Faster Width-dependent Algorithm for Mixed Packing and Covering LPs

Digvijay Boob Georgia Tech Atlanta, GA digvi jaybb40@gatech.edu Saurabh Sawlani Georgia Tech Atlanta, GA sawlani@gatech.edu

Di Wang* Google AI Atlanta, GA wadi@google.com

Abstract

In this paper, we give a faster width-dependent algorithm for mixed packingcovering LPs. Mixed packing-covering LPs are fundamental to combinatorial optimization in computer science and operations research. Our algorithm finds a $1 + \varepsilon$ approximate solution in time $O(Nw/\varepsilon)$, where N is number of nonzero entries in the constraint matrix, and w is the maximum number of nonzeros in any constraint. This algorithm is faster than Nesterov's smoothing algorithm which requires $O(N\sqrt{nw}/\varepsilon)$ time, where n is the dimension of the problem. Our work utilizes the framework of area convexity introduced in [Sherman-FOCS'17] to obtain the best dependence on ε while breaking the infamous ℓ_{∞} barrier to eliminate the factor of \sqrt{n} . The current best width-independent algorithm for this problem runs in time $O(N/\varepsilon^2)$ [Young-arXiv-14] and hence has worse running time dependence on ε . Many real life instances of mixed packing-covering problems exhibit small width and for such cases, our algorithm can report higher precision results when compared to width-independent algorithms. As a special case of our result, we report a $1 + \varepsilon$ approximation algorithm for the densest subgraph problem which runs in time $O(md/\varepsilon)$, where m is the number of edges in the graph and d is the maximum graph degree.

1 Introduction

Mixed packing and covering linear programs (LPs) are a natural class of LPs where coefficients, variables, and constraints are non-negative. They model a wide range of important problems in combinatorial optimization and operations research. In general, they model any problem which contains a limited set of available resources (packing constraints) and a set of demands to fulfill (covering constraints).

Two special cases of the problem have been widely studied in literature: pure *packing*, formulated as $\max_x \{b^T x \mid Px \leq p\}$; and pure *covering*, formulated as $\min_x \{b^T x \mid Cx \geq c\}$ where P, p, C, c, b are all non-negative. These are known to model fundamental problems such as maximum bipartite graph matching, minimum set cover, etc. [LN93]. Algorithms to solve packing and covering LPs have also been applied to great effect in designing flow control systems [BBR04], scheduling problems [PST95], zero-sum matrix games [Nes05] and in mechanism design [ZN01]. In this paper, we study the mixed packing and covering (MPC) problem, formulated as checking the feasibility of the set: $\{x \mid Px \leq p, Cx \geq c\}$, where P, C, p, c are non-negative. We say that x is an ε -approximate solution to MPC if it belongs to the relaxed set $\{x \mid Px \leq (1 + \varepsilon)p, Cx \geq (1 - \varepsilon)c\}$. MPC is a generalization of pure packing and pure covering, hence it is applicable to a wider range of problems such as multi-commodity flow on graphs [You01, She17], non-negative linear systems and X-ray tomography [You01].

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

^{*}Work done when author was at Georgia Tech.

General LP solving techniques such as the interior point method can approximate solutions to MPC in as few as $O(\log(1/\varepsilon))$ iterations - however, they incur a large per-iteration cost. In contrast, iterative approximation algorithms based on first-order optimization methods require $poly(1/\varepsilon)$ iterations, but the iterations are fast and in most cases are conducive to efficient parallelization. This property is of utmost importance in the context of ever-growing datasets and the availability of powerful parallel computers, resulting in much faster algorithms in relatively low-precision regimes.

1.1 Previous work

In literature, algorithms for the MPC problem can be grouped into two broad categories: widthdependent and width-independent. Here, width is an intrinsic property of a linear program which typically depends on the dimensions and the largest entry of the constraint matrix, and is an indication of the range of values any constraint can take. In the context of this paper and the MPC problem, we define w_P and w_C as the maximum number of non-zeros in any constraint in P and C respectively.

We define the width of the LP as $w \stackrel{\text{def}}{=} \max(w_P, w_C)$.

One of the first approaches used to solve LPs was Langrangian-relaxation: replacing hard constraints with loss functions which enforce the same constraints indirectly. Using this approach, Plotkin, Schmoys and Tardos [PST95], and Grigoriadis and Khachiyan [GK96] obtained width-dependent polynomial-time approximation algorithms for MPC. Luby and Nisan [LN93] gave the first width-dependent parallelizable algorithm for pure packing and pure covering, which ran in $\tilde{O}(\varepsilon^{-4})$ parallel time, and $\tilde{O}(N\varepsilon^{-4})$ total work. Here, *parallel time* (sometimes termed as *depth*) refers to the longest chain of dependent operations, and *work* refers to the total number of operations in the algorithm.

Young [You01] extended this technique to give the first width-independent parallel algorithm for MPC in $\tilde{O}(\varepsilon^{-4})$ parallel time, and $\tilde{O}(md\varepsilon^{-2})$ total work². Young [You14] later improved his algorithm to run using total work $O(N\varepsilon^{-2})$. Mahoney *et al.* [MRWZ16] later gave an algorithm with a faster parallel run-time of $\tilde{O}(\varepsilon^{-3})$.

The other most prominent approach in literature towards solving an LP is by converting it into a smooth function [Nes05], and then applying general first-order optimization techniques [Nes05, Nes12]. Although the dependence on ε from using first-order techniques is much improved, it usually comes at the cost of sub-optimal dependence on the input size and width. For the MPC problem, Nesterov's accelerated method [Nes12], as well as Bienstock and Iyengar's adaptation [BI06] of Nesterov's smoothing [Nes05], give rise to algorithms with runtime linearly depending on ε^{-1} , but with far from optimal dependence on input size and width. For pure packing and pure covering problems, however, Allen-Zhu and Orrechia [AO19] were the first to incorporate Nesterov-like acceleration while still being able to obtain near-linear width-independent runtimes, giving a $\tilde{O}(N\varepsilon^{-1})$ time algorithm for the packing problem. For the covering problem, they gave a $\tilde{O}(N\varepsilon^{-1.5})$ time algorithm, which was then improved to $\tilde{O}(N\varepsilon^{-1})$ by [WRM16]. Importantly, however, the above algorithms do not generalize to MPC.

1.2 Our contributions

We give the best parallel width-dependent algorithm for MPC, while only incurring a linear dependence on ε^{-1} in the parallel runtime and total work. Additionally, the total work has near-linear dependence on the input-size. Formally, we state our main theorem as follows.

Theorem 1.1. There exists a parallel ε -approximation algorithm for the mixed packing covering problem, which runs in $\widetilde{O}(w \cdot \varepsilon^{-1})$ parallel time, while performing $\widetilde{O}(w \cdot N \cdot \varepsilon^{-1})$ total work, where N is the total number of non-zeros in the constraint matrices, and w is the width of the given LP.

Table 1 compares the running time of our algorithm to previous works solving this problem.

Sacrificing width independence for faster convergence with respect to precision proves to be a valuable trade-off for several combinatorial optimization problems which naturally have a low width. Prominent examples of such problems which are not pure packing or covering problems include *multicommodity flow* and *densest subgraph*, where the width is bounded by the degree of a vertex. In a large number of real-world graphs, the maximum vertex degree is usually small, hence our

 $^{^{2}}d$ here is the maximum number of constraints that any variable appears in.

	Parallel Runtime	Total Work	Comments
Young [You01]	$\widetilde{O}(\varepsilon^{-4})$	$\widetilde{O}(md\varepsilon^{-2})$	d is column-width
Bienstock and Iyengar [BI06]		$\widetilde{O}(n^{2.5}w_P^{1.5}w\varepsilon^{-1})$	width-dependent
Nesterov [Nes12]	$\widetilde{O}(w\sqrt{n}\varepsilon^{-1})$	$\widetilde{O}(w \cdot N\sqrt{n}\varepsilon^{-1})$	width-dependent
Young [You14]	$\widetilde{O}(\varepsilon^{-4})$	$\widetilde{O}(N\varepsilon^{-2})$	
Mahoney et al. [MRWZ16]	$\widetilde{O}(\varepsilon^{-3})$	$\widetilde{O}(N\varepsilon^{-3})$	
This paper	$\widetilde{O}(w\varepsilon^{-1})$	$\widetilde{O}(wN\varepsilon^{-1})$	width-dependent

Table 1: Comparison of runtimes of ε -approximation algorithms for the mixed packing covering problem.

algorithm proves to be much faster when we want high-precision solutions. We explicitly show that this result directly gives the fastest algorithm for the densest subgraph problem on low-degree graphs in Appendix C.

2 Notation and Definitions

For any integer q, we represent using $\|\cdot\|_q$ the q-norm of any vector. We represent the infinity-norm as $\|\cdot\|_{\infty}$. We denote the infinity-norm ball (sometimes called the ℓ_{∞} ball) as the set $\mathcal{B}^n_{\infty}(r) \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n : \|x\|_{\infty} \leqslant r\}$. The nonnegative part of this ball is denoted as $\mathcal{B}^n_{+,\infty}(r) = \{x \in \mathbb{R}^n : x \ge \mathbf{0}_n, \|x\|_{\infty} \leqslant r\}$. For radius r = 1, we drop the radius specification and use the short notation \mathcal{B}^n_{∞} and $\mathcal{B}^n_{+,\infty}$. We denote the extended simplex of dimension k as $\Delta^+_k \stackrel{\text{def}}{=} \{x \in \mathbb{R}^k : \sum_{i=1}^k x_i \leqslant 1\}$. For any $y \ge \mathbf{0}_k$, $\operatorname{proj}_{\Delta^+_k}(y) = y/\|y\|_1$ if $\|y\|_1 \ge 1$. Further, for any set K, we represent its interior, relative interior and closure as $\operatorname{int}(K)$, $\operatorname{relint}(K)$ and $\operatorname{cl}(K)$, respectively. The function exp is applied to a vector element wise.

For any matrix A, we use $\operatorname{nnz}(A)$ to denote the number of nonzero entries in it. We use $A_{i,:}$ and $A_{:,j}$ to refer to the *i*th row and *j*th column of A respectively. We use notation A_{ij} (or $A_{i,j}$ alternatively) to denote an element in the *i*-th row and *j*-th column of matrix A. $||A||_{\infty}$ denotes the operator norm $||A||_{\infty \to \infty} \stackrel{\text{def}}{=} \sup_{x \neq 0} \frac{||Ax||_{\infty}}{||x||_{\infty}}$. For a symmetric matrix A and an antisymmetric matrix B, we define an operator \geq_i as $A \geq_i B \Leftrightarrow \begin{bmatrix} A & -B \\ B & A \end{bmatrix}$ is positive semi-definite.

We formally define an ε -approximate solution to the mixed packing-covering (MPC) problem as follows.

Definition 2.1. We say that x is an ε -approximate solution of the mixed packing-covering problem if x satisfies $x \in \mathcal{B}^n_{+,\infty}$, $Px \leq (1 + \varepsilon)\mathbf{1}_p$ and $Cx \geq (1 - \varepsilon)\mathbf{1}_c$.

Here, $\mathbf{1}_k$ denotes a vectors of 1's of dimension k for any integer k. The saddle point problem on two sets $x \in X$ and $y \in Y$ can be defined as follows:

$$\min_{x \in X} \max_{y \in Y} \mathcal{L}(x, y), \tag{1}$$

where $\mathcal{L}(x, y)$ is some bilinear form between x and y. For this problem, we define the *primal-dual* gap function as $\sup_{(\bar{x}, \bar{y}) \in X \times Y} \mathcal{L}(x, \bar{y}) - \mathcal{L}(\bar{x}, y)$. This gap function can be used as measure of accuracy of the above saddle point solution.

Definition 2.2. We say that $(x, y) \in X \times Y$ is an ε -optimal solution for (1) if $\sup_{(\bar{x}, \bar{y}) \in X \times Y} \mathcal{L}(x, \bar{y}) - \mathcal{L}(\bar{x}, y) \leq \varepsilon$.

3 Technical overview

The mixed packing-covering (MPC) problem is formally defined as follows.

Given two nonnegative matrices $P \in \mathbb{R}^{p \times n}$, $C \in \mathbb{R}^{c \times n}$, find an $x \in \mathbb{R}^n$, $x \ge 0$, $||x||_{\infty} \le 1$ such that $Px \le \mathbf{1}_p$ and $Cx \ge \mathbf{1}_c$ if it exists, otherwise report infeasibility.

Note that the vector of 1's on the right hand side of the packing and covering constraints can be obtained by simply scaling each constraint appropriately. We also assume that each entry in the matrices P and C is at most 1. This assumption, and subsequently the ℓ_{∞} constraints on x also cause no loss of generality³.

We reformulate MPC as a saddle point problem, as defined in Section 2;

$$\lambda^* \stackrel{\text{def}}{=} \min_{x \in \mathcal{B}^n_{+,\infty}} \max_{y \in \Delta^+_c, \ z \in \Delta^+_p} L(x, y, z), \tag{2}$$

where $L(x, y, z) := \begin{bmatrix} y^T \ z^T \end{bmatrix} \begin{bmatrix} P & -\mathbf{1}_p \\ -C & \mathbf{1}_c \end{bmatrix} \begin{bmatrix} x \\ 1 \end{bmatrix}$. The relation between the two formulations is shown in Section 4. For the rest of the paper, we focus on the saddle point formulation (2).

 $\eta(x) \stackrel{\text{def}}{=} \max_{y \in \Delta_c^+, z \in \Delta_p^+} L(x, y, z)$ is a piecewise linear convex function. Assuming oracle access to this "inner" maximization problem, the "outer" problem of minimizing $\eta(x)$ can be performed using first order methods like mirror descent, which are suitable when the underlying problem space is the unit ℓ_{∞} ball. One drawback of this class of methods is that their rate of convergence, which is standard for non-accelerated first order methods on non-differentiable objectives, is $O(\frac{1}{\varepsilon^2})$ to obtain an ε -approximate minimizer x of η which satisfies $\eta(x) \leq \eta^* + \varepsilon$, where η^* is the optimal value. This means that the algorithm needs to access the inner maximization oracle $O(\frac{1}{\varepsilon^2})$ times, which can become prohibitively large in the high precision regime.

Note that even though η is a piecewise linear non-differentiable function, it is not a black box function, but a maximization linear functions in x. This structure can be exploited using Nesterov's smoothing technique [Nes05]. In particular, $\eta(x)$ can be approximated by choosing a strongly convex³ function $\phi : \Delta_p^+ \times \Delta_c^+ \to \mathbb{R}$ and considering

$$\widetilde{\eta}(x) = \max_{y \in \Delta_c^+, z \in \Delta_p^+} L(x, y, z) - \phi(y, z).$$

This strongly convex regularization yields that $\tilde{\eta}$ is a Lipschitz-smooth⁴ convex function. If L is the constant of Lipschitz smoothness of $\tilde{\eta}$ then application of any of the accelerated gradient methods in literature will converge in $O(\sqrt{L/\varepsilon})$ iterations. Moreover, it can also be shown that in order to construct a smooth ε -approximation $\tilde{\eta}$ of η , the Lipschitz smoothness constant L can be chosen to be of the order $O(1/\varepsilon)$, which in turn implies an overall convergence rate of $O(1/\varepsilon)$. In particular, Nesterov's smoothing achieves an oracle complexity of $O((||P||_{\infty} + ||C||_{\infty})D_x \max\{D_y, D_z\}\varepsilon^{-1})$, where where D_x , D_y and D_z denote the sizes of the ranges of their respective regularizers which are strongly convex functions. D_y and D_z can be made of the order of $\log p$ and $\log c$, respectively. However, D_x can be problematic since x belongs to an ℓ_{∞} ball. More on this will soon follow.

Nesterov's dual extrapolation algorithm[Nes07] gives a very similar complexity but is a different algorithm in that it directly addresses the saddle point formulation (2) rather than viewing the problem as optimizing a non-smooth function η . The final convergence for the dual extrapolation algorithm is given in terms of the *primal-dual gap* function of the saddle point problem (2). This algorithms views the saddle point problem as solving variational inequality for an appropriate monotone operator in joint domain (x, y, z). Moreover, as opposed to smoothing techniques which only regularize the dual, this algorithm regularizes both primal and dual parts (*joint regularization*), hence is a different scheme altogether.

Note that for both schemes mentioned above, the maximization oracle itself has an analytical expression which involves matrix-vector multiplication. Hence each call to the oracle incurs a sequential run-time of nnz(P) + nnz(C). Then, overall complexity for both schemes is of order $O((nnz(P) + nnz(C))(||P||_{\infty} + ||C||_{\infty})D_x \max\{D_y, D_z\}\varepsilon^{-1}).$

³This transformation can be achieved by adapting techniques from [WRM16] while increasing dimension of the problem up to a logarithmic factor. Details of this fact are in Appendix B in the full version of this paper. For the purpose of the main text, we work with this assumption.

⁴Definitions of Lipschitz-smoothness and strong convexity can be found in many texts in nonlinear programming and machine learning. e.g. [Bub14]. Intuitively, f is Lipschitz-smooth if the rate of change of ∇f can be bounded by a quantity known as the "constant of Lipschitz smoothness".

The ℓ_{∞} barrier

Note that the first method, i.e., Nesterov's smoothing technique has known lower bounds due to [GN15] (see Corollary 1 in their paper). According to their result, the framework of Nesterov's smoothing has a known limitation since it only regularizes the dual variables. As opposed to this, Nesterov's dual extrapolation regularizes both primal and dual variables, and has potential to skip the earlier mentioned lower bounds of [GN15]. However, the complexity result of this method involves a D_x term, which denotes the range of a convex function over the domain of x. The following lemma states a lower bound for this range in case of ℓ_{∞} balls.

Lemma 3.1. Any strongly convex function has a range of at least $\Omega(\sqrt{n})$ on any ℓ_{∞} ball.

Since $D_x = \Omega(\sqrt{n})$ for each member function of this wide class, there is no hope of eliminating this \sqrt{n} factor using techniques involving explicit use of strong convexity.

So, the goal now is to find a joint regularization function with a small range over ℓ_{∞} balls, but still act as good enough regularizers to enable accelerated convergence of the descent algorithm. In pursuit of breaking this ℓ_{∞} barrier, we draw inspiration from the notion of *area convexity* introduced by Sherman [She17]. Area convexity is a weaker notion than strong convexity, however, it is still strong enough to ensure that accelerated first order methods still go through when using area convex regularizers. Since this is a weaker notion than strong convexity, we can construct area convex functions which have range of $O(n^{o(1)})$ on ℓ_{∞} ball.

First, we define area convexity, and then go on to mention its relevance to the saddle point problem (2).

Area convexity is a notion defined in context of a matrix $A \in \mathbb{R}^{a \times b}$ and a convex set $K \subseteq \mathbb{R}^{a+b}$. Let $M_A \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{0}_{b \times b} & -A^T \\ A & \mathbf{0}_{a \times a} \end{bmatrix}$.

Definition 3.2 ([She17]). A function ϕ is *area convex* with respect to a matrix A on a convex set K iff for any $t, u, v \in K$, ϕ satisfies $\phi\left(\frac{t+u+v}{3}\right) \leq \frac{1}{3}\left(\phi(t)+\phi(u)+\phi(v)\right) - \frac{1}{3\sqrt{3}}(v-u)^T M_A(u-t)$.

To understand the definition above, let us first look at the notion of strong convexity. ϕ is said to be strongly convex if for any two points $t, u, \frac{1}{2}(\phi(t) + \phi(u))$ exceeds $\phi(\frac{1}{2}(t+u))$ by an amount proportional to $||t - u||_2^2$. Definition 3.2 generalizes this notion in context of matrix A for any three points x, y, z. ϕ is area-convex on set K if for any three points $t, u, v \in K$, we have $\frac{1}{3}(\phi(t) + \phi(u) + \phi(v))$ exceeds $\phi(\frac{1}{3}(t+u+v))$ by an amount proportional to the area of the triangle defined by the convex hull of t, u, v.

Consider the case that points t, u, v are collinear. For this case, the area term (i.e., the term involving M_A) in Definition 3.2 is 0 since matrix M_A is antisymmetric. In this sense, area convexity is even weaker than strict convexity. Moreover, the notion of area is parameterized by matrix A. To see a specific example of this notion of area, consider $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ and $t, u, v \in \mathbb{R}^2$. Then, for all possible permutations of t, u, v, the area term takes a value equal to $\pm(t_1(u_2 - v_2) + u_1(v_2 - t_2) + v_1(t_2 - u_2))$. Since the condition holds irrespective of the permutation so we must have that $\phi(\frac{t+u+v}{3}) \leq \frac{1}{3}(\phi(t) + \phi(u) + \phi(v)) - \frac{1}{3\sqrt{3}}|t_1(u_2 - v_2) + u_1(v_2 - t_2) + v_1(t_2 - u_2)|$. But note

that area of triangle formed by points t, u, v is equal to $\frac{1}{2}|t_1(u_2 - v_2) + u_1(v_2 - t_2) + v_1(t_2 - u_2)|$. Hence the area term is just a high dimensional matrix based generalization of the area of a triangle.

Coming back to the saddle point problem (2), we need to pick a suitable area convex function ϕ on the set $\mathcal{B}_{+,\infty}^n \times \Delta_p^+ \times \Delta_c^+$. Since ϕ is defined on the joint space, it has the property of joint regularization vis a vis (2). However, we need an additional parameter: a suitable matrix M_A . The choice of this matrix is related to the bilinear form of the *primal-dual gap function* of (2). We delve into the technical details of this in Section 4, however, we state that the matrix is composed of P, Cand some additional constants. The algorithm we state exactly follows Nesterov's dual extrapolation method described earlier. One notable difference is that in [Nes07], they consider joint regularization by a strongly convex function which does not depend on the problem matrices P, C but only on the constraint set $\mathcal{B}_{+,\infty}^n \times \Delta_p^+ \times \Delta_c^+$. Our area convex regularizer, on the other hand, is tailor made for the particular problem matrices P, C as well as the constraint set.

4 Area Convexity for Mixed Packing Covering LPs

In this section, we present our technical results and algorithm for the MPC problem, with the end goal of proving Theorem 1.1. First, we relate an $(1 + \varepsilon)$ -approximate solution to the saddle point problem to an ε -approximate solution to MPC. Next, we present some theoretical background towards the goal of choosing and analyzing an appropriate area-convex regularizer in the context of the saddle point formulation, where the key requirement of the area convex function is to obtain a provable and efficient convergence result. Finally, we explicitly show an area convex function which is generated using a simple "gadget" function. We show that this area convex function satisfies all key requirements and hence achieves the desired accelerated rate of convergence. This section closely follows [She17], in which the author chooses an area convex function specific to the undirected multicommodity flow problem. Due to space constraints, we relegate almost all proofs to Appendix A (in the full version) and simply include pointers to proofs in [She17] when it is directly applicable.

4.1 Saddle Point Formulation for MPC

Consider the saddle point formulation in (2) for MPC. Given a feasible primal-dual feasible solution pair (x, y, z) and $(\bar{x}, \bar{y}, \bar{z})$ for (2), we denote w = (x, u, y, z) and $\bar{w} = (\bar{x}, \bar{u}, \bar{y}, \bar{z})$ where $u, \bar{u} \in \mathbb{R}$. Then, we define a function $Q : \mathbb{R}^{n+1+p+c} \times \mathbb{R}^{n+1+p+c} \to \mathbb{R}$ as

$$Q(w, \bar{w}) \stackrel{\text{def}}{=} \begin{bmatrix} \bar{y}^T \ \bar{z}^T \end{bmatrix} \begin{bmatrix} P & -\mathbf{1}_p \\ -C & \mathbf{1}_c \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} - \begin{bmatrix} y^T \ z^T \end{bmatrix} \begin{bmatrix} P & -\mathbf{1}_p \\ -C & \mathbf{1}_c \end{bmatrix} \begin{bmatrix} \bar{x} \\ \bar{u} \end{bmatrix}$$

Note that if $u = \overline{u} = 1$, then

$$\sup_{\bar{w}\in\mathcal{W}}Q(w,\bar{w}) = \sup_{\bar{x}\in\mathcal{B}^n_+, \infty, \bar{y}\in\Delta^+_p, \bar{z}\in\Delta^+_c}L(x,\bar{y},\bar{z}) - L(\bar{x},y,z)$$

is precisely the primal-dual gap function defined in Section 2. Notice that if (x^*, y^*, z^*) is a saddle point of (2), then we have

$$L(x^*, y, z) \leq L(x^*, y^*, z^*) \leq L(x, y^*, z^*)$$

for all $x \in \mathcal{B}^n_{+,\infty}, y \in \Delta^+_p, z \in \Delta^+_c$. From above equation, it is clear that $Q(w, w^*) \ge 0$ for all $w \in \mathcal{W}$ where $\mathcal{W} \stackrel{\text{def}}{=} \mathcal{B}^n_{+,\infty} \times \{1\} \times \Delta^+_p \times \Delta^+_c$ and $w^* = (x^*, 1, y^*, z^*) \in \mathcal{W}$. Moreover, $Q(w^*, w^*) = 0$. This motivates the following accuracy measure of the candidate approximate solution w.

Definition 4.1. We say that $w \in W$ is an ε -optimal solution of (2) iff

$$\sup_{\bar{w}\in\mathcal{W}}Q(w,\bar{w})\leqslant\varepsilon$$

Remark 4.2. Recall the definition of M_A for a matrix A in Section 3. We can rewrite $Q(w, \bar{w}) = \bar{w}^T J w$ where $J = M_H$ and

$$H = \begin{bmatrix} P & -\mathbf{1}_p \\ -C & \mathbf{1}_c \end{bmatrix} \quad \Rightarrow \quad J := \begin{bmatrix} \mathbf{0}_{n \times n} & \mathbf{0}_{n \times 1} & -P^T & C^T \\ \mathbf{0}_{1 \times n} & 0 & \mathbf{1}_p^T & -\mathbf{1}_c^T \\ P & -\mathbf{1}_p & \mathbf{0}_{p \times p} & \mathbf{0}_{p \times c} \\ -C & \mathbf{1}_c & \mathbf{0}_{c \times p} & \mathbf{0}_{c \times c} \end{bmatrix}.$$

Thus, the gap function in Definition 4.1 can be written in the bilinear form $\sup_{\bar{w} \in \mathcal{W}} \bar{w}^T J w$.

Lemma 4.3 relates the ε -optimal solution of (2) to the ε -approximate solution to MPC.

Lemma 4.3. Let (x, y, z) satisfy $\sup_{(\bar{x}, \bar{y}, \bar{z}) \in \mathcal{B}^n_{+,\infty} \times \Delta^+_p \times \Delta^+_c} L(x, \bar{y}, \bar{z}) - L(\bar{x}, y, z) \leq \varepsilon$. Then either 1. x is an ε -approximate solution of MPC, or 2. y, z satisfy $y^T (P\bar{x} - \mathbf{1}_p) + z^T (-C\bar{x} + \mathbf{1}_c) > 0$ for all $\bar{x} \in \mathcal{B}^n_{+,\infty}$.

This lemma states that in order to find an ε -approximate solution of MPC, it suffices to find ε -optimal solution of (2). Henceforth, we will focus on ε -optimality of the saddle point formulation (2).

4.2 Area Convexity with Saddle Point Framework

Here we state some useful lemmas which help in determining whether a differentiable function is area convex. We start with the following remark which follows from the definition of area convexity (Definition 3.2).

Remark 4.4. If ϕ is area convex with respect to A on a convex set K, and $\overline{K} \subseteq K$ is a convex set, then ϕ is area convex with respect to A on \overline{K} .

The following two lemmas from [She17] provide the key characterization of area convexity.

Lemma 4.5. Let $A \in \mathbb{R}^{2 \times 2}$ symmetric matrix. $A \ge_{i} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \Leftrightarrow A \ge 0$ and $det(A) \ge 1$.

Lemma 4.6. Let ϕ be twice differentiable on the interior of convex set K, i.e., int(K).

- 1. If ϕ is area convex with respect to A on int(K), then $d^2\phi(x) \geq_i M_A$ for all $x \in int(K)$.
- 2. If $d^2\phi(x) \ge_i M_A$ for all $x \in int(K)$, then ϕ is area convex with respect to $\frac{1}{3}A$ on int(K). Moreover, if ϕ is continuous on cl(K), then ϕ is area convex with respect to $\frac{1}{3}A$ on cl(K).

In order to handle the operator \geq_i (recall from Section 2), we state some basic but important properties of this operator, which will come in handy in later proofs.

Remark 4.7. For symmetric matrices A and C and antisymmetric matrices B and D,

- 1. If $A \geq_i B$ then $A \geq_i (-B)$.
- 2. If $A \geq_i B$ and $\lambda \geq 0$ then $\lambda A \geq_i \lambda B$.
- 3. If $A \geq_i B$ and $C \geq_i D$ then $A + C \geq_i (B + D)$.

Having laid a basic foundation for area convexity, we now focus on its relevance to solving the saddle point problem (2). Considering Remark 4.2, we can write the gap function criterion of optimality in terms of bilinear form of the matrix J. Suppose we have a function ϕ which is area convex with respect to H on set W. Then, consider the following *jointly-regularized* version of the bilinear form:

$$\widetilde{\eta}(w) := \sup_{\bar{w} \in \mathcal{W}} \bar{w}^T J w - \phi(\bar{w}).$$
(3)

Similar to Nesterov's dual extrapolation, one can attain $O(1/\varepsilon)$ convergence of accelerated gradient descent for function $\tilde{\eta}(w)$ in (3) over variable w. In order to obtain gradients of $\tilde{\eta}(w)$, we need access to $\operatorname{argmax}_{\bar{w}\in W} \bar{w}^T J w - \phi(\bar{w})$. However, it may not be possible to find an exact maximizer in all cases. Again, one can get around this difficulty by instead using an approximate optimization oracle of the problem in (3).

Definition 4.8. A δ -optimal solution oracle (OSO) for $\phi : W \to \mathbb{R}$ takes input a and outputs $w \in W$ such that

$$a^T w - \phi(w) \ge \sup_{\bar{w} \in \mathcal{W}} a^T \bar{w} - \phi(\bar{w}) - \delta.$$

Given Φ as a δ -OSO for a function ϕ , consider the following algorithm (Algorithm 4.2):

Algorithm 1 Area Convex Mixed Packing Covering (AC-MPC)

Initialize $w_0 = (\mathbf{0}_n, 1, \mathbf{0}_{p+c})$ for $t = 0, \dots, T$ do $w_{t+1} \leftarrow w_t + \Phi(Jw_t + 2J\Phi(Jw_t))$ end for

For Algorithm 4.2, [She17] shows the following:

Lemma 4.9. Let $\phi : \mathcal{W} \to [-\rho, 0]$. Suppose ϕ is area convex with respect to $2\sqrt{3}H$ on \mathcal{W} . Then for $J = M_H$ and for all $t \ge 1$ we have $w_t/t \in \mathcal{W}$ and,

$$\sup_{\bar{w}\in\mathcal{W}}\bar{w}J\tfrac{w_t}{t}\leqslant\delta+\tfrac{\rho}{t}.$$

In particular, in $\frac{\rho}{\varepsilon}$ iterations, Algorithm 4.2 obtain $(\delta + \varepsilon)$ -solution of the saddle point problem (2).

The analysis of this lemma closely follows the analysis of Nesterov's dual extrapolation.

Note that, each iteration consists of O(1) matrix-vector multiplications, O(1) vector additions, and O(1) calls to the approximate oracle. Since the former two are parallelizable to $O(\log n)$ depth, the same remains to be shown for the oracle computation to complete the proof of the run-time in Theorem 1.1.

Recall from the discussion in Section 3 that the critical bottleneck of Nesterov's method is that diameter of the ℓ_{∞} ball is $\Omega(\sqrt{n})$, which is achieved even in the Euclidean ℓ_2 norm. This makes ρ in Lemma 4.9 to also be $\Omega(\sqrt{n})$, which can be a major bottleneck for high dimensional LPs, which are commonplace among real-world applications.

Although, on the face of it, area convexity applied to the saddle point formulation (2) has a similar framework to Nesterov's dual extrapolation, the challenge is to construct a ϕ for which we can overcome the above bottleneck. Particularly, there are three key challenges to tackle:

1. We need to show that existence of a function ϕ that is area convex with respect to H on \mathcal{W} .

2. $\phi : \mathcal{W} \to [-\rho, 0]$ should be such that ρ is not too large.

3. There should exist an efficient δ -OSO for ϕ .

In the next subsection, we focus on these three aspects in order to complete our analysis.

4.3 Choosing an area convex function

First, we consider a simple 2-D gadget function and prove a "nice" property of this gadget. Using this gadget, we construct a function which can be shown to be area convex using the aforementioned property of the gadget.

Let $\gamma_{\beta} : \mathbb{R}^2_+ \to \mathbb{R}$ be a function parameterized by β defined as

$$\gamma_{\beta}(a,b) = ba \log a + \beta b \log b.$$

Lemma 4.10. Suppose $\beta \ge 2$. Then $d^2\gamma_{\beta}(a,b) \ge \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ for all $a \in (0,1]$ and b > 0.

Now, using the function γ_{β} , we construct a function ϕ and use the sufficiency criterion provided in Lemma 4.6 to show that ϕ is area convex with respect to J on W. Note that our set of interest W is not full-dimensional, whereas Lemma (4.6) is only stated for int and not for relint. To get around this difficulty, we consider a larger set $\overline{W} \supset W$ such that \overline{W} is full dimensional and ϕ is area convex on \overline{W} . Then we use Remark 4.4 to obtain the final result, i.e., area convexity of ϕ .

Theorem 4.11. Let
$$w = (x, u, y, z)$$
 and define

$$\begin{split} \phi(w) &\stackrel{\text{def}}{=} \sum_{i=1}^{p} \sum_{j=1}^{n} P_{ij} \gamma_{p_i}(x_j, y_i) + \sum_{i=1}^{p} \gamma_2(u, y_i) + \sum_{i=1}^{c} \sum_{j=1}^{n} C_{ij} \gamma_{c_i}(x_j, z_i) + \sum_{i=1}^{c} \gamma_2(u, z_i), \\ \text{where } p_i &= 2 * \frac{\|P\|_{\infty}}{\|P_{i,:}\|_1} \text{ and } c_i = 2 * \frac{\|C\|_{\infty}}{\|C_{i,:}\|_1}, \text{ then } \phi \text{ is area convex with respect to } \frac{1}{3} \begin{bmatrix} P & -\mathbf{1}_p \\ -C & \mathbf{1}_c \end{bmatrix} \text{ on } \\ \text{set } \overline{\mathcal{W}} &:= \mathcal{B}_{+,\infty}^{n+1}(1) \times \Delta_p^+ \times \Delta_c^+. \text{ In particular, it also implies } 6\sqrt{3}\phi \text{ is area convex with respect to } \\ 2\sqrt{3} \begin{bmatrix} P & -\mathbf{1}_p \\ -C & \mathbf{1}_c \end{bmatrix} \text{ on set } \mathcal{W}. \end{split}$$

Theorem 4.11 addresses the first part of the key three challenges. Next, Lemma 4.12 shows an upper bound on the range of ϕ .

Lemma 4.12. Function $\phi : \mathcal{W} \to [-\rho, 0]$ then $\rho = O(||P||_{\infty} \log p + ||C||_{\infty} \log c)$.

Finally, we need an efficient δ -OSO. Consider the following alternating minimization algorithm.

Algorithm 2 δ -OSO for ϕ

Input $a \in \mathbb{R}^{n+1}$, $a^1 \in \mathbb{R}^p$, $a^2 \in \mathbb{R}^c$, $\delta > 0$ Initialize $(x^0, u^0) \in \mathcal{B}^n_{+,\infty} \times \{1\}$ arbitrarily. for $k = 1, \dots, K$ do $(y^k, z^k) \leftarrow \underset{y \in \Delta_c^+, z \in \Delta_p^+}{\operatorname{argmax}} y^T a^1 + z^T a^2 - \phi(x^{k-1}, u^{k-1}, y, z)$ $(x^k, u^k) \leftarrow \underset{(x,u) \in \mathcal{B}_{+,\infty}^n \times \{1\}}{\operatorname{argmax}} [x^T u] a - \phi(x, u, y^k, z^k)$ end for

[Bec15] shows the following convergence result.

Lemma 4.13. For $\delta > 0$, Algorithm 2 is a δ -OSO for ϕ which converges in $O(\log \frac{1}{\delta})$ iterations.

We show that for our chosen ϕ , we can perform the two argmax computations in each iteration of Algorithm 2 analytically in time $O(\operatorname{nnz}(P) + \operatorname{nnz}(C))$, and hence we obtain a δ -OSO which takes $O((\operatorname{nnz}(P) + \operatorname{nnz}(C) \log \frac{1}{\delta}))$ total work. Parallelizing matrix-vector multiplications eliminates the dependence on nnz(P) and nnz(C), at the cost of another log(N) term.

Lemma 4.14. Each argmax in Algorithm 2 can be computed as follows:

 $\begin{aligned} x^{k} &= \min\{\exp\left(\frac{a}{P^{T}y^{k}+C^{T}z^{k}}-1\right), \mathbf{1}_{n}\} \text{ for all } j \in [n].\\ y^{k} &= \operatorname{proj}_{\Delta_{p}^{+}}\left(\exp\left\{\frac{1}{2(\|P\|_{\infty}+1)}(a^{1}-Px^{k-1}\log x^{k-1})\right\}\right)\\ z^{k} &= \operatorname{proj}_{\Delta_{c}^{+}}\left(\exp\left\{\frac{1}{2(\|C\|_{\infty}+1)}(a^{2}-Cx^{k-1}\log x^{k-1})\right\}\right)\end{aligned}$ In particular, we can compute x^k, y^k, z^k in O(nnz(P) + nnz(C)) work and $O(\log N)$ parallel time.

Acknowledgements

We thank Richard Peng for many important pointers and discussions.

References

- [AO19] Zeyuan Allen-Zhu and Lorenzo Orecchia. Nearly linear-time packing and covering LP solvers - achieving width-independence and -convergence. Math. Program., 175(1-2):307-353, 2019.
- [BBR04] Yair Bartal, John W. Byers, and Danny Raz. Fast, distributed approximation algorithms for positive linear programming with applications to flow control. SIAM J. Comput., 33(6):1261-1279, 2004.
- [Bec15] Amir Beck. On the convergence of alternating minimization for convex programming with applications to iteratively reweighted least squares and decomposition schemes. SIAM Journal on Optimization, 25(1):185–209, 2015.
- [BGM14] Bahman Bahmani, Ashish Goel, and Kamesh Munagala. Efficient primal-dual graph algorithms for mapreduce. In Algorithms and Models for the Web Graph - 11th International Workshop, WAW 2014, Beijing, China, December 17-18, 2014, Proceedings, pages 59-78, 2014.
 - [BI06] Daniel Bienstock and Garud Iyengar. Approximating fractional packings and coverings in o(1/epsilon) iterations. SIAM J. Comput., 35(4):825-854, 2006.
- [Bub14] Sébastien Bubeck. Theory of convex optimization for machine learning. arXiv preprint arXiv:1405.4980, 15, 2014.
- [Cha00] Moses Charikar. Greedy approximation algorithms for finding dense components in a graph. In Proceedings of the Third International Workshop on Approximation Algorithms for Combinatorial Optimization, APPROX '00, pages 84-95, Berlin, Heidelberg, 2000.
- [GK96] Michael D. Grigoriadis and Leonid G. Khachiyan. Approximate minimum-cost multicommodity flows in õ(epsilon⁻²knm) time. Math. Program., 75:477–482, 1996.
- [GN15] Cristóbal Guzmán and Arkadi Nemirovski. On lower complexity bounds for large-scale smooth convex optimization. Journal of Complexity, 31(1):1 – 14, 2015.

- [LN93] Michael Luby and Noam Nisan. A parallel approximation algorithm for positive linear programming. In Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA, pages 448–457, 1993.
- [MRWZ16] Michael W. Mahoney, Satish Rao, Di Wang, and Peng Zhang. Approximating the solution to mixed packing and covering lps in parallel o(epsilon^{-3}) time. In 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy, pages 52:1–52:14, 2016.
 - [Nes05] Yurii Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, 2005.
 - [Nes07] Yurii Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. *Math. Program.*, 109(2-3):319–344, 2007.
 - [Nes12] Yurii Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
 - [PST95] Serge A. Plotkin, David B. Shmoys, and Éva Tardos. Fast approximation algorithms for fractional packing and covering problems. *Math. Oper. Res.*, 20(2):257–301, 1995.
 - [She17] Jonah Sherman. Area-convexity, 1_∞ regularization, and undirected multicommodity flow. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 452–460, 2017.
- [WRM16] Di Wang, Satish Rao, and Michael W. Mahoney. Unified acceleration method for packing and covering problems via diameter reduction. In 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy, pages 50:1–50:13, 2016.
 - [You01] Neal E. Young. Sequential and parallel algorithms for mixed packing and covering. In 42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA, pages 538–546, 2001.
 - [You14] Neal E. Young. Nearly linear-time approximation schemes for mixed packing/covering and facility-location linear programs. *CoRR*, abs/1407.3015, 2014.
 - [ZN01] Edo Zurel and Noam Nisan. An efficient approximate allocation algorithm for combinatorial auctions. In Proceedings 3rd ACM Conference on Electronic Commerce (EC-2001), Tampa, Florida, USA, October 14-17, 2001, pages 125–136, 2001.