidford, A., Song, Z., Wang, D.: Minimum cost flows, MDPs, and L1-regression in nearly linear time for dense instances. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 859–869 (2021)
38. Vavasis, S.A., Ye, Y.: A primal-dual interior point method whose running time depends only on the constraint matrix. *Math. Program.* **74**(1), 79–120 (1996)
39. Wallacher, C.: A Generalization of the Minimum-mean Cycle Selection Rule in Cycle Canceling Algorithms. Unpublished Manuscript. Institute für Angewandte Mathematik, Technische Universität Braunschweig (1989)
40. Wayne, K.D.: A polynomial combinatorial algorithm for generalized minimum cost flow. In: *Mathematics of Operations Research*, pp. 445–459 (2002)
41. Wallacher, C., Zimmermann, U.T.: A polynomial cycle canceling algorithm for submodular flows. *Math. Program.* **86**(1), 1–15 (1999)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.