Online Covering with Sum of ℓ_q-Norm Objectives*

Viswanath Nagarajan[1] and Xiangkun Shen[2]

1 **University of Michigan, Ann Arbor, MI, USA**
`viswa@umich.edu`
2 **University of Michigan, Ann Arbor, MI, USA**
`xkshen@umich.edu`

Abstract

We consider fractional online covering problems with ℓ_q-norm objectives. The problem of interest is of the form $\min\{f(x) : Ax \geq 1, x \geq 0\}$ where $f(x) = \sum_e c_e \|x(S_e)\|_{q_e}$ is the weighted sum of ℓ_q-norms and A is a non-negative matrix. The rows of A (i.e. covering constraints) arrive online over time. We provide an online $O(\log d + \log \rho)$-competitive algorithm where $\rho = \frac{\max a_{ij}}{\min a_{ij}}$ and d is the maximum of the row sparsity of A and $\max |S_e|$. This is based on the online primal-dual framework where we use the dual of the above convex program. Our result expands the class of convex objectives that admit good online algorithms: prior results required a monotonicity condition on the objective f which is not satisfied here. This result is nearly tight even for the linear special case. As direct applications we obtain (i) improved online algorithms for non-uniform buy-at-bulk network design and (ii) the first online algorithm for throughput maximization under ℓ_p-norm edge capacities.

1998 ACM Subject Classification F.1.2 Modes of Computation

Keywords and phrases online algorithm, covering/packing problem, convex, buy-at-bulk, throughput maximization

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.12

1 Introduction

The online primal-dual approach is a widely used approach for online problems. This involves solving a discrete optimization problem online as follows (i) formulate a linear programming relaxation and obtain a primal-dual online algorithm for it; (ii) obtain an online rounding algorithm for the resulting fractional solution. While this is similar to a linear programming (LP) based approach for offline optimization problems, a key difference is that solving the LP relaxation in the online setting is highly non-trivial. (Recall that there are general polynomial time algorithms for solving LPs offline.) So there has been a lot of effort in obtaining good online algorithms for various classes of LPs: see [1, 13, 23] for pure covering LPs, [13] for pure packing LPs and [5] for certain mixed packing/covering LPs. Such online LP solvers have been useful in obtaining online algorithms for various problems, eg. set cover [2], facility location [1], machine scheduling [5], caching [8] and buy-at-bulk network design [21].

Recently, [6] initiated a systematic study of online fractional covering and packing with *convex* objectives; see also the full versions [7, 11, 15]. These papers obtained good online algorithms for a large class of fractional convex covering problems. They also demonstrated the utility of this approach via many applications that could not be solved using just online LPs. However these results were limited to convex objectives $f : \mathbb{R}^n_+ \to \mathbb{R}_+$ satisfying a

* A full version of the paper is available at `https://arxiv.org/abs/1705.02194`.

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).
Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;
Article No. 12; pp. 12:1–12:12
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

monotone gradient property, i.e. $\nabla f(z) \geq \nabla f(y)$ pointwise for all $z, y \in \mathbb{R}^n$ with $z \geq y$. There are however many natural convex functions that do not satisfy such a gradient monotonicity condition. Note that this condition requires the Hessian $\nabla^2 f(x)$ to be pointwise non-negative in addition to convexity which only requires $\nabla^2 f(x)$ to be positive semidefinite.

In this paper, we focus on convex functions f that are sums of different ℓ_q-norms. This is a canonical class of convex functions with non-monotone gradients and prior results are not applicable; see Section 1.1 for a more detailed comparison. We show that sum of ℓ_q-norm functions admit a logarithmic competitive online algorithm. This result is nearly tight because there is a logarithmic lower bound even for online covering LPs (which corresponds to an ℓ_1 norm objective). We also provide two applications of our result (i) improved competitive ratios (by two logarithmic factors) for some online non-uniform buy-at-bulk problems studied in [21], and (ii) the first online algorithm for throughput maximization with ℓ_p-norm edge capacities (the competitive ratio is logarithmic which is known to be best possible even in the special case of individual edge capacities).

Given that we achieve log-competitive online algorithms for sums of ℓ_q-norms, a natural question is whether such a result holds for all norms. Recall that any norm is a convex function. It turns out that a log-competitive algorithm is not possible for general norms. This follows from a result in [7] which shows an $\Omega(q \log d)$ lower bound for minimizing the objective $\|Bx\|_q$ under covering constraints (where B is a non-negative matrix). It is still an interesting open question to identify the correct competitive ratio for general norm functions.

1.1 Our Results and Techniques

We consider the online covering problem

$$\min \left\{ \sum_{e=1}^{r} c_e \|x(S_e)\|_{q_e} \;:\; Ax \geq \mathbf{1},\, x \in \mathbb{R}^n_+ \right\}, \tag{1}$$

where each $S_e \subseteq [n] := \{1, 2, \cdots n\}$, $q_e \geq 1$, $c_e \geq 0$ and A is a non-negative $m \times n$ matrix. For any $S \subseteq [n]$ and $q \geq 1$ we use the standard notation $\|x(S)\|_q = \left(\sum_{i \in S} x_i^q\right)^{1/q}$. We also consider the dual of this convex program, which is the following packing problem:

$$\max \left\{ \sum_{k=1}^{m} y_k \;:\; A^T y = \mu,\, \sum_{e=1}^{r} \mu_e = \mu,\, \|\mu_e(S_e)\|_{p_e} \leq c_e \,\forall e \in [r],\, y \geq 0 \right\}. \tag{2}$$

The values p_e above satisfy $\frac{1}{p_e} + \frac{1}{q_e} = 1$; so $\|\cdot\|_{p_e}$ is the dual norm of $\|\cdot\|_{q_e}$. This dual can be derived from (1) using Lagrangian duality.

Our framework captures the classic setting of packing/covering LPs when $r = n$ and for each $e \in [n]$ we have $S_e = \{e\}$ and $q_e = 1$. Our main result is:

▶ **Theorem 1.** *There is an $O(\log d + \log \rho)$-competitive online algorithm for (1) and (2) where the covering constraints in (1) and variables y in (2) arrive over time. Here d is the maximum of the row-sparsity of A and $\max_{e=1}^{r} |S_e|$ and $\rho = \frac{\max\{a_{ij}\}}{\min\{a_{ij}\}}$.*

We note that this bound is also the best possible, even in the linear case [13]. For just the covering problem, a better $O(\log d)$ bound is known in the linear case [23] as well as for convex functions with monotone gradients [6].

The algorithm in Theorem 1 is the natural extension of the primal-dual approach for online LPs [13]. We use the gradient $\nabla f(x)$ at the current primal solution x as the cost function, and use this to define a multiplicative update for the primal. Simultaneously, the dual solution y is increased additively. This algorithm is in fact identical to the one in [6] for convex functions with monotone gradients. See Algorithm 1 for the formal description. The contribution of this paper is in the analysis of this algorithm, which requires new ideas to deal with non-monotone gradients.

Limitations of previous approaches [6]

Recall that the general convex covering problem is

$$\min \left\{ f(x) \,:\, Ax \geq \mathbf{1},\, x \in \mathbb{R}^n_+ \right\},$$

where $f : \mathbb{R}^n_+ \to \mathbb{R}_+$ is a convex function. Its dual is:

$$\max \left\{ \sum_{k=1}^{m} y_k - f^*(\mu) \,:\, A^T y = \mu,\, y \geq 0 \right\},$$

where $f^*(\mu) = \max_{x \in \mathbb{R}^n_+} \{\mu^T x - f(x)\}$ is the Fenchel conjugate of f. When f is the sum of ℓ_q-norms, these primal-dual convex programs reduce to (1) and (2).

We restrict the discussion of prior techniques to functions f with $\max_{x \in \mathbb{R}^n_+} \frac{x^T \nabla f(x)}{f(x)} \leq 1$ because this condition is satisfied by sums of ℓ_q norms.[1] At a high level, the analysis in [6] uses the gradient monotonicity to prove a *pointwise upper bound* $A^T y \leq \nabla f(\bar{x})$ where $\bar{x}$ is the final primal solution. This allows them to lower bound the dual objective by $\sum_{k=1}^m y_k$ because $f^*(\nabla f(\bar{x})) \leq 0$ for any $\bar{x}$ (see Lemma 4(d) in [6]). Moreover, proving the pointwise upper bound $A^T y \leq \nabla f(\bar{x})$ is similar to the task of showing dual feasibility in the *linear* case [13, 23] where $\nabla f(\bar{x})$ corresponds to the (fixed) primal cost coefficients.

Below we give a simple example with an ℓ_q-norm objective where the pointwise upper bound $A^T y \leq \nabla f(\bar{x})$ is not satisfied by the online primal-dual algorithm unless the dual solution y is scaled down by a large (i.e. polynomial) factor. This means that one cannot obtain a sub-polynomial competitive ratio for (1) using this approach directly.

Consider an instance with objective function $f(x) = \|x\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$. There are $m = \sqrt{n}$ covering constraints, where the k^{th} constraint is $\sum_{i=k(m-1)+1}^{km} x_i \geq 1$. Note that each variable appears in only one constraint. Let P be the value of primal objective and D be the value of dual objective at any time. Suppose that the rate of increase of the primal objective is at most α times that of the dual; α corresponds to the competitive ratio in the online primal-dual algorithm. Upon arrival of any constraint k, it follows from the primal updates that all the variables $\{x_i\}_{i=k(m-1)+1}^{km}$ increase from 0 to $\frac{1}{m}$. So the increase in P due to constraint k is $(\sqrt{k} - \sqrt{k-1})\frac{1}{\sqrt{m}}$ for iteration k. This means that the increase in D is at least $\frac{1}{\alpha}(\sqrt{k} - \sqrt{k-1})\frac{1}{\sqrt{m}}$, and so $y_k \geq \frac{1}{\alpha}(\sqrt{k} - \sqrt{k-1})\frac{1}{\sqrt{m}}$. Finally, since $\bar{x} = \frac{1}{m}\mathbf{1}$, we know that $\nabla f(\bar{x}) = \frac{1}{m}\mathbf{1}$ (recall $n = m^2$). On the other hand, $(A^T y)_1 = y_1 \geq \frac{1}{\alpha\sqrt{m}}$. Therefore, in order to guarantee $A^T y \leq \nabla f(\bar{x})$ we must have $\alpha \geq \sqrt{m} = n^{1/4}$.

Our approach

First, we show that by duplicating variables and using an online separation oracle approach (as in [1]) one can ensure that the sets $\{S_e\}_{e=1}^r$ are disjoint. This allows for a simple

[1] The result in [6] also applies to other convex functions with monotone gradients, but the competitive ratio depends exponentially on $\max_{x \in \mathbb{R}^n_+} \frac{x^T \nabla f(x)}{f(x)}$.

expression for ∇f which is useful in the later analysis. Then we utilize the specific form of the primal-dual convex programs (1) and (2) and an explicit expression for ∇f to show that the dual y is approximately feasible. In particular we show that $\|y^T A(S_e)\|_{p_e} \leq O(\log d\rho) \cdot c_e$ for each $e \in [r]$; here $A(S_e)$ denotes the submatrix of A with columns from S_e. Note that this is a weaker requirement than upper bounding $A^T y$ pointwise by $\nabla f(\bar{x})$.

In order to bound $\|y^T A(S_e)\|_{p_e}$, we analyze each $e \in [r]$ separately. We partition the steps of the algorithm into *phases* where phase j corresponds to steps where $\Phi_e = \sum_{i\in S_e} x_i^{q_e} \approx \theta^j$; here $\theta > 1$ is a parameter that depends on q_e. The number of phases can be bounded using the fact that Φ_e is monotonically increasing. By triangle inequality we upper bound $\|y^T A(S_e)\|_{p_e}$ by $\sum_j \|y_{(j)}^T A(S_e)\|_{p_e}$ where $y_{(j)}$ denotes the dual variables that arrive in phase j. And in each phase j, we can upper bound $\|y_{(j)}^T A(S_e)\|_{p_e}$ using the differential equations for the primal and dual updates.

Applications

We also provide two applications of Theorem 1.

Non-uniform multicommodity buy-at-bulk. This is a well-studied network design problem in the offline setting [16, 17]. For its online version, the first poly-logarithmic competitive ratio was obtained recently in [21]. A key step in this result was a fractional online algorithm for a certain mixed packing-covering LP. We improve the competitive ratio of this step from $O(\log^3 n)$ to $O(\log n)$ which leads to a corresponding improvement in the final result of [21]. See Theorem 6.

Throughput maximization with ℓ_p-norm capacities. The online problem of maximizing throughput subject to edge capacities is well studied and a tight logarithmic competitive ratio is known [4, 13]. We consider the generalization where instead of individual edge capacities, we can have capacity constraints on subsets as follows. A ℓ_p-norm capacity of c for some subset S of edges means that the ℓ_p-norm of the loads on edges of S must be at most c. We show that one can obtain a randomized log-competitive algorithm even in this setting, which generalizes the case with edge-capacities. See Theorem 7.

1.2 Related Work

The online primal-dual framework for linear programs [14] is fairly well understood. Tight results are known for the class of packing and covering LPs [13, 23], with competitive ratio $O(\log d)$ for covering LPs and $O(\log d\rho)$ for packing LPs; here d is the row-sparsity and ρ is the ratio of the maximum to minimum entries in the constraint matrix. Such LPs are very useful because they correspond to the LP relaxations of many combinatorial optimization problems. Combining the online LP solver with suitable online rounding schemes, good online algorithms have been obtained for many problems, eg. set cover [2], group Steiner tree [1], caching [8] and ad-auctions [12]. Online algorithms for LPs with mixed packing and covering constraints were obtained in [5]; the competitive ratio was improved in [6]. Such mixed packing/covering LPs were also used to obtain an online algorithm for capacitated facility location [5]. A more complex mixed packing/covering LP was used recently in [21] to obtain online algorithms for non-uniform buy-at-bulk network design: as an application of our result, we obtain a simpler and better (by two log-factors) online algorithm for this problem.

There have also been a number of results utilizing the online primal-dual framework with *convex* objectives for specific problems, eg. matching [19], caching [25], energy-efficient scheduling [18, 22] and welfare maximization [10, 24]. All of these results involve separable convex/concave functions. Recently, [6] considered packing/covering problems with general (non-separable) convex objectives, but (as discussed previously) this result requires a monotone gradient assumption on the convex function. The sum of ℓ_q-norm objectives considered in this paper does not satisfy this condition. While our primal-dual algorithm is identical to [6], we need new techniques in the analysis.

All the results above (as well as ours) involve convex objectives and linear constraints. We note that [20] obtained online primal-dual algorithms for certain semidefinite programs (i.e. involving non-linear constraints). While both our result and [20] generalize packing/covering LPs, they are not directly comparable.

We also note that online algorithms with ℓ_q-norm objectives have been studied previously for many scheduling problems, eg. [3, 9]. These results use different approaches and are not directly comparable to ours. More recently [7] used ideas from the online primal-dual approach in an online algorithm for unrelated machine scheduling with ℓ_p-norm objectives as well as startup costs. However, the algorithm in [7] was tailored to their scheduling setting and we do not currently see a connection between our result and [7].

2 Preliminaries

Recall the primal covering problem (1) and its dual packing problem (2). In the online setting, the constraints in the primal and variables in the dual arrive over time. We need to maintain monotonically increasing primal (x) and dual (y) solutions. The following lemma shows that one can assume that the sets $\{S_e\}_{e=1}^r$ are *disjoint* without loss of generality. We defer the proof of the lemma to the full version. This leads to a much simpler expression for ∇f that will be used in Section 3.

▶ **Lemma 2.** *If there is a poly-time α-competitive algorithm for instances with disjoint S_e, then there is a poly-time $O(\alpha)$-competitive algorithm for general instances.*

Henceforth we will assume that the sets $\{S_e\}_{e=1}^r$ are disjoint. The dual program (2) in this case reduces to:

$$\max \left\{ \sum_{k=1}^{m} y_k \;:\; A^T y = \mu,\; \|\mu(S_e)\|_{p_e} \le c_e \, \forall e \in [r],\; y \ge 0 \right\}. \tag{3}$$

It is easy to see that weak duality holds (Lemma 3). Strong duality also holds because (1) satisfies Slater's condition; however we do not use this fact.

▶ **Lemma 3.** *For any pair of feasible solutions x to (1) and (y, μ) to (3), we have*

$$\sum_{e=1}^{r} c_e \|x(S_e)\|_{q_e} \ge \sum_{k=1}^{m} y_k.$$

Proof. This follows from the following inequalities:

$$\sum_{k=1}^{m} y_k = y^T \mathbf{1} \le y^T A x = \mu^T x \le \sum_{e=1}^{r} \sum_{i \in S_e} \mu_i \cdot x_i \le \sum_{e=1}^{r} \|\mu(S_e)\|_{p_e} \cdot \|x(S_e)\|_{q_e} \le \sum_{e=1}^{r} c_e \cdot \|x(S_e)\|_{q_e}.$$

The first inequality is by primal feasibility; the second and last are by dual feasibility; the fourth is by Hölder's inequality. ◀

When the k^{th} request $\sum_{i=1}^{n} a_{ki}x_i \geq 1$ arrives
Let τ be a continuous variable denoting the current time.;
while *the constraint is unsatisfied, i.e.,* $\sum_{i=1}^{n} a_{ki}x_i < 1$ **do**
  For each i with $a_{ki} > 0$, increase x_i at rate $\frac{\partial x_i}{\partial \tau} = \frac{a_{ki}x_i + \frac{1}{d}}{\nabla_i f(x)} = \frac{a_{ki}x_i+\frac{1}{d}}{c_e x_i^{q_e-1}} \|x(S_e)\|_{q_e}^{q_e-1}$;
  Increase y_k at rate $\frac{\partial y_k}{\partial \tau} = 1$;
  Set $\mu = A^T y$;
end

Algorithm 1: Algorithm for ℓ_q-norm packing/covering.

3 Algorithm and analysis

Let $f(x) = \sum_{e=1}^{r} c_e \|x(S_e)\|_{q_e}$ denote the primal objective in (1).

In order to ensure that the gradient ∇f is positive, the primal solution x starts off as $\delta \cdot \mathbf{1}$ where $\delta > 0$ is arbitrarily small. So we assume that the initial primal value is zero.

It is clear that the algorithm maintains a feasible and monotonically non-decreasing primal solution x. The dual solution (y, μ) is also monotonically non-decreasing, but not necessarily feasible. We will show that (y, μ) is $O(\log \rho d)$-approximately feasible, i.e. the packing constraints in (3) are violated by at most an $O(\log \rho d)$ factor.

▶ **Lemma 4.** *The primal objective $f(x)$ is at most twice the dual objective $\sum_{k=1}^{m} y_k$.*

Proof. We will show that the rate of increase of the primal is at most twice that of the dual. Consider the algorithm upon the arrival of some constraint k. Then

$$\frac{df(x)}{d\tau} = \sum_{i:a_{ki}>0} \nabla_i f(x) \cdot \frac{\partial x_i}{\partial \tau} = \sum_{i:a_{ki}>0} (a_{ki}x_i + \frac{1}{d}) \leq 2.$$

The inequality comes from the fact that (i) the process for the kth constraint is terminated when $\sum_i a_{ki}x_i = 1$ and (ii) the number of non-zeroes in constraint k is at most d. Also it is clear that the dual objective increases at rate one, which finishes the proof. ◀

▶ **Lemma 5.** *The dual solution (y, μ) is $O(\log \rho d)$-approximately feasible, i.e.*

$$\|\mu(S_e)\|_{p_e} \leq O(\log \rho d) \cdot c_e, \qquad \forall e \in [r].$$

Proof. Fix any $e \in [r]$. When $q_e = 1$ the analysis is simple (details are deferred to the full version). Here we only consider $q_e > 1$ which is the main part of the analysis. In order to prove the desired upper bound on $\|\mu(S_e)\|_{p_e}$ we use a potential function $\Phi = \sum_{i \in S_e}(x_i^{q_e})$. Let phase zero denote the period where $\Phi \leq \zeta := (\frac{1}{\max\{a_{ij}\} \cdot d^2})^{q_e}$, and for each $j \geq 1$, phase j is the period where $\theta^{j-1} \cdot \zeta \leq \Phi < \theta^j \cdot \zeta$. Here $\theta > 1$ is a parameter depending on q_e that will be determined later. Note that $\Phi \leq d(\frac{1}{\min\{a_{ij}\}})^{q_e}$ as variable x_i will never be increased beyond $1/\min_{j=1}^{m} a_{ij}$. So the number of phases is at most $3q_e \cdot \log(d\rho)/\log\theta$. Next, we bound the increase in $\|\mu(S_e)\|_{p_e}$ for each phase separately.

For any phase, we have the following equalities

$$\frac{\partial x_i}{\partial \tau} = \frac{a_{ki}x_i + \frac{1}{d}}{c_e x_i^{q_e-1}} \|x(S_e)\|_{q_e}^{q_e-1}, \quad \frac{\partial y_k}{\partial \tau} = 1, \quad \frac{\partial \mu_i}{\partial \tau} = a_{ki}$$

$$\Rightarrow d\mu_i = \frac{c_e \, a_{ki} \, x_i^{q_e-1}}{(\sum_{j \in S_e} x_j^{q_e})^{1-\frac{1}{q_e}} (a_{ki}x_i + \frac{1}{d})} dx_i \tag{4}$$

Phase zero. Suppose that each x_i increases to α_i in phase zero. From (4) we have

$$d\mu_i \leq \frac{d\, c_e\, a_{ki}\, x_i^{q_e-1}}{(\sum_{j\in S_e} x_j^{q_e})^{1-\frac{1}{q_e}}} dx_i \qquad \Rightarrow \qquad \frac{1}{d\, c_e\, a_{ki}} d\mu_i \leq \frac{x_i^{q_e-1}}{(\sum_{j\in S_e} x_j^{q_e})^{1-\frac{1}{q_e}}} dx_i.$$

This means that the increase $\Delta\mu_i$ in μ_i can be bounded as:

$$\frac{1}{d\, c_e\, a_{ki}} \Delta\mu_i \leq \int_\delta^{\alpha_i} \frac{x_i^{q_e-1}}{(\sum_{j\in S_e} x_j^{q_e})^{1-\frac{1}{q_e}}} dx_i \leq \int_0^{\alpha_i} 1 dx_i \leq \alpha_i.$$

Since in phase zero, $\Phi \leq (\frac{1}{\max\{a_{ij}\}\cdot d^2})^{q_e}$, we know that each $\alpha_i \leq \frac{1}{\max\{a_{ij}\}\cdot d^2}$. So $\Delta\mu_i \leq \frac{c_e}{d}$ and at the end of phase zero, we have $\|\mu(S_e)\|_{p_e} \leq \|\mu(S_e)\|_1 \leq c_e$. The last inequality is because $d \geq \max_e |S_e|$.

Phase $j \geq 1$. Let Φ_0 and Φ_1 be the value of Φ at the beginning and end of this phase respectively. In phase j, suppose that each x_i increases from s_i to t_i. Then,

$$d\mu_i = \frac{c_e\, a_{ki}\, x_i^{q_e-1}}{(\sum_{j\in S_e} x_j^{q_e})^{1-\frac{1}{q_e}} (a_{ki}x_i + \frac{1}{d})} dx_i \leq \frac{c_e x_i^{q_e-2}}{(\sum_{j\in S_e} x_j^{q_e})^{1-\frac{1}{q_e}}} dx_i\,.$$

So the increase $\Delta\mu_i$ in μ_i during this phase is:

$$\Delta\mu_i \leq \int_{s_i}^{t_i} \frac{c_e x_i^{q_e-2}}{(\sum_{j\in S_e} x_j^{q_e})^{1-\frac{1}{q_e}}} dx_i.$$

Note that variables $x_{i'}$ for $i' \neq i$ can also increase in this phase: so we cannot directly bound the above integral. This is precisely where the potential Φ is useful. We know that throughout this phase, $\sum_{i\in S_e} x_i^{q_e} \geq \Phi_0$. So,

$$\Delta\mu_i \leq c_e \int_{s_i}^{t_i} \frac{x_i^{q_e-2}}{\Phi_0^{1-\frac{1}{q_e}}} dx_i = c_e \frac{t_i^{q_e-1} - s_i^{q_e-1}}{(q_e-1)\Phi_0^{1-\frac{1}{q_e}}} = c_e \frac{t_i^{q_e-1} - s_i^{q_e-1}}{(q_e-1)\Phi_0^{\frac{1}{p_e}}}.$$

Above we used the assumption that $q_e > 1$ in evaluating the integral. Now,

$$\begin{aligned}(\Delta\mu_i)^{p_e} &\leq \frac{c_e^{p_e}}{(q_e-1)^{p_e}\,\Phi_0} \cdot \left(t_i^{q_e-1} - s_i^{q_e-1}\right)^{p_e} \leq \frac{c_e^{p_e}}{(q_e-1)^{p_e}\,\Phi_0} \cdot \left(t_i^{(q_e-1)p_e} - s_i^{(q_e-1)p_e}\right) \\ &= \frac{c_e^{p_e}}{(q_e-1)^{p_e}\,\Phi_0} \cdot (t_i^{q_e} - s_i^{q_e})\ .\end{aligned}$$

The first inequality above uses the fact that $(z_1+z_2)^{p_e} \geq z_1^{p_e} + z_2^{p_e}$ for any $p_e \geq 1$ and $z_1, z_2 \geq 0$, with $z_1 = s_i^{q_e-1}$ and $z_2 = t_i^{q_e-1} - s_i^{q_e-1}$. The last equality uses $\frac{1}{p_e} + \frac{1}{q_e} = 1$.

We can now bound

$$\sum_{i\in S_e} (\Delta\mu_i)^{p_e} \leq \frac{c_e^{p_e}}{(q_e-1)^{p_e}\,\Phi_0} \cdot \sum_{i\in S_e} (t_i^{q_e} - s_i^{q_e}) = \frac{c_e^{p_e}}{(q_e-1)^{p_e}\,\Phi_0} (\Phi_1 - \Phi_0) \leq \frac{c_e^{p_e}}{(q_e-1)^{p_e}} (\theta - 1)\,.$$

Let $\mu_j \in \mathbb{R}^{|S_e|}$ denote the increase in $\mu(S_e)$ during phase j. It follows from the above that $\|\mu_j\|_{p_e} \leq \frac{c_e}{q_e-1}(\theta-1)^{1/p_e}$.

Combining across phases. Note that $\mu = \sum_{j\geq 0} \mu_j$. By triangle inequality, we have

$$\|\mu\|_{p_e} \leq \sum_{j\geq 0} \|\mu_j\|_{p_e} \leq c_e + \sum_{j\geq 1} \|\mu_j\|_{p_e} \leq c_e \left(1 + \frac{3q_e(\theta-1)^{1/p_e}}{(q_e-1)\log\theta} \cdot \log(d\rho)\right). \tag{5}$$

To complete the proof we show next that for any $q_e > 1$, there is some choice of $\theta > 1$ such that the right-hand-side above is $O(\log(d\rho)) \cdot c_e$.

- If $q_e \geq 2$ then setting $\theta = 2$, we have $\frac{3q_e}{(q_e-1)}(\theta-1)^{1/p_e}/\log\theta \leq 6$.
- If $1 < q_e < 2$ then set $\theta = 1 + (q_e-1)^{-\epsilon p_e}$, where $\epsilon = \frac{1}{-\log(q_e-1)} > 0$. We have

$$\frac{(\theta-1)^{\frac{1}{p_e}}}{\log\theta} \leq \frac{(\theta-1)^{\frac{1}{p_e}}}{\log(q_e-1)^{-\epsilon p_e}} = \frac{(q_e-1)^{-\epsilon}}{\log(q_e-1)^{-\epsilon p_e}} = \frac{(q_e-1)^{-\epsilon}}{-\epsilon p_e \log(q_e-1)} = \frac{(q_e-1)^{-\epsilon}}{p_e} = \frac{2}{p_e}.$$

  The first inequality above uses that $\theta - 1 = (q_e-1)^{-\epsilon p_e} > 1$. Thus we have

$$\frac{3q_e(\theta-1)^{1/p_e}}{(q_e-1)\log\theta} \leq \frac{6q_e}{(q_e-1)p_e} = 6,$$

  where the last equality uses $\frac{1}{p_e} + \frac{1}{q_e} = 1$.

So in either case we have that the right-hand-side of (5) is at most $(1 + 6\log(d\rho)) \cdot c_e$. ◀

Combining Lemmas 3, 4 and 5, we obtain Theorem 1.

4 Applications

4.1 Online Buy-at-Bulk Network Design

In the non-uniform buy-at-bulk problem, we are given a directed graph $G = (V, E)$ with a monotone subadditive cost function $g_e : \mathbb{R}_+ \to \mathbb{R}_+$ on each edge $e \in E$ and a collection $\{(s_i, t_i)\}_{i=1}^m$ of m source/destination pairs. The goal is to find an $s_i - t_i$ path P_i for each $i \in [m]$ such that the objective $\sum_{e\in E} g_e(load_e)$ is minimized; here $load_e$ is the number of paths using e. An equivalent view of this problem involves two costs c_e and ℓ_e for each edge $e \in E$ and the objective $\sum_{e\in\cup P_i} c_e + \sum_{e\in E} \ell_e \cdot load_e$. In the online setting, the pairs (s_i, t_i) arrive over time and we need to decide on the path P_i immediately after the i^{th} pair arrives. Recently, [21] gave a modular online algorithm for non-uniform buy-at-bulk with competitive ratio $O(\alpha\beta\gamma \cdot \log^5 n)$ where:

- α is the "junction tree" approximation ratio,
- β is the integrality gap of the natural LP for single-sink instances,
- γ is the competitive ratio of an online algorithm for single-sink instances.

See [21] for more details. One of the main components in this result was an $O(\log^3 n)$-competitive fractional online algorithm for a certain mixed packing/covering LP. Here we show that Theorem 1 can be used to provide a better (and tight) $O(\log n)$-competitive ratio. This leads to the following improvement:

▶ **Theorem 6.** *There is an $O(\alpha\beta\gamma \cdot \log^3 n)$-competitive ratio for non-uniform buy-at-bulk, where α, β, γ are as above.*

The LP relaxation. Let $\mathcal{T} = \{s_i, t_i : i \in [m]\}$ denote the set of all sources/destinations. For each $i \in [m]$ and root $r \in V$ variable z_{ir} denotes the extent to which both s_i and t_i route to/from r. For each $r \in V$ and $e \in E$, variable x_{er} denotes the extent to which edge e is used in the routing to root r. For each $r \in V$ and $u \in \mathcal{T}$, variables $\{f_{r,u,e} : e \in E\}$ represent a flow between r and u. [21] relied on solving the following LP:

$$\begin{aligned}
\min \quad & \sum_{r\in V}\sum_{e\in E} c_e \cdot x_{e,r} \;+\; \sum_{r\in V}\sum_{e\in E} \ell_e \cdot \sum_{u\in\mathcal{T}} f_{r,u,e} \\
\text{s.t.} \quad & \sum_{r\in V} z_{ir} \ge 1, \qquad \forall i \in [m] \\
& \{f_{r,s_i,e} : e \in E\} \text{ is a flow from } s_i \text{ to } r \text{ of } z_{ir} \text{ units}, \qquad \forall r \in V,\ i \in [m] \\
& \{f_{r,t_i,e} : e \in E\} \text{ is a flow from } r \text{ to } t_i \text{ of } z_{ir} \text{ units}, \qquad \forall r \in V,\ i \in [m] \\
& f_{r,u,e} \le x_{e,r}, \qquad \forall u \in \mathcal{T},\ e \in E \\
& x, f, z \ge 0
\end{aligned}$$

The online algorithm in [21] for this LP has competitive ratio $O(D \cdot \log n)$ w.r.t. the optimal integral solution; here D is an upper bound on the length of any $s_i - t_i$ path (note that D can be as large as n). Using a height reduction operation, they could ensure that $D = O(\log n)$ while incurring an additional $O(\log n)$-factor loss in the objective. This lead to the $O(\log^3 n)$ factor for the fractional online algorithm. Here we provide an improved $O(\log n)$-competitive algorithm for this LP which does not require any bound on the path-lengths.

For any $r \in V$ and $u \in \mathcal{T}$, let $\mathsf{MC}(r, u)$ denote the $u - r$ (resp. $r - u$) minimum cut in the graph with edge capacities $\{f_{r,u,e} : e \in E\}$ if u is a source (resp. destination). By the max-flow min-cut theorem, it follows that $z_{ir} \le \min\{\mathsf{MC}(r, s_i), \mathsf{MC}(r, t_i)\}$. Using this, we can combine the first three constraints of the above LP into the following:

$$\sum_{r\in V} \min\{\mathsf{MC}(r, s_i), \mathsf{MC}(r, t_i)\} \;\ge\; 1, \qquad \forall i \in [m].$$

For a fixed $i \in [m]$, this constraint is equivalent to the following. For each $r \in V$, pick either an $s_i - r$ cut (under capacities $f_{r,s_i,\star}$) or an $r - t_i$ cut (under capacities $f_{r,t_i,\star}$), and check if the total cost of these cuts is at least 1. This leads to the following reformulation that eliminates the x and z variables.

$$\begin{aligned}
\min \quad & \sum_{r\in V}\sum_{e\in E} c_e \cdot \left(\max_{u\in\mathcal{T}} f_{r,u,e}\right) \;+\; \sum_{r\in V}\sum_{e\in E} \ell_e \cdot \sum_{u\in\mathcal{T}} f_{r,u,e} \\
\text{s.t.} \quad & \sum_{r\in R_s} f_{r,s_i}(S_r) \;+\; \sum_{r\in R_t} f_{r,t_i}(T_r) \ge 1, \qquad \forall i \in [m],\ \forall (R_s, R_t) \text{ partition of } V, \\
& \qquad \forall S_r : s_i - r \text{ cut}, \forall r \in R_s,\ \forall T_r : r - t_i \text{ cut}, \forall r \in R_t \\
& f \ge 0.
\end{aligned}$$

Note that $\ell_{\log(n)}$-norm is a constant approximation for ℓ_∞. Therefore we can reformulate the above objective function (at the loss of a constant factor) as the sum of $\ell_{\log(n)}$ and ℓ_1 norms. Our fractional solver applies to this convex covering problem, and yields an $O(\log n)$-competitive ratio (note that $\rho = 1$ for this instance). In order to get a polynomial running time, we can use the natural “separation oracle” approach to produce violated covering constraints.

Each iteration above runs in polynomial time since the minimum cuts can be computed in polynomial time. In order to bound the number of iterations, consider the potential $\psi = \sum_{e\in E}(f_{r,s_i,e} + f_{r,t_i,e})$. Note that $0 \le \psi \le 2|E|$ and each iteration increases ψ by at least $\frac{1}{2}$. So the number of iterations is at most $4|E|$.

When the i^{th} request (s_i, t_i) arrives
repeat
For each $r \in V$, compute $\mathsf{MC}(r, s_i)$ and $\mathsf{MC}(r, t_i)$ and the respective cuts S_r and T_r;
Let $R_s = \{r \in V \,:\, \mathsf{MC}(r, s_i) \leq \mathsf{MC}(r, t_i)\}$ and $R_t = V \setminus R_s$;
Run Algorithm 1 with constraint $\sum_{r \in R_s} f_{r,s_i}(S_r) + \sum_{r \in R_t} f_{r,t_i}(T_r) \geq 1$;
until $\sum_{r \in V} \min\{\mathsf{MC}(r, s_i)\,,\, \mathsf{MC}(r, t_i)\} \geq \frac{1}{2}$;

Algorithm 2: Separation Oracle Based Algorithm for Buy-at-Bulk.

4.2 Throughput Maximization with ℓ_p-norm Capacities

The online problem of maximizing multicommodity flow was studied in [4, 13]. In this problem, we are given a directed graph with edge capacities $u(e)$. Requests (s_i, t_i) arrive in an online fashion. The algorithm should choose a path between s_i and t_i and allocate a bandwidth of 1 on the path to serve request i. The total bandwidth allocated on any edge is not allowed to exceed its capacity. This is the simplest version of the multicommodity routing problem. Here we consider an extension with ℓ_p-norm capacity constraints on subsets of edges. This can be used to model situations where edges are provided by multiple agents. Each agent j owns a subset S_j of edges and it requires the ℓ_{p_j}-norm of the bandwidths of these edges to be at most c_j. In this section we assume the S_j are disjoint. Our result also applies to general S_j via a reduction to disjoint instances.

▶ **Theorem 7.** *Assume that $c_j = \Omega(\log m) \cdot |S_j|^{1/p_j}$ for each j. Then there is a randomized $O(\log m)$-competitive online algorithm for throughput maximization with ℓ_p-norm capacities, where m is the number of edges in the graph.*

We note that a similar "high capacity" assumption is also needed in the linear special case [4, 13] where each $|S_j| = 1$.

In a fractional version of the problem, a request can be satisfied by several paths and the allocation of bandwidth can be in range $[0, 1]$ instead of being restricted from $\{0, 1\}$. For request (s_i, t_i), let $\mathcal{P}_i$ be the set of simple paths between s_i and t_i. Variable $f_{i,P}$ is defined to be the amount of flow on the path P for request (s_i, t_i). The total profit of algorithm is the (fractional) number of requests served and the performance is measured with respect to the maximum number of requests that could be served if the requests are known beforehand. We describe the problem as a packing problem:

$$\max \quad \sum_i \sum_{P \in \mathcal{P}_i} f_{i,P} \tag{6}$$

$$\text{s.t.} \quad \sum_{P \in \mathcal{P}_i} f_{i,P} \leq 1, \qquad \forall i \tag{7}$$

$$\sum_i \sum_{P \in \mathcal{P}_i : e \in P} f_{i,P} = \mu_e, \qquad \forall e \tag{8}$$

$$\|\mu(S_j)\|_{p_j} \leq c_j, \qquad \forall j \tag{9}$$

$$f \in \mathbb{R}^n_+$$

Note that single edge capacity is a special case of (9) with $|S_j| = 1$ and any p_j. The corresponding primal problem is the following.

$$\min \quad \sum_j c_j \|x(S_j)\|_{q_j} + \sum_i z_i \tag{10}$$

$$\text{s.t.} \quad z_i + \sum_{e \in P} x_e \geq 1, \qquad \forall i,\ P \in \mathcal{P}_i \tag{11}$$

$$x, z \in \mathbb{R}^n_+$$

where z is the dual variable of constraints (7) and x is the dual variable of constraints (8) to (9). This formulation falls in the form of (1). Combining our algorithm with online rounding techniques, we can prove Theorem 7. A formal proof appears in the full version.

References

1 Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. A general approach to online network optimization problems. *ACM Transactions on Algorithms*, 2(4):640–660, 2006.

2 Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. The online set cover problem. *SIAM J. Comput.*, 39(2):361–370, 2009.

3 Baruch Awerbuch, Yossi Azar, Edward F. Grove, Ming-Yang Kao, P. Krishnan, and Jeffrey Scott Vitter. Load balancing in the $\mathrm{l_p}$ norm. In *FOCS*, pages 383–391, 1995.

4 Baruch Awerbuch, Yossi Azar, and Serge Plotkin. Throughput-competitive on-line routing. In *FOCS*, pages 32–40. IEEE, 1993.

5 Yossi Azar, Umang Bhaskar, Lisa K. Fleischer, and Debmalya Panigrahi. Online mixed packing and covering. In *SODA*, 2013.

6 Yossi Azar, Niv Buchbinder, T.-H. Hubert Chan, Shahar Chen, Ilan Reuven Cohen, Anupam Gupta, Zhiyi Huang, Ning Kang, Viswanath Nagarajan, Joseph Naor, and Debmalya Panigrahi. Online algorithms for covering and packing problems with convex objectives. In *FOCS*, pages 148–157, 2016.

7 Yossi Azar, Ilan Reuven Cohen, and Debmalya Panigrahi. Online covering with convex objectives and applications. *CoRR*, abs/1412.3507, 2014. URL: `http://arxiv.org/abs/1412.3507`.

8 Nikhil Bansal, Niv Buchbinder, and Joseph Naor. Randomized competitive algorithms for generalized caching. *SIAM J. Comput.*, 41(2):391–414, 2012.

9 Nikhil Bansal and Kirk Pruhs. Server scheduling to balance priorities, fairness, and average quality of service. *SIAM J. Comput.*, 39(7):3311–3335, 2010.

10 Avrim Blum, Anupam Gupta, Yishay Mansour, and Ankit Sharma. Welfare and profit maximization with production costs. In *FOCS*, pages 77–86, 2011.

11 Niv Buchbinder, Shahar Chen, Anupam Gupta, Viswanath Nagarajan, and Joseph Naor. Online packing and covering framework with convex objectives. *CoRR*, abs/1412.8347, 2014.

12 Niv Buchbinder, Kamal Jain, and Joseph Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*, pages 253–264, 2007.

13 Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Math. Oper. Res.*, 34(2):270–286, 2009.

14 Niv Buchbinder and Joseph (Seffi) Naor. The design of competitive online algorithms via a primal-dual approach. *Found. Trends Theor. Comput. Sci.*, 3(2-3):93–263, 2007.

15 T.-H. Hubert Chan, Zhiyi Huang, and Ning Kang. Online convex covering and packing problems. *CoRR*, abs/1502.01802, 2015.

16 Moses Charikar and Adriana Karagiozova. On non-uniform multicommodity buy-at-bulk network design. In *STOC*, pages 176–182. ACM, 2005.

17 Chandra Chekuri, Mohammad Taghi Hajiaghayi, Guy Kortsarz, and Mohammad R. Salavatipour. Approximation algorithms for nonuniform buy-at-bulk network design. *SIAM J. Comput.*, 39(5):1772–1798, 2010.

18 Nikhil R. Devanur and Zhiyi Huang. Primal dual gives almost optimal energy efficient online algorithms. In *SODA*, pages 1123–1140, 2014.

19 Nikhil R. Devanur and Kamal Jain. Online matching with concave returns. In *STOC*, pages 137–144. ACM, 2012.

20 Noa Elad, Satyen Kale, and Joseph (Seffi) Naor. Online semidefinite programming. In *ICALP*, pages 40:1–40:13, 2016.

21 Alina Ene, Deeparnab Chakrabarty, Ravishankar Krishnaswamy, and Debmalya Panigrahi. Online buy-at-bulk network design. In *FOCS*, pages 545–562, 2015.

22 Anupam Gupta, Ravishankar Krishnaswamy, and Kirk Pruhs. Online primal-dual for non-linear optimization with applications to speed scaling. In *WAOA*, pages 173–186, 2012.

23 Anupam Gupta and Viswanath Nagarajan. Approximating sparse covering integer programs online. *Math. Oper. Res.*, 39(4):998–1011, 2014.

24 Zhiyi Huang and Anthony Kim. Welfare maximization with production costs: A primal dual approach. In *SODA*, pages 59–72, 2015.

25 Ishai Menache and Mohit Singh. Online caching with convex costs: Extended abstract. In *SPAA*, pages 46–54, 2015.