

A Simple Boosting Framework for Transshipment

Goran Zuzic   

Google Research, Zürich, Switzerland

ETH Zürich, Switzerland

Abstract

Transshipment is an important generalization of both the shortest path problem and the optimal transport problem. The task asks to route a demand using a flow of minimum cost over (uncapacitated) edges. Transshipment has recently received extensive attention in theoretical computer science as it is the centerpiece of all modern theoretical breakthroughs in parallel and distributed (approximate) shortest-path computation, a classic and well-studied problem.

The key advantage of transshipment over shortest paths is the so-called *boosting* property: one can often boost a crude approximate solution to a (near-optimal) $(1 + \varepsilon)$ -approximate solution. However, our understanding of this phenomenon is limited: it is not clear which approximators can be boosted. Moreover, all current boosting frameworks are built with a specific type of approximator in mind and are relatively complicated.

The main takeaway of our paper is conceptual: *any* black-box oracle that computes an approximate *dual* solution can be boosted to an $(1 + \varepsilon)$ -approximator. This decouples and simplifies all known near-optimal $(1 + \varepsilon)$ -approximate transshipment and shortest paths results: they all (implicitly) construct approximate dual solutions and boost them.

We provide a very simple analysis based on the multiplicative weights framework. Furthermore, to keep the paper completely self-contained, we provide a new (and arguably much simpler) analysis of multiplicative weights that leverages well-known optimization tools to bypass the ad-hoc calculations used in the standard analyses.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Theory of computation → Mathematical optimization; Theory of computation → Mixed discrete-continuous optimization; Theory of computation → Shortest paths; Theory of computation → Parallel algorithms

Keywords and phrases mixed continuous-discrete optimization, boosting, multiplicative weights, theoretical computer science, shortest path, parallel algorithms

Digital Object Identifier 10.4230/LIPIcs.ESA.2023.104

Related Version *Full Version:* <https://arxiv.org/abs/2110.11723>

Funding *Goran Zuzic:* Work presented in this paper was partially performed while at ETH Zürich. Supported in part by NSF grants CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808, a Sloan Research Fellowship, and funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 949272).

Acknowledgements The author would like to thank Bernhard Haeupler, Patrik Pavic, and Richard Peng for helpful discussions about the paper. The author would also like to thank the anonymous reviewers for their helpful suggestions that significantly improved the quality of the paper.

1 Introduction

Suppose we are given a weighted graph $G = (V, E)$ and a *demand vector* $d \in \mathbb{R}^V$ satisfying $\sum_{v \in V} d(v) = 0$, where $d(v)$ denotes the number of units of some (single) commodity that are either available (if $d(v) > 0$) or required (if $d(v) < 0$) at the node v . *Transshipment* asks to distribute the available units of the commodity until it perfectly matches the requirement. The goal is to minimize the total cost of movement, where moving a single unit over an edge e has a cost of $w(e)$ ($w(e) \geq 0$ is the weight of e).



© Goran Zuzic;

licensed under Creative Commons License CC-BY 4.0

31st Annual European Symposium on Algorithms (ESA 2023).

Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 104; pp. 104:1–104:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

104:2 A Simple Boosting Framework for Transshipment

Transshipment is a strong generalization of multiple problems, including the $s - t$ shortest path problem, optimal transport, or the assignment problem on metric spaces.

► **Example 1.** Given two nodes s, t in a weighted graph G , the shortest path can be modeled as transshipment for the demand $d(v) := \mathbb{I}[v = s] - \mathbb{I}[v = t]$.

► **Example 2.** Let (V, d) be a metric space and let $A, B \subseteq V$ be two disjoint subsets. The minimum cost perfect matching between A and B (the so-called assignment problem [14]) can be modeled as transshipment on the complete bipartite graph $(A \cup B, A \times B)$ with weights $w(\{a, b\}) = d(a, b)$ and the demand vector $d(v) := \mathbb{I}[v \in A] - \mathbb{I}[v \in B]$. This also models the Wasserstein distance (also known as the earth mover's distance or optimal transport [18]).

Perhaps surprisingly, transshipment has proven to be extremely useful for uncovering the distance structure (i.e., shortest paths) of a graph. Indeed, the problem has been the centerpiece of all near-optimal approaches to the single-source shortest path problem in the parallel and distributed settings [5, 10, 2, 19, 15].

The key property that differentiates transshipment from other similar problems like shortest path is the so-called *boosting* property – one can boost a crude, say $\text{poly}(\log n)$ -approximate solution, to a near-optimal $(1 + \varepsilon)$ -approximate solution. This conceptually reduces $(1 + \varepsilon)$ -transshipment (and shortest path computation) to approximate transshipment. However, not all approximators can be boosted and a more principled understanding of which approaches are susceptible to boosting is required. Moreover, the current boosting algorithms are coupled together with the specific approximators they use, making them non-modular, complicated, and hard to reuse.

The main takeaway of our paper is conceptual: *any* black-box oracle that computes a α -approximate *dual* solution can be boosted to a $(1 + \varepsilon)$ -approximate dual solution. This significantly simplifies current transshipment results by decoupling them into two independent questions: (1) how to obtain an approximate dual solution (which is often model-specific), and (2) how to boost this approximate solution (which can be reused). The scope of this paper is to develop a simple framework for the latter question.

We provide a very simple algorithm and analysis based on the *multiplicative weights framework*. Furthermore, to keep the paper completely self-contained, we provide a new (and arguably much simpler) analysis of multiplicative weights that leverages well-known optimization tools to bypass the ad-hoc calculations used in the standard analyses. (deferred to the full version).

We now provide several examples that show how prior approaches all (implicitly) construct approximate dual and then boost them.

► **Example 3.** Sherman [17] gave the first sequential almost-linear¹ $(1 + \varepsilon)$ -transshipment algorithm. The main insight was the construction of a so-called *linear cost approximator*, which is a linear operator R (i.e., matrix) such that $\|Rd\|_1$ approximates the optimal transshipment cost in the sense that $\text{OPT}(d) \leq \|Rd\|_1 \leq n^{o(1)}\text{OPT}(d)$ for all demands $d \in \mathbb{R}^V$. Their paper uses linear cost approximators with subgradient descent to show one can obtain a $(1 + \varepsilon)$ -approximate solution. We provide a conceptual decoupling and reinterpretation of their paper: one can use any linear cost approximator R to directly obtain an approximate dual solution, which can, in turn, be boosted to an $(1 + \varepsilon)$ -approximate solution via our framework.

¹ We refer to $m\text{poly}(\log n)$ as near-linear and $m^{1+o(1)}$ as almost-linear runtimes.

► **Example 4.** Haeupler and Li [8] solve $n^{o(1)}$ -approximate transshipment in the distributed setting and leave the possibility of boosting to an $(1 + \varepsilon)$ -approximation as the main open problem, which would have yielded important consequences in the distributed setting. Our paper provides a partial explanation to why their approach was not susceptible to boosting: their approach, based on low-stretch spanning trees, only computes a *primal* solution (i.e., an approximate flow), whereas a dual solution is required. A dual-based solution was later recently developed by Rozhon et al. [15].

► **Example 5.** Other successful $(1 + \varepsilon)$ -transshipment approaches either approximate the solution by solving the original problem on a spanner [5], or by constructing a linear cost approximator on an emulator of the original graph [10, 2] (an emulator of G is a graph H whose distance structure multiplicatively approximates the one of G , see Section 4.1). We show all of these approaches can be reinterpreted as obtaining an approximate dual solution.

Comparison with Becker et al. [5]. The paper contributed the first polylog-competitive existentially-optimal shortest path algorithm in the distributed setting (up to $\tilde{O}(1)$ -factors). Crucially, they develop a boosting framework for transshipment which, similar to this paper, uses an approximate dual solver to construct a near-optimal solution. The main drawbacks of their solver are that (1) the analysis of [5] is quite involved, stemming from it being based on projected gradient descent, and (2) as written, the interface of the [5] solver relies on solving a modified version of transshipment which is harder to interpret and work with than the original one. As stated in the journal version of [5], their interface can be significantly simplified (by working with projections), but this degrades the runtime to have an α^4 -dependency w.r.t. the approximation quality α of the approximator (we provide an α^2 -dependency) and requires non-explicit modifications to the solver that might be difficult for non-experts.² On the other hand, the approach presented in our paper has several drawbacks compared to [5], such as: (1) our solver requires a guess on the optimal solution, which is obtained using binary search, while their solver does not need adapting the internal parameters during the optimization process, and (2) our dual-only solver needs to perform extra steps to return a feasible primal solution. However, independent of the drawbacks, we believe the user-friendly interface, better runtime, and a simpler analysis make our conceptual contribution worthwhile.

Potential impact. Ultimately, we hope that this paper will encourage an ongoing effort to simplify deep algorithmic results that use continuous optimization tools. Such an effort would potentially yield a dual benefit: it would both lower the barrier to entry for newcomers by conceptually simplifying the current approaches, as well as help to transfer the modern theoretically-optimal algorithms into real-world state-of-the-art by allowing practitioners to independently combine the theoretical ideas with the many heuristics necessary for an algorithm to perform well in practice.


Organization of the paper. We present a model-oblivious boosting framework for transshipment in Section 3 and apply it in Section 4 to simplify previous results. These applications are loosely grouped by the method of computing the approximate dual solution: Section 4.1 presents results when the approximate solution is computed on a spanner or emulator (i.e.,

² Specifically, they require the returned dual solution be orthogonal to the demand vector. However, as they note in the journal version of their paper, this issue can be mitigated by projecting a general solution to the space of vectors orthogonal to the demand: this can be shown to work with a loss in approximation factor and time complexity if one appropriately initializes the solution.

on graphs that approximate the original metric). Section 4.2 presents results that compute the dual solution via (aforementioned) linear cost approximators. Finally, the full version of the paper gives a simple and self-contained analysis of multiplicative weights.

2 Preliminaries

Graph Notation. Let $G = (V, E)$ be an undirected graph and let $n := |V|, m := |E|$. It is often convenient to direct E consistently. For simplicity and without loss of generality, we assume that $V = \{v_1, v_2, \dots, v_n\}$ and define $\vec{E} = \{(v_i, v_j) \mid (v_i, v_j) \in E, i < j\}$. We identify E and \vec{E} by the obvious bijection. We chose this orientation for simplicity and concreteness: arbitrarily changing the orientations does not influence the results (if done consistently). We denote with $B \in \{-1, 0, 1\}^{V \times \vec{E}}$ the node-edge incidence matrix of G , which for any $v \in V$ and $e = (s, t) \in \vec{E}$ assigns $B_{s,e} = 1$, $B_{t,e} = -1$, and $B_{u,e} = 0$ when $u \notin \{s, t\}$. A weight or length function w assigns each edge $e \in \vec{E}$ a weight $w(e) > 0$. The weight function can also be interpreted as a diagonal *weight matrix* $W \in \mathbb{R}_{\geq 0}^{\vec{E} \times \vec{E}}$ which assigns $W_{e,e} = w(e) \geq 1$ for any $e \in \vec{E}$ (and 0 on all off-diagonal entries). In this paper, it is often more convenient to specify weighted graphs via $G \cong (B, W)$, i.e., by specifying its matrices B and W as defined above.



■ **Figure 1** A simple graph G and its corresponding node-edge incidence matrix B .


Flows and Transshipment (TS). A *demand* is a $d \in \mathbb{R}^V$. We say a demand is *proper* if $\sum_{v \in V} d_v = 0$. A *flow* is a vector $f \in \mathbb{R}^{\vec{E}}$, where $f(\vec{e}) > 0$ if flowing in the direction of the arc \vec{e} and negative otherwise. A flow f *routes* a demand d if $Bf = d$. It is easy to see only proper demands are routed by flows. The cost of a flow f is $\|Wf\|_1$. For a weighted graph G and a given proper demand d the *transshipment problem* (or TS, for short) asks to find a flow f_d^* of minimum-cost among flows that route d . In other words, the tuple (B, W, d) specifies a transshipment instance. When the underlying graph $G \cong (B, W)$ is clear from the context, we define $\|d\|_{\text{OPT}} := \|Wf_d^*\|_1$ to denote the cost of the optimal flow for routing demand d . The transshipment problem naturally admits the following LP formulation and its dual. The primal asks us to optimize over all *flows* $f \in \mathbb{R}^{\vec{E}}$, while the dual asks us to optimize over all vectors $\phi : \mathbb{R}^V$, which we refer to as *potentials*.

$$\text{Primal: } \min \|Wf\|_1 : Bf = d, \quad \text{Dual: } \max \langle d, \phi \rangle : \|W^{-1}B^T\phi\|_\infty \leq 1. \quad (1)$$

Scalar products are denoted as $\langle x, y \rangle = x^T \cdot y$. Finally, we assume the weights and demands are polynomially-bounded, hence $\|d\|_{\text{OPT}} \leq n^{O(1)}$. Any feasible primal and dual values provide an upper and lower bound on $\|d\|_{\text{OPT}}$, formally stated in the following well-known result.

► **Fact 6.** Let $f \in \mathbb{R}^{\vec{E}}$ and $\phi \in \mathbb{R}^V$ be any feasible primal and dual solution: i.e., $Bf = d$ and $\|W^{-1}B^T\phi\|_\infty \leq 1$. Then $\langle d, \phi \rangle \leq \|d\|_{\text{OPT}} \leq \|Wf\|_1$.

For example, consider the s - t shortest path subproblem where $d(v) = \mathbb{I}[v = s] - \mathbb{I}[v = t]$ specifies the demand. One optimal solution to the primal/dual pair is to set $f(\vec{e})$ to 1 iff \vec{e} is on some fixed shortest path from s to t ; $\phi(v)$ is set to the distance in G from $t \in V$. Note that in this case the primal and dual objectives are equal, and correspond to the weight of the shortest path from s to t .



■ **Figure 2** A example transshipment graph with its exact solution. The original graph is unit weight $w_e = 1$ and undirected. The demand d is non-zero at four nodes. The optimal primal flow f is depicted in blue and is non-zero for four edges. One of many optimal vectors ϕ is depicted in red. The optimal value of the solution is $\text{OPT} = 4$.

Asymptotic Notation. We use \tilde{O} to hide polylogarithmic factors in n , i.e., $\tilde{O}(1) = \text{polylog } n$.

Algorithmic model and basic vector operations. To facilitate both simplicity and generality, we specify our algorithms using high-level operations. Specifically, in a unit operation, we can perform the following so-called **basic vector operations**: (1) assign vectors in \mathbb{R}^n or \mathbb{R}^m to variables, (2) add two (vector) variables together, (3) apply any scalar function $\lambda : \mathbb{R} \rightarrow \mathbb{R}$ to each component of a vector separately, and (4) compute matrix-vector products with matrices B, B^T, W , and W^{-1} . Note that each basic vector operation can be near-optimally compiled into standard parallel/distributed models. In PRAM: each operation can be performed in $\tilde{O}(1)$ depth and near-linear work. In the standard distributed model of computation CONGEST [12] basic vector operations can be computed in a single round of distributed computation (where the variables are stored in the obvious distributed fashion).

Multiplicative weights (MW) framework is a powerful meta-algorithm that allows for (among other things) solving various optimization tasks by reducing them to simpler (so-called “linearized”) versions of the original problem [3]. For the purposes of this paper, we define the following pair of tasks.

► **Definition 7.** Let $\gamma \in \mathbb{R}$ be a scalar, $A \in \mathbb{R}^{m \times n}$ be a matrix, and $b \in \mathbb{R}^n$ be a vector. We define the following **Feasibility task**:

$$\exists x \in \mathbb{R}^n \mid \|Ax\|_\infty + \langle b, x \rangle \leq \gamma. \tag{2}$$

Given additionally a vector $p \in \mathbb{R}^m$ satisfying $\|p\|_1 \leq 1$ and an accuracy parameter $\varepsilon > 0$, we also define the following **Linearized task**:

$$\exists x \in \mathbb{R}^n \mid \langle p, Ax \rangle + \langle b, x \rangle \leq \gamma - \varepsilon. \tag{3}$$

Note that if the Feasibility task (Equation (2)) is feasible, then the Linearized task (Equation (3)) is also feasible for every p (satisfying $\|p\|_1 \leq 1$) and every $\varepsilon > 0$.

104:6 A Simple Boosting Framework for Transshipment

Suppose we want to solve some fixed Feasibility task (Equation (2)) and assume we know how to solve the accompanying (typically much easier!) Linearized task (Equation (3), for any p and ε) via some *black-box Oracle*. Then, there exists a simple algorithm that computes a solution to the Feasibility task by repeatedly querying the Oracle with different values of p that satisfy $\|p\|_1 \leq 1$ (the accuracy ε stays fixed); the oracle is assumed to return a feasible solution x for each queried Linearized task.

We define the **width of the Oracle** $\rho > 0$ to be (any upper bound on) the largest width of a solution $\|Ax\|_\infty$ that can be returned by the Oracle, i.e., $\rho \geq \|Ax\|_\infty$ during the course of the algorithm. Oracles with larger widths need to be queried more times, hence we aim to construct Oracles with their width being as small as possible. The following Theorem 8 and Algorithm 1 give a solver for the Feasibility task (Equation (2)) assuming the Oracle. We defer the proof to the full version of the paper.

► **Theorem 8.** *Let (A, b, γ) be a feasible Feasibility task (Equation (2)) and fix $\varepsilon > 0$. Suppose we have access to an Oracle that will solve the accompanying Linearized task Equation (3) specified by $(A, b, \gamma, p, \varepsilon)$ for any $\|p\|_1 \leq 1$. Then, Algorithm 1 constructs a feasible solution for Equation (2) and queries the Oracle at most $4\varepsilon^{-2}\rho^2 \ln(2m)$ times. Here, $\rho > 0$ is the width of the Oracle.*

■ **Algorithm 1** Solver for the Feasibility task using an oracle for the Linearized task.

-
1. **Input:** Feasibility task $(A \in \mathbb{R}^{m \times n}, b \in \mathbb{R}^n, \gamma \in R)$ and $\varepsilon > 0$.
 2. Initialize $x_* \leftarrow \vec{0} \in \mathbb{R}^n$ and $\beta := \varepsilon/(2\rho^2)$.
 3. For $t = 1, \dots, T$ rounds, where $T := 4\varepsilon^{-2}\rho^2 \ln(2m)$:
 - a. Let $q \leftarrow \begin{bmatrix} A \\ -A \end{bmatrix} x_* \in \mathbb{R}^{2m}$.
 - b. Let $q'_i \leftarrow \exp(\beta q_i)_i$ for $i \in [2m]$.
 - c. Let $p_t \leftarrow (1/\sum_{i=1}^{2m} q'_i)(q'_i - q'_{i+m})$. (Normalization and flattening.)
 - d. MW outputs $p_t \in \mathbb{R}^m$ to Oracle. (Note that $\|p_t\|_1 \leq 1$.)
 - e. Oracle returns a solution $x_t \in \mathbb{R}^n$ to the Linearized task w.r.t. p_t . (ρ must be set large enough such that $\|Ax\|_\infty \leq \rho$.)
 - f. We update $x_* \leftarrow x_* + x_t$.
 4. MW outputs $(1/T) \cdot x_* \in \mathbb{R}^n$.
-

3 A Boosting Framework for Transshipment

We describe how to compute an $(1 + \varepsilon)$ -approximate solution for transshipment given only a black-box oracle which computes an α -approximate dual solution. This oracle is called the dual-only α -approximator (where the dual is defined as in Equation (1)).

► **Definition 9** (Dual-Only Approximator). *Let $G \cong (B, W)$ be a weighted graph. A dual-only α -approximator for transshipment over G is an oracle which, given any proper demand $d \in \mathbb{R}^V$, outputs a dual solution $\phi \in \mathbb{R}^V$ satisfying the following properties:*

- (Dual feasibility) $\|W^{-1}B^T\phi\|_\infty \leq 1$.
- (Approximation guarantee): $\langle d, \phi \rangle \geq \frac{1}{\alpha} \|d\|_{\text{OPT}}$.

Note that, directly from its definition, a dual-only α -approximator can be used to obtain an α -approximate value of the solution.

We typically want $\text{poly}(\log n)$ -approximators or $n^{o(1)}$ -approximators that can be constructed and queried efficiently. However, for pedagogical purposes, we first show that the minimum spanning tree (MST) is a non-trivial n -approximator.

► **Example 10 (MST).** Let $d \in \mathbb{R}^V$ be an arbitrary proper demand. The MST T of G can be used as a simple n -approximator for transshipment. First, root the MST T in an arbitrary $r \in V$ and assume without loss of generality (up to re-orientation of edges in E) that all edges point from parent to child nodes in this rooted tree. Next, let f_T be the unique flow supported on T that perfectly routes d . We now define $\phi \in \mathbb{R}^V$ by saying $\phi(r) := 0$ and proceeding in a top-to-bottom order. For each parent-child tree-edge $e = (p, c)$, set $\phi(c) := \phi(p) - \text{sign}(f_T(e))w(e)$. Now, by construction, we have $\|Wf_T\|_1 = \langle \phi, d \rangle$ (decompose the flow f_T into a positive combination of oriented paths such that all of them cross each tree edge with the same orientation; the claim is true for each one of them and, therefore, for their sum). Furthermore, $\|W^{-1}B^T\phi\|_\infty \leq n$ by the following argument: consider each edge $e = \{u, v\} \in E_G$ and consider the unique $u - v$ tree-path. This path is composed of at most n edges, and each one of them have weight at most $w(e)$ (since T is the MST). Hence, $|\phi(u) - \phi(v)| \leq w(e) \cdot n$, which is equivalent to the claim above. Finally, defining $\phi_* := \frac{1}{n}\phi$ we get a feasible dual solution that is n -approximate: Using Fact 6 we have $\|d\|_{\text{OPT}} \leq \|Wf_T\|_1 = \langle \phi, d \rangle = n \cdot \langle \phi_*, d \rangle$, as required.

We now show the central claim of our framework: given a dual-only α -approximator we can leverage the multiplicative weights framework (Definition 7) to provide feasible potentials (i.e., a dual solution) $\phi \in \mathbb{R}^V$ that are $(1 + \varepsilon)$ -approximate, i.e., $\langle d, \phi \rangle \geq \frac{1}{1+\varepsilon} \|d\|_{\text{OPT}}$. The existence of the boosting procedure is formalized in Lemma 11, while the explicit algorithm is deferred to the full version of the paper.

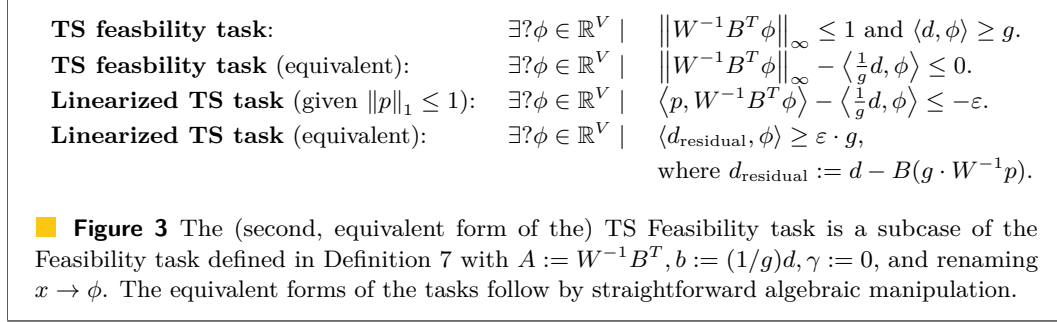
As an assumption to simplify the exposition, we can safely assume that we know the value of $g := \|d\|_{\text{OPT}}$ (up to a multiplicative $1 + \varepsilon$) as this value can be “guessed” via a standard binary search since our method will either certify that $\|d\|_{\text{OPT}} \leq (1 + \varepsilon)g$ (telling us our guess g is too low), or will otherwise construct a feasible solution ϕ with $\langle d, \phi \rangle \geq g$ (telling us we can increase our guess).

► **Lemma 11.** *Let (B, W, d) be a transshipment instance and let $\varepsilon > 0$. Given any $g \geq 0$, and any dual-only α -approximator, there is a $4\varepsilon^{-2}\alpha^2 \ln(2m)$ -round algorithm that, in each round, queries the approximator once and performs $O(1)$ basic vector operations. At termination, the algorithm either:*

1. *outputs (feasible) potentials $\phi_* \in \mathbb{R}^V$ satisfying $\|W^{-1}B^T\phi_*\|_\infty \leq 1$ and $\langle d, \phi_* \rangle \geq g$, or,*
2. *detects that $\|d\|_{\text{OPT}} \leq (1 + \varepsilon)g$; indeed, it outputs an (infeasible) flow $f_* \in \mathbb{R}^E$ satisfying $\|Wf_*\|_1 \leq g$ and $\|d - Bf_*\|_{\text{OPT}} \leq \varepsilon g$.*

Remark. If one is only concerned about finding an $(1 + \varepsilon)$ -approximate dual solution, one can completely ignore the infeasible flow that is being outputted and simply use the fact that the second condition guarantees $\|d\|_{\text{OPT}} \leq (1 + \varepsilon)g$, which is sufficient for binary search. Furthermore, we note that any such (infeasible) flow f_* satisfying the above properties guarantees $\|d\|_{\text{OPT}} \leq (1 + \varepsilon)g$ by the following argument. First, by definition of $\|d - Bf_*\|_{\text{OPT}} \leq \varepsilon g$, there exists a flow f_{residual} that routes $d - Bf_*$ and has cost $\|Wf_{\text{residual}}\|_1 \leq \varepsilon g$. Then, the flow $f_* + f_{\text{residual}}$ routes demand d (since $Bf_* + Bf_{\text{residual}} = Bf_* + d - Bf_* = d$) and has cost at most $\|Wf_*\|_1 + \|Wf_{\text{residual}}\|_1 \leq g + \varepsilon g = (1 + \varepsilon)g$, implying that $\|d\|_{\text{OPT}} \leq (1 + \varepsilon)g$.

Proof. First, finding potentials $\phi \in \mathbb{R}^V$ satisfying $\|W^{-1}B^T\phi\|_\infty \leq 1$ and $\langle d, \phi \rangle \geq g$ is equivalent to finding potentials $\exists \phi \in \mathbb{R}^V \mid \|W^{-1}B^T\phi\|_\infty - \langle \frac{1}{g}d, \phi \rangle \leq 0$ (one direction is immediate, the other direction follows by the fact that we can scale ϕ such that $\langle d, \phi \rangle = g$). Therefore, it is sufficient to solve the following so-called TS Feasibility task (see Figure 3).



We apply the MW framework by interpreting the TS Feasibility task as a Feasibility Task in the sense of Definition 7, solving it via Algorithm 1 where we have to implement the Oracle.

First, we note that the TS Feasibility task directly corresponds to a $(A := W^{-1}B^T, b := -\frac{1}{g}d, \gamma := 0)$ -Feasibility task. We aim to implement the Oracle for the corresponding Linearized TS task with a small width ρ . To recap, ρ is the maximum value of $\|W^{-1}B^T\phi\|_\infty$ ever returned by the Oracle – we later determine that setting $\rho := \alpha$ suffices.

The Oracle, upon receiving p by Algorithm 1, queries the dual-only α -approximator with the (so-called) residual demand $d_{\text{residual}} := d - B(g \cdot W^{-1}p)$. Intuitively, we can interpret p , or more specifically $g \cdot W^{-1}p$, as the “current” iterate of the final flow solution. Specifically, $\|g \cdot W^{-1}p\|_1 \leq g$, i.e., it has a small cost since g is a guess for OPT. If the residual demand can be routed with a small cost of at most εg (which can be estimated via the approximator), it means that $\|d_{\text{residual}}\| \leq \varepsilon g$, hence $f_* := g \cdot W^{-1}p$ satisfies the second output condition.

The approximator, being asked to route d_{residual} , returns the α -approximate feasible dual, i.e., a vector $\phi_{\text{residual}} \in \mathbb{R}^V$ satisfying $\langle d_{\text{residual}}, \phi_{\text{residual}} \rangle \geq \frac{1}{\alpha} \|d_{\text{residual}}\|_{\text{OPT}}$ and $\|W^{-1}B^T\phi_{\text{residual}}\|_\infty \leq 1$. The Oracle outputs $\phi' := \alpha \cdot \phi_{\text{residual}}$. Note that the width of the oracle is exactly $\|W^{-1}B\phi'\|_\infty = |\alpha| \|W^{-1}B\phi_{\text{residual}}\|_\infty \leq \alpha \cdot 1 = \alpha$.

Either $\langle d_{\text{residual}}, \phi' \rangle \geq \varepsilon \cdot g$, and the Oracle successfully solves the Linearized TS task by returning ϕ_{residual} , in which case the MW loop continues. If this is always the case, Algorithm 1 outputs ϕ_* satisfying $\|W^{-1}B^T\phi_*\|_\infty \leq 1$ and $\langle d, \phi_* \rangle \geq g$, as required. Regarding the width of the solution, we have that $\|W^{-1}B^T\phi'\|_\infty \leq \alpha$, hence setting $\rho := \alpha$ suffices, leading to at most $4\varepsilon^{-2}\alpha^2 \ln 2m$ rounds of the algorithm.

On the other hand, if this is (ever) not the case, we say the Oracle fails. In this case, at the moment of failure, we define $f_* := g \cdot W^{-1}p$ and have that $\langle d_{\text{residual}}, \phi' \rangle \leq \varepsilon \cdot g$. Since $\phi' = \alpha \phi_{\text{residual}}$, we have $\langle d_{\text{residual}}, \phi_{\text{residual}} \rangle \leq \frac{\varepsilon}{\alpha} \cdot g$. Since ϕ_{residual} is an α -approximate dual w.r.t. d_{residual} , we have that $\|d - B(g \cdot W^{-1}p)\|_{\text{OPT}} = \|d_{\text{residual}}\|_{\text{OPT}} \leq \alpha \cdot \frac{\varepsilon}{\alpha} g = \varepsilon \cdot g$. Therefore, f_* satisfies the second condition and we are done. ◀

The full algorithm is deferred to the full version.

Reducing the residual error of the primal. While the booster of Lemma 11 returns a feasible $(1 + \varepsilon)$ -approximate dual solution, it does not return a feasible primal solution (i.e., flow). However, this issue can be resolved by repeatedly routing the residual demand $d - Bf_*$

until the cost of routing the residual demand drops to an insignificant $1/\text{poly}(n)$ -fraction of the original cost, at which point any trivial reparation scheme suffices (like routing along the MST). See the full version of the paper for more details. Combining the above result with binary searching the guess g and with the residual error reduction (but without the model-specific trivial routing), we immediately yield the following result (proof deferred to the full version).

► **Corollary 12.** *Let (B, W, d) be a transshipment instance. Given any $1/2 \geq \varepsilon > 0$, $C > 0$ and dual-only α -approximator, there is an $\tilde{O}(C \cdot \varepsilon^{-2} \alpha^2)$ -round algorithm computing (both):*

- a feasible dual ϕ_* satisfying $(1 + \varepsilon)^{-1} \|d\|_{\text{OPT}} \leq \langle d, \phi \rangle \leq \|d\|_{\text{OPT}}$, and,
- an (infeasible) primal f_* satisfying $\|Wf\|_1 \leq (1 + \varepsilon) \|d\|_{\text{OPT}}$ and $\|d - Bf_*\|_{\text{OPT}} \leq n^{-C} \|d\|_{\text{OPT}}$.

In each round, the algorithm performs $O(1)$ basic vector operations and queries to the approximator.

4 Applications

In this section, we show how to apply the boosting framework of Section 3 to simplify and decouple several landmark results in the parallel and distributed settings. First, we describe results that approximate transshipment by solving it on a compact graph representation called a spanner or emulator (Section 4.1). Then, we describe results that use linear cost approximators (Section 4.2).

4.1 Approximating via spanners and emulators

A β -approximate **emulator** of a graph $G = (V, E_G)$ is a weighted graph $H = (V, E_H)$ on the same vertex set where the distances are approximated with a distortion of β ; i.e., $\text{dist}_G(u, v) \leq \text{dist}_H(u, v) \leq \beta \cdot \text{dist}_G(u, v)$ for all $u, v \in V$. A spanner is simply an emulator that is a subgraph of G , i.e., where $E_H \subseteq E_G$, making it particularly well-studied in some settings.

Approximating with emulators is conceptually straightforward: faced with a transshipment instance on G , we (approximately) solve the instance on H , which yields an approximate solution on G . This is captured by the following result.

► **Theorem 13.** *Let H be a β -approximate emulator of G . Any dual-only α -approximator on H is a dual-only $(\alpha \cdot \beta)$ -approximator on G .*

Proof. Fix a demand d on G . Querying the H -approximator, we obtain a dual solution ϕ_H satisfying $\|W_H^{-1} B_H^T \phi_H\|_\infty \leq 1$; we also know an accompanying primal solution f_H exists with $W_H(f_H) \leq \alpha \langle d, \phi_H \rangle$.

Primal solution. We construct a flow f_G in G as follows. For each edge $e \in E_H$ we know, due to $\text{dist}_G(u, v) \leq \text{dist}_H(u, v)$, that there exists a path in G of length at most $w_H(e)$; we add $f_H(e)$ amount of flow along this path. It is easy to check that, f_G routes d (i.e., $B_G f_G = d$, hence it is feasible) and that $W_G(f_G) \leq W_H(f_H)$, hence $\|d\|_{\text{OPT}(G)} \leq \alpha \langle d, \phi_H \rangle$.

Dual solution. Let $\phi_G := \frac{1}{\beta} \phi_H$. Note that $\|d\|_{\text{OPT}(G)} \leq (\alpha\beta) \cdot \langle d, \phi_G \rangle$, hence it is sufficient to deduce $\|W_G^{-1} B_G^T \phi_G\|_\infty \leq 1$. Since ϕ_H is feasible in H , we have for each $e' = \{u', v'\} \in E_H$ that $(B_H^T \phi_H)_{e'} = |\phi_H(u') - \phi_H(v')| \leq \frac{|\phi_H(u') - \phi_H(v')|}{\beta} = \frac{w_H(u', v')}{\beta}$. Fix an edge $e = \{u, v\} \in E_G$; since $\text{dist}_H(u, v) \leq \beta \cdot \text{dist}_G(u, v)$ there exists a path

104:10 A Simple Boosting Framework for Transshipment

($u = p'_0, p'_1, p'_2, \dots, p'_k = v$) in H of length at most $\beta \cdot w_G(e)$. Therefore, we can deduce that $\|W_G^{-1}B_G^T\phi_G\|_\infty \leq 1$ in the following way: $|(W_G^{-1}B_G^T\phi_G)_e| = \frac{|\phi_G(u) - \phi_G(v)|}{w_G(e)} \leq \frac{\sum_{i=1}^T |\phi_G(p'_{i-1}) - \phi_G(p'_i)|}{w_G(e)} \leq \frac{\sum_{i=1}^T w_H(p'_{i-1}, p'_i)}{\beta w_G(e)} \leq \frac{\beta w_G(e)}{\beta w_G(e)} = 1.$ ◀

Remark. There are a few immediate extensions to the above proof. Given a primal-dual approximator (one that returns both a primal and a dual) on a *spanner*, we can immediately obtain a primal-dual approximator on G since the returned primal f_H is also a feasible primal in G . A similar property holds for emulators, but one would need to provide a mapping that embeds each edge $e \in E_G$ into (paths of) H that are of length at most $\beta \cdot w(e)$ in order to construct the flow f_G on G .

Application: TS in Broadcast congested clique [5]. Using algorithms from prior work, a Broadcast congested clique can compute an $\tilde{O}(1)$ -approximate Baswana-Sen [4] spanner H in $\tilde{O}(1)$ rounds. The edges of such a spanner are naturally partitioned into n parts of size $\tilde{O}(1)$, where each part is associated with a unique node, and that node knows the edges in its part. Therefore, the spanner can be made global knowledge in $\tilde{O}(1)$ rounds using broadcasts. Therefore, each node can solve a transshipment instance on H , providing an $\tilde{O}(1)$ -approximator for the original graph via Theorem 13, culminating in an $\tilde{O}(\varepsilon^{-2})$ -round solution for $(1 + \varepsilon)$ -transshipment.

Application: existentially-optimal SSSP in Broadcast CONGEST [5]. Consider the single-source shortest path (SSSP) problem where each node wants to compute $(1 + \varepsilon)$ -approximate from some source $s \in V$. From prior work, we can compute an *overlay graph* $G' = (V', E')$ where $V' \subseteq V$ and $|V'| = \tilde{O}(\varepsilon^{-1}\sqrt{n})$ such that the SSSP task on G reduces to SSSP on G' , and G' can be computed in $\tilde{O}(D + \varepsilon^{-1}\sqrt{n})$ rounds. As was shown in [5], an SSSP instance can be solved by solving $\tilde{O}(1)$ transshipment instances (the details are non-trivial and out of scope of this paper), hence the problem reduces to solving TS on G' . However, any T -round Broadcast congested clique algorithm can be simulated on G' in $T \cdot O(D + |V'|) = T \cdot \tilde{O}(D + \varepsilon^{-1}\sqrt{n})$ rounds of Broadcast CONGEST: we simulate a single round by constructing a BFS tree on G (of depth $O(D)$ and in $O(D)$ rounds), and then pipelining all $|V'|$ messages (that are to be broadcasted in the current round) to the root and then down to all other nodes, taking $O(D + |V'|)$ rounds in Broadcast CONGEST per round of Broadcast congested clique. Combining with the Broadcast congested clique result, we obtain an $\tilde{O}(\varepsilon^3)(D + \sqrt{n})$ -round algorithm.

Application: near-optimal TS in PRAM [2]. The paper introduces a concept called *low-hop emulator* $H = (V, E_H)$ of $G = (V, E)$ satisfying (i) H is an $\tilde{O}(1)$ -approximate emulator of G , (ii) $|E_H| = \tilde{O}(n)$, and (iii) $\text{dist}_H^{O(\log \log n)}(u, v) = \text{dist}_G(u, v)$, i.e., every (exact) shortest path in H has at most $O(\log \log n)$ hops (edges). Moreover, low-hop emulators can be computed in PRAM in $\tilde{O}(1)$ depth and $\tilde{O}(m)$ work. Low hop emulators are particularly useful since Property (iii) implies that one can compute (exact) SSSP on them in $\tilde{O}(1)$ depth and $\tilde{O}(n)$ work (e.g., using $O(\log \log n)$ rounds of Bellman-Ford). The ability to compute exact SSSP enables each node of H to be embedded into ℓ_1 space of dimension $\tilde{O}(1)$ with (worst-case) distortion $\tilde{O}(1)$ (via so-called Bourgain's embedding [6] via $\tilde{O}(1)$ SSSP oracle calls). Since H is an emulator of G , the same embedding is an $\tilde{O}(1)$ -distortion embedding of G . Using Theorem 13, this reduces $(1 + \varepsilon)$ -TS to finding an $\tilde{O}(1)$ -approximator in ℓ_1 space. This can be done in $\tilde{O}(1)$ depth and $\tilde{O}(n)$ work using linear cost approximators (explained in Section 4.2) by utilizing the so-called *randomly shifted grids* method [9]. This culminates in an $\tilde{O}(\varepsilon^{-2})$ depth and $\tilde{O}(\varepsilon^{-2}m)$ work $(1 + \varepsilon)$ -transshipment algorithm.

4.2 Approximating by linear cost approximators

A particularly successful type of approximator for transshipment has been the linear cost approximator. The successes of such an approximator include the first $m^{1+o(1)}$ algorithm for transshipment in the centralized model [17] and the first $\tilde{O}(m)$ -work and $\tilde{O}(1)$ -depth parallel shortest path algorithm [2, 10].

► **Definition 14.** *An α -approximate linear cost approximator for a weighted graph G is a $k \times n$ matrix P , such that, for any proper demand d it holds that $\|d\|_{\text{OPT}} \leq \|Pd\|_1 \leq \alpha \|d\|_{\text{OPT}}$.*

Our insight is that one can immediately convert a linear cost approximator P to a dual-only approximator. Note that the sign function is applied entry-wise to a vector.

► **Theorem 15.** *Let P be an α -approximate linear cost approximator. Consider the function $\phi(d)$ that maps a demand d to $\phi(d) := \frac{1}{\alpha} P^T \text{sign}(Pd)$. Then, ϕ is a dual-only α -approximator.*

Proof. Let $G \cong (B, W)$ be the underlying graph. First, we show that the following subclaim about a linear-algebraic guarantee that characterizes P : we have that $\|yPBW^{-1}\|_\infty \leq \alpha$ over all $\|y\|_\infty \leq 1$. Specifically, for each oriented edge $\vec{e} \in \vec{E}$, consider how P approximates the cost of routing a unit from the head to the tail of \vec{e} . Formally, we define the demand $d_{\vec{e}}$ to be $d_{\vec{e}}(x) := \mathbb{I}[x = s] - \mathbb{I}[x = t]$ for an edge $\vec{e} = (s, t) \in \vec{E}$. Clearly, $\|d_{\vec{e}}\|_{\text{OPT}} \leq w(e)$, hence it is necessary that $\|Pd_{\vec{e}}w(e)^{-1}\|_1 \leq \alpha$. Furthermore, it is easy to see that the columns of B are exactly $d_{\vec{e}}$ over all $\vec{e} \in \vec{E}$, hence each column of PBW^{-1} has ℓ_1 -norm at most α . This is equivalent to $\|y^T PBW^{-1}\|_\infty \leq \alpha$ over all $\|y\|_\infty \leq 1$. This proves the subclaim.

We now prove the complete result. Let $y := \text{sign}(Pd)$ and $\phi(d) := \frac{1}{\alpha} P^T y$. Since, $\|d\|_{\text{OPT}} \leq \|Pd\|_1$, there must exist a flow f satisfying d such that $\|d\|_{\text{OPT}} \leq \|Wf\|_1 \leq \|Pd\|_1$. We verify all properties Definition 9. (Dual feasibility) $\|W^{-1}B^T\phi(d)\|_\infty = \frac{1}{\alpha} \|W^{-1}B^T P^T y\|_\infty \leq \frac{1}{\alpha} \cdot \alpha = 1$ via the subclaim. (Approximation guarantee) $\langle d, \phi(d) \rangle = \frac{1}{\alpha} \langle Pd, y \rangle = \frac{1}{\alpha} \langle Pd, \text{sign}(Pd) \rangle = \frac{1}{\alpha} \|Pd\|_1 \geq \frac{1}{\alpha} \|d\|_{\text{OPT}}$. ◀

Having a dual-only α -approximator that can be evaluated in M time, we construct (via Corollary 12) an $\tilde{O}(\varepsilon^{-2}\alpha^2 \cdot M)$ time $(1 + \varepsilon)$ -approximate algorithm for transshipment.

► **Corollary 16.** *Let P be an α -approximate linear cost approximator on a weighted graph G and suppose that we can evaluate matrix-vector multiplications with P and P^T (and other basic vector operations) in M time. Given any TS instance, there is an $\tilde{O}(\varepsilon^{-2}\alpha^2 M)$ -time algorithm that computes a $(1 + \varepsilon)$ -approximate primal-dual pair (f, ϕ) satisfying the properties listed in Corollary 12.*

Application: almost-optimal sequential TS [17]. The goal is to construct $\varepsilon^{-2}m^{1+o(1)}$ -time $(1 + \varepsilon)$ -TS solver in the sequential setting. Following Corollary 16, it is sufficient to construct a $n^{o(1)}$ -approximate linear cost approximator P , which is accomplished as follows. Each vertex of a weighted graph G is embedded into ℓ_1 space of dimension $O(\log^2 n)$ with (worst-case) distortion $O(\log n)$ (via so-called Bourgain's embedding [6] in $\tilde{O}(m)$ sequential time). Then, the dimension of the embedding is reduced to $d := O(\sqrt{\log n})$ via a simple Johnson-Lindenstrauss projection [7], increasing the distortion of the embedding to $\exp(O(d)) = n^{o(1)}$. Finally, the paper constructs a $O(\log^{1.5} n)$ -approximate linear cost approximator in this (virtual) ℓ_1 space of dimension d that can be evaluated efficiently, leading to a $\exp(O(d)) \cdot O(\log^{1.5} n) = n^{o(1)}$ -approximate linear cost approximator in G , which yields the result. **Approximator in ℓ_1 space:** We give a short cursory description on how to construct the approximator P . Re-scale and round the ℓ_1 space such that all coordinates are integral.

104:12 A Simple Boosting Framework for Transshipment

Then, each point x calculates the distance $c(x)$ to the closest point with all-even coordinates. Then, x uniformly spreads its demand $d(x)$ among all points with all-even coordinates that are of distance exactly $c(x)$ to x . Finally, repeat the algorithm on points with all-even coordinates (delete other points, divide all coordinates by 2). After $O(\log n)$ iterations, the entire remaining demand will be supported on 2^d vertices of the hypercube, which can be routed to a common vertex yielding a $O(d)$ approximation. It can be shown that the cost incurred by spreading the demand at any particular step $O(d)$ -approximates the optimal solution, and that the optimal solution does not increase in-between two steps, leading to a $O(d \log n) = O(\log^{1.5} n)$ -approximate linear cost approximator. **Efficiency:** Evaluating the approximator requires computing the demands at each step in the above algorithm. Evaluating even the first step requires $n2^d$ time since each point x sends its demand to (potentially) $2^d = n^{o(1)}$ closest all-even points. Therefore, the dimension of the embedding is reduced to $O(\sqrt{\log n})$. Moreover, the paper (implicitly) claims this approximator in ℓ_1 can be evaluated in $m^{1+o(1)}$ time. Finally, we remark that the approximator does not yield a flow in the original graph in any meaningful way, (i.e., it only approximates costs), confirming that it is dual-only. Together, we solve $(1 + \varepsilon)$ -TS in $\varepsilon^{-2}m^{1+o(1)}$ time.

Application: near-optimal TS in PRAM [10]. The goal is to solve $(1 + \varepsilon)$ -TS in $\tilde{O}(1)$ depth and $\tilde{O}(m)$ work in PRAM. The paper constructs an $\tilde{O}(1)$ -approximate linear cost approximator P with sparsity $\tilde{O}(m)$, meaning it can be evaluated in $\tilde{O}(1)$ depth and $\tilde{O}(m)$ work, which would yield the result. To do so, the paper follows [17] by embedding G in ℓ_1 space with distortion $\tilde{O}(1)$ and dimension $d := \tilde{O}(1)$ and then uses the randomly shifted grids methods of [9] to approximate the cost in this virtual space. **Approximator in ℓ_1 space:** We define a *randomly shifted grid* of scale W to be the set $W(\mathbb{Z}^d + u) \subseteq \mathbb{R}^d$, where each coordinate of $u \in \mathbb{R}^d$ is uniformly drawn from $[0, 1)$ (i.e., one obtains a randomly shifted grid by taking all integral d -dimensional points, randomly translating them along each axis, then multiplying all coordinates by W). Initially, set $W \leftarrow \tilde{O}(1)$. The routing works by sampling $s := \tilde{O}(1)$ randomly shifted grids of scale W and, for each grid, each point x sends $1/s$ of its demand $d(x)$ to the closest point in the grid. The scale W is increased by a polylogarithmic factor and the algorithm is repeated for $O(\log n)$ steps until all demand is supported on a hypercube, at which point it can be $O(d)$ -approximated by aggregating it at a single vertex. It can be shown that the cost incurred by routing the demand at any particular step $\tilde{O}(1)$ -approximates the optimal solution, and that the optimal solution increases only by a multiplicative $1 + 1/\text{poly}(\log n)$ factor, hence after $O(\log n)$ iterations we obtain an $\tilde{O}(1)$ -approximate linear cost approximator P that has sparsity $\tilde{O}(m)$. **Vertex reduction framework:** On its face, the above approach simply shows that in order to get $(1 + \varepsilon)$ -transshipment (and $(1 + \varepsilon)$ -shortest paths, as arduously shown in the paper), it is sufficient to find an $\tilde{O}(1)$ -distortion ℓ_1 -embedding. However, to find an ℓ_1 -embedding, one needs $\tilde{O}(1)$ -approx shortest paths (with some additional technical requirements concerning the violation of the triangle inequality). To resolve this cycle, the paper goes through the vertex reduction framework of [11, 13] which reduces the number of vertices by a polylogarithmic factor, recursively solves transshipment, lifts the solution to the original graph, and repairs it using the boosting framework. The details are out-of-scope.

Future work. The ideas used for solving transshipment have historically paralleled the ideas used for solving maximum flow problems. Adding to the connection between these two problems, approximate solutions to maximum flow can also be boosted in a similar way to transshipment [16] via linear cost approximators (called *congestion approximators*). However,

no framework that can handle black-box approximators has been developed – creating such a framework would conceptually simplify the task of designing approximate maximum flow solutions. Furthermore, both transshipment and maximum flow are special cases of the so-called ℓ_p -norm flow, which also seems to support boosting [1].

References

- 1 Deeksha Adil, Rasmus Kyng, Richard Peng, and Sushant Sachdeva. Iterative refinement for ℓ_p -norm regression. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1405–1424. SIAM, 2019.
- 2 Alexandr Andoni, Clifford Stein, and Peilin Zhong. Parallel approximate undirected shortest paths via low hop emulators. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 322–335. ACM, 2020.
- 3 Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory Comput.*, 8(1):121–164, 2012.
- 4 Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Structures & Algorithms*, 30(4):532–563, 2007.
- 5 Ruben Becker, Andreas Karrenbauer, Sebastian Krinninger, and Christoph Lenzen. Near-Optimal Approximate Shortest Paths and Transshipment in Distributed and Streaming Models. In *31st International Symposium on Distributed Computing (DISC)*, volume 91, pages 7:1–7:16, 2017.
- 6 Jean Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985.
- 7 Sanjoy Dasgupta and Anupam Gupta. An elementary proof of the johnson-lindenstrauss lemma. *International Computer Science Institute, Technical Report*, 22(1):1–5, 1999.
- 8 Bernhard Haeupler and Jason Li. Faster distributed shortest path approximations via shortcuts. *arXiv preprint*, 2018. [arXiv:1802.03671](https://arxiv.org/abs/1802.03671).
- 9 Piotr Indyk and Nitin Thaper. Fast image retrieval via embeddings. In *3rd international workshop on statistical and computational theories of vision*, volume 2(3), page 5, 2003.
- 10 Jason Li. Faster parallel algorithm for approximate shortest path. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 308–321. ACM, 2020.
- 11 Aleksander Madry. Fast approximation algorithms for cut-based problems in undirected graphs. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 245–254, 2010.
- 12 David Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.
- 13 Richard Peng. Approximate undirected maximum flows in $O(m \text{polylog}(n))$ time. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1862–1867, 2016.
- 14 Lyle Ramshaw and Robert E Tarjan. On minimum-cost assignments in unbalanced bipartite graphs. *HP Labs, Palo Alto, CA, USA, Tech. Rep. HPL-2012-40R1*, 20, 2012.
- 15 Václav Rozhoň, Christoph Grunau, Bernhard Haeupler, Goran Zuzic, and Jason Li. Undirected $(1+\epsilon)$ -shortest paths via minor-aggregates: Near-optimal deterministic parallel & distributed algorithms. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*, 2022.
- 16 Jonah Sherman. Nearly maximum flows in nearly linear time. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 263–269, 2013.

104:14 A Simple Boosting Framework for Transshipment

- 17 Jonah Sherman. Generalized preconditioning and undirected minimum-cost flow. In *Proceedings of the 2017 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 772–780, 2017.
- 18 Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- 19 Goran Zuzic, Goramoz Goranci, Mingquan Ye, Bernhard Haeupler, and Xiaorui Sun. Universally-optimal distributed shortest paths and transshipment via graph-based l1-oblivious routing. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2022.