Domain decomposition for entropy regularized optimal transport

Mauro Bonafini¹ · Bernhard Schmitzer¹

Received: 31 January 2020 / Revised: 12 March 2021 / Accepted: 23 May 2021 /

Published online: 19 November 2021

© The Author(s) 2021

Abstract

We study Benamou’s domain decomposition algorithm for optimal transport in the entropy regularized setting. The key observation is that the regularized variant converges to the globally optimal solution under very mild assumptions. We prove linear convergence of the algorithm with respect to the Kullback–Leibler divergence and illustrate the (potentially very slow) rates with numerical examples. On problems with sufficient geometric structure (such as Wasserstein distances between images) we expect much faster convergence. We then discuss important aspects of a computationally efficient implementation, such as adaptive sparsity, a coarse-to-fine scheme and parallelization, paving the way to numerically solving large-scale optimal transport problems. We demonstrate efficient numerical performance for computing the Wasserstein-2 distance between 2D images and observe that, even without parallelization, domain decomposition compares favorably to applying a single efficient implementation of the Sinkhorn algorithm in terms of runtime, memory and solution quality.

Mathematics Subject Classification 65K10 · 49M27

1 Introduction

1.1 Motivation

(Computational) optimal transport Optimal transport is a fundamental optimization problem with applications in various branches of mathematics. Let μ and ν be prob-

✉ Bernhard Schmitzer
schmitzer@cs.uni-goettingen.de

Mauro Bonafini
bonafini@cs.uni-goettingen.de

¹ Campus Institute Data Science, University of Göttingen, Göttingen, Germany

ability measures over spaces X and Y and let $\Pi(\mu, \nu)$ be the set of transport plans, i.e. probability measures on $X \times Y$ with μ and ν as first and second marginal. Further, let $c : X \times Y \rightarrow \mathbb{R}$ be a *cost function*. The Kantorovich formulation of optimal transport is then given by

$$\inf \left\{ \int_{X \times Y} c(x, y) \, d\pi(x, y) \mid \pi \in \Pi(\mu, \nu) \right\}. \quad (1.1)$$

We refer to the monographs [21, 26] for a thorough introduction and historical context. Due to its geometric intuition and robustness it is becoming particularly popular in image and data analysis. Therefore, one main challenge is the development of efficient numerical methods, which has seen immense progress in recent years, such as solvers for the Monge–Ampère equation [3], semi-discrete methods [13, 16], entropic regularization [9], and multi-scale methods [18, 24]. An introduction to computational optimal transport, an overview on available efficient algorithms, and applications can be found in [19].

Domain decomposition Benamou introduced a domain decomposition algorithm for Wasserstein-2 optimal transport on $\mathbb{R}^d$ [2], based on Brenier’s polar factorization [4]. In the language of (1.1) the algorithm works as follows: Let X and Y be subsets of $\mathbb{R}^d$ and $c(x, y) = \|x - y\|^2$ the squared Euclidean distance. Let now (X_1, X_2, X_3) be a partition of X and write $X_{\{1,2\}} := X_1 \cup X_2$, $X_{\{2,3\}} := X_2 \cup X_3$ and let $\pi^{(0)} \in \Pi(\mu, \nu)$ be some initial feasible transport plan. The next iterate $\pi^{(1)}$ is then obtained by optimizing $\pi^{(0)}$ on the set $X_{\{1,2\}} \times Y$, while keeping it fixed on $X_3 \times Y$. Next, $\pi^{(2)}$ is obtained by optimizing $\pi^{(1)}$ on $X_{\{2,3\}} \times Y$, while keeping it fixed on $X_1 \times Y$. These two steps are then repeated. This is illustrated (for entropic optimal transport) in Fig. 1.

In [2] it is shown that this algorithm converges to the globally optimal solution when μ is Lebesgue absolutely continuous and (X_1, X_2, X_3) satisfy a ‘convex overlap’ condition, which roughly states that ‘if a function is convex on $X_{\{1,2\}}$ and $X_{\{2,3\}}$, then it must be convex on X ’. Further, it is shown that this method can easily be extended to more than three cells, thus leading to a parallelizable algorithm. Unfortunately, it is also demonstrated that the discretized method does not converge to a global minimizer (for $d > 1$). When the discretization is refined, it is shown that the optimal solution is recovered in the limit. But this requires an increasing number of discretization points in the three sets (X_1, X_2, X_3) and thus numerically obtaining a good approximation of the globally optimal solution remains challenging with this method.

1.2 Outline and contribution

Entropic domain decomposition The key observation in this article is that for entropic optimal transport the domain decomposition algorithm converges to the unique globally optimal solution for arbitrary (measurable) bounded costs c , as soon as $\mu(X_2) > 0$. In particular, this covers the case where X and Y are discrete and finite. Therefore, entropic smoothing is not only a useful tool in solving the optimal transport sub-problems arising in domain decomposition, it also facilitates conver-

Fig. 1 Iterates of the domain decomposition algorithm for entropic optimal transport in one dimension (yellow represents high, blue low mass density). $\mu = \nu$ is the uniform probability measure on $X = Y = [0, 1]$, equipped with cost function $c(x, y) = \|x - y\|^2$. The optimal coupling is a blurred ‘diagonal’ measure (approximately equal to $\pi^{(5)}$). We initialize with an ‘anti-diagonal’ feasible $\pi^{(0)}$, i.e. the order must be approximately reversed. X is divided into three cells (dashed lines, first panel) and $\pi^{(\ell)}$ is alternately optimized over $X_{\{1,2\}}$ and $X_{\{2,3\}}$. Iterates $\pi^{(1)}$ to $\pi^{(5)}$ were obtained by optimizing the previous iterate over the red rectangle (color figure online)

gence of the whole algorithm. While the convergence rate can be very slow in suitable worst-case examples, we show empirically that it is fast on problems with sufficient geometric structure, thus leading to a practical and efficient parallel numerical method for large-scale problems.

Preliminaries and definition of the algorithm We start by establishing the basic mathematical setting and notation in Section 2.1. Some basic facts about entropic optimal transport are recalled in Section 2.2. In Sect. 3 the entropic domain decomposition algorithm is defined and a simple convergence proof is sketched.

Theoretical and empirical worst-case convergence rate Section 4 is dedicated to proving that the iterates converge to the unique optimal solution linearly with respect to the Kullback–Leibler divergence. We first give a proof for the three-cell example discussed above (Sect. 4.1) and then generalize it to more general decompositions (Sect. 4.2). Our results apply to arbitrary (measurable) bounded cost functions c , marginal measures μ and ν , and we make minimal assumptions on the decomposition structure (it must be ‘connected’ in a suitable sense, Definition 7, in particular convergence only requires $\mu(X_2) > 0$ in the three-cell example). As the entropic regularization parameter tends to zero, our rate bound tends to one exponentially. This is reminiscent of the convergence rate for the standard Sinkhorn algorithm established in [10] with respect to Hilbert’s projective metric. While we do not expect that our bounds on the convergence rate are tight, we demonstrate in Sect. 4.3 with some numerical (near) worst-case examples that they capture the qualitative behaviour of the algorithm on difficult problems.

Relation to parallel sorting The slow convergence rates from Sect. 4 do not suggest that the algorithm is efficient. But these slow rates rely on maliciously designed counter-examples. In practice usually much more geometric structure is available, similar to the setting originally considered by Benamou [2]. While a detailed analysis of the convergence rate in these cases is beyond the scope of this article, to provide some intuition, we discuss in Sect. 5 the relation of the domain decomposition algorithm for optimal transport with the odd-even transposition parallel sorting algorithm

[14, Exercise 37] in one dimension. This algorithm is known to converge in $O(N)$ iterations where N is the number of cells that the domain is partitioned into.

Practical implementation and geometric large-scale examples In Sect. 6 we provide a practical version of the domain decomposition algorithm, leading to an efficient numerical method for large scale problems with geometric structure. We discuss how the memory footprint can be reduced, how one can handle approximate solutions to the sub-problems as obtained by the Sinkhorn algorithm, and how to combine the algorithm with the ε -scaling heuristic and a multi-scale scheme, see [23].

Numerical experiments and comparison with single Sinkhorn algorithm The efficiency of the implementation is then demonstrated numerically by computing the Wasserstein-2 distance between images in Sect. 7. We report accurate approximate primal and dual solutions with low entropic regularization artifacts after a logarithmic number (w.r.t. the marginal size) of domain decomposition iterations. On a standard desktop computer with 6 cores a high-quality approximate optimal transport between two mega-pixel images is computed in approximately 4 minutes.

A comparison to a single Sinkhorn algorithm, when applied to the full problem, is given in Sect. 7.4. We find that the domain decomposition approach has several key advantages. First, it allows more extensive parallelization. The standard Sinkhorn algorithm is already parallel in the sense that each half-iteration consists essentially of a matrix-vector product that can be parallelized. But the results of this operation must be communicated to all workers after each half-iteration, thus only allowing local parallelization such as a single GPU. In the domain decomposition variant, small sub-problems are solved independently and their results must only be communicated upon completion, thus allowing parallelization over multiple machines.

Second, domain decomposition allows to reduce the memory footprint. The naive Sinkhorn algorithm requires storage of the full kernel matrix, the size of which grows quadratically with the marginal size. This can be avoided by using efficient heuristics, such as Gaussian convolutions or pre-factored heat kernels [25] but these methods only work in particular settings and with sufficiently high regularization. Alternatively, with adaptive kernel truncation [23] the memory demand is approximately linear in the marginal size. But to retain numerical stability, the truncation parameter may not be chosen too aggressively, thus still making memory demand a practical constraint. In the domain decomposition variant, only the kernel matrices for the sub-problems that are currently being re-solved are needed; for all other sub-problems only information about their marginals is kept. In addition, via parallelization this may also be distributed over several computers.

In our numerical experiments we find that in comparable runtime (even without parallelization) the domain decomposition algorithm obtains more accurate primal and dual iterates by requiring only approximately 11% of the memory of the single Sinkhorn solver. Dealing only with small sub-problems at a time also makes it easier to handle numerical delicacies associated with small entropic regularization parameters. The domain decomposition method therefore more reliably solves large problems.

2 Background

2.1 Setting and notation

- X and Y are compact metric spaces. We assume compactness to avoid overly technical arguments while covering the numerically relevant setting.
- For a compact metric space Z the set of finite signed Radon measures over Z is denoted by $\mathcal{M}(Z)$. The subsets of non-negative and probability measures are denoted by $\mathcal{M}_+(Z)$ and $\mathcal{M}_1(Z)$. The Radon norm of $\mu \in \mathcal{M}(Z)$ is denoted by $\|\mu\|_{\mathcal{M}(Z)}$ and one has $\|\mu\|_{\mathcal{M}(Z)} = \mu(Z)$ for $\mu \in \mathcal{M}_+(Z)$. We often simply write $\|\mu\|$ for the norm when the meaning is clear from context.
- For $\mu \in \mathcal{M}_+(Z)$ we denote by $L^1(Z, \mu)$, $L^\infty(Z, \mu)$ the usual function spaces and add a subscript $+$ to denote the subsets of μ -a.e. non-negative functions. The corresponding norms are denoted by $\|\cdot\|_{L^1(Z, \mu)}$ and $\|\cdot\|_{L^\infty(Z, \mu)}$ but we often merely write $\|\cdot\|_1$ and $\|\cdot\|_\infty$ (or even $\|\cdot\|$ for the latter) when space and measure are clear from context.
- For $\mu \in \mathcal{M}_+(Z)$ and a measurable function $u : Z \rightarrow \mathbb{R}$ we denote by $\langle \mu, u \rangle$ the integration of u with respect to μ , and define the normalized integral

$$\int_Z u \, d\mu := \frac{1}{\mu(Z)} \int_Z u \, d\mu = \frac{\langle \mu, u \rangle}{\mu(Z)}. \quad (2.1)$$

- For $\mu \in \mathcal{M}_+(Z)$ and measurable $S \subset Z$ the restriction of μ to S is denoted by $\mu \llcorner S$. Similarly, for a measurable function $u : Z \rightarrow \mathbb{R}$, set $u \llcorner S(x) = u(x)$ if $x \in S$, and 0 otherwise.
- The maps $P_X : \mathcal{M}_+(X \times Y) \rightarrow \mathcal{M}_+(X)$ and $P_Y : \mathcal{M}_+(X \times Y) \rightarrow \mathcal{M}_+(Y)$ denote the projections of measures on $X \times Y$ to their marginals, i.e.

$$(P_X \pi)(S_X) := \pi(S_X \times Y) \quad \text{and} \quad (P_Y \pi)(S_Y) := \pi(X \times S_Y)$$

for $\pi \in \mathcal{M}_+(X \times Y)$, $S_X \subset X$, $S_Y \subset Y$ measurable.

- For (measurable) functions $a : X \rightarrow \mathbb{R} \cup \{\infty\}$, $b : Y \rightarrow \mathbb{R} \cup \{\infty\}$, the functions $a \oplus b, a \otimes b : X \times Y \rightarrow \mathbb{R} \cup \{\infty\}$ are given by

$$(a \oplus b)(x, y) = a(x) + b(y), \quad (a \otimes b)(x, y) = a(x) \cdot b(y).$$

For two measures $\mu \in \mathcal{M}_+(X)$, $\nu \in \mathcal{M}_+(Y)$ their product measure on $X \times Y$ is denoted by $\mu \otimes \nu$.

2.2 Entropic optimal transport

We recall in this Section the entropic regularization approach to optimal transport and collect the main existence and characterization results which will be needed in this article.

Definition 1 (*Kullback–Leibler divergence*) Let Z be a compact metric space. For $\mu \in \mathcal{M}(Z)$, $\nu \in \mathcal{M}_+(Z)$ the *Kullback–Leibler divergence* (or relative entropy) of μ w.r.t. ν is given by

$$\text{KL}(\mu|\nu) := \begin{cases} \int_Z \varphi\left(\frac{d\mu}{d\nu}\right) d\nu & \text{if } \mu \ll \nu, \mu \geq 0, \\ +\infty & \text{else,} \end{cases}$$

$$\text{with } \varphi(s) := \begin{cases} s \log(s) - s + 1 & \text{if } s > 0, \\ 1 & \text{if } s = 0, \\ +\infty & \text{else.} \end{cases}$$

The Fenchel–Legendre conjugate of φ is given by $\varphi^*(z) = \exp(z) - 1$.

Definition 2 (*Entropy regularized optimal transport*) Let $\mu \in \mathcal{M}_1(X)$, $\nu \in \mathcal{M}_1(Y)$ and $\hat{\mu} \in \mathcal{M}_+(X)$ and $\hat{\nu} \in \mathcal{M}_+(Y)$ such that

$$\|\hat{\mu}\| = \|\hat{\nu}\| > 0, \quad \hat{\mu} \ll \mu, \quad \hat{\nu} \ll \nu, \quad \frac{d\hat{\mu}}{d\mu} \in L^\infty(X, \mu), \quad \frac{d\hat{\nu}}{d\nu} \in L^\infty(Y, \nu).$$

Similar to above, denote by

$$\Pi(\hat{\mu}, \hat{\nu}) := \{\pi \in \mathcal{M}_+(X \times Y) \mid P_X \pi = \hat{\mu}, P_Y \pi = \hat{\nu}\}$$

the set of *transport plans* between $\hat{\mu}$ and $\hat{\nu}$. The condition $\|\hat{\mu}\| = \|\hat{\nu}\|$ ensures that the set is non-empty. For a lower-semicontinuous cost function $c \in L_+^\infty(X \times Y, \mu \otimes \nu)$ the Kantorovich optimal transport problem between $\hat{\mu}$ and $\hat{\nu}$ is then given by

$$\inf \left\{ \int_{X \times Y} c(x, y) d\pi(x, y) \mid \pi \in \Pi(\hat{\mu}, \hat{\nu}) \right\}. \quad (2.2)$$

Existence of minimizers in this setting is provided, for instance, by [26, Theorem 4.1]. For a regularization parameter $\varepsilon > 0$ define

$$k := \exp(-c/\varepsilon) \in L_+^\infty(X \times Y, \mu \otimes \nu), \quad \text{and} \quad K := k \cdot (\mu \otimes \nu). \quad (2.3)$$

In this article we will frequently exploit that k is bounded by $0 < \exp(-\|c\|_\infty)/\varepsilon \leq k \leq 1$. Then the entropic optimal transport problem, regularized with respect to $\mu \otimes \nu$ is given by

$$\inf \{ \varepsilon \text{KL}(\pi|K) \mid \pi \in \Pi(\hat{\mu}, \hat{\nu}) \}. \quad (2.4)$$

Problem (2.4) can be solved (approximately) with the Sinkhorn algorithm (Remark 1). In addition, in our article, entropic smoothing will ensure convergence of the domain decomposition scheme. We refer to [19, Chapter 4] for an overview on entropic optimal transport and its numerical implications. The (Γ) -convergence of (2.4) to (2.2) as $\varepsilon \rightarrow 0$ has been shown in [6, 7, 15] in various settings and with different strategies.

We now collect some results about entropic optimal transport required in this article. A proof is given in the “Appendix”.

Proposition 1 (Optimal entropic transport couplings)

- (i) (2.4) has a unique minimizer $\pi^* \in \Pi(\hat{\mu}, \hat{\nu})$.
- (ii) There exist measurable $u^* : X \rightarrow \mathbb{R}_+$, $v^* : Y \rightarrow \mathbb{R}_+$ such that $\pi^* = u^* \otimes v^* \cdot K$. u^* and v^* are unique μ -a.e. and ν -a.e. up to a positive re-scaling $(u^*, v^*) \rightarrow (\lambda \cdot u^*, \lambda^{-1} \cdot v^*)$ for $\lambda > 0$.
- (iii) The couple (u^*, v^*) satisfies the integral equations

$$u^*(x) \int_Y k(x, y') v^*(y') d\nu(y') = \frac{d\hat{\mu}}{d\mu}(x) \quad \text{and} \quad v^*(y) \int_X k(x', y) u^*(x') d\mu(x') = \frac{d\hat{\nu}}{d\nu}(y) \quad (2.5)$$

for μ -a.e. $x \in X$ and ν -a.e. $y \in Y$.

- (iv) $u^* \in L_+^\infty(X, \mu)$, $v^* \in L_+^\infty(Y, \nu)$, $\|u^*\|_{L^1(X, \mu)} > 0$, $\|v^*\|_{L^1(Y, \nu)} > 0$,

$$\frac{1}{\|v^*\|_1} \frac{d\hat{\mu}}{d\mu}(x) \leq u^*(x) \leq \frac{\exp(\|c\|_\infty/\varepsilon)}{\|v^*\|_1} \frac{d\hat{\mu}}{d\mu}(x) \quad \text{for } \mu\text{-a.e. } x \in X, \quad (2.6)$$

$$\frac{1}{\|u^*\|_1} \frac{d\hat{\nu}}{d\nu}(y) \leq v^*(y) \leq \frac{\exp(\|c\|_\infty/\varepsilon)}{\|u^*\|_1} \frac{d\hat{\nu}}{d\nu}(y) \quad \text{for } \nu\text{-a.e. } y \in Y \quad (2.7)$$

and $\log u^* \in L^1(X, \hat{\mu})$, $\log v^* \in L^1(Y, \hat{\nu})$.

Definition 3 (Dual problem) In addition to (2.4), we will consider a corresponding dual problem which, for the purpose of this article, is best stated as

$$\sup \{J(u, v) \mid u \in L_+^\infty(X, \mu), v \in L_+^\infty(Y, \nu)\} \quad (2.8)$$

with $J : L_+^\infty(X, \mu) \times L_+^\infty(Y, \nu) \mapsto [-\infty, \infty)$ given by

$$J : (u, v) \mapsto \int_X \log u d\hat{\mu} + \int_Y \log v d\hat{\nu} - \int_{X \times Y} u \otimes v dK + \|K\|. \quad (2.9)$$

This is a reparametrization of the usual dual problem where one uses $\varepsilon \log u$ and $\varepsilon \log v$ as variables instead (cf. [23]). Again, we briefly collect a few helpful results. A proof is given in the “Appendix”.

Proposition 2 (Duality for regularized optimal transport)

- (i) The functional J , (2.9), is well defined.
- (ii) For $u \in L_+^\infty(X, \mu)$, $v \in L_+^\infty(Y, \nu)$, $\pi \in \Pi(\hat{\mu}, \hat{\nu})$ one has $J(u, v) \leq \text{KL}(\pi | K)$.
- (iii) A transport plan $\pi^* \in \Pi(\hat{\mu}, \hat{\nu})$ of the form $\pi^* = u^* \otimes v^* \cdot K$ for $u^* \in L_+^\infty(X, \mu)$, $v^* \in L_+^\infty(Y, \nu)$ is optimal for (2.4). In that case, (u^*, v^*) are optimal for (2.8). Maximizers for (2.8) exist.

Remark 1 (*Sinkhorn algorithm*) Problems (2.4) and (2.8) can be solved (approximately) with the Sinkhorn algorithm. For some initial $u^{(0)} \in L_+^\infty(X, \mu)$ (not identical to zero), it is given for $\ell = 0, 1, 2, \dots$ by

$$v^{(\ell)}(y) := \frac{\frac{d\hat{\nu}}{d\nu}(y)}{\int_X k(x, y) u^{(\ell)}(x) d\mu(x)}, \quad u^{(\ell+1)}(x) := \frac{\frac{d\hat{\mu}}{d\mu}(x)}{\int_Y k(x, y) v^{(\ell)}(y) d\nu(y)}.$$

We refer to these two steps as Y - and X -iteration. It is quickly verified that a Y -iteration corresponds to a partial optimization of J , (2.9), over v for fixed u . Conversely, an X -iteration is obtained by optimizing over u for fixed v . Hence the sequence $(u^{(\ell)}, v^{(\ell)})_\ell$ is a maximizing sequence of (2.9). Complementarily, the sequence of measures $\pi^{(\ell)} := (u^{(\ell)} \otimes v^{(\ell)}) \cdot K$ converges (under suitable assumptions) to the solution of (2.4). In general one has $P_Y \pi^{(\ell)} = v$ but $P_X \pi^{(\ell)} \neq \mu$. Conversely, for $\pi^{(\ell+1/2)} := (u^{(\ell+1)} \otimes v^{(\ell)}) \cdot K$ one has $P_X \pi^{(\ell+1/2)} = \mu$ but $P_Y \pi^{(\ell+1/2)} \neq v$. Again, for a thorough overview we refer to [19, Chapter 4]. A computationally efficient implementation is discussed in [23], which we use as a reference in this article.

3 Entropic domain decomposition algorithm

Throughout this article we will be concerned with solving

$$\min \{ \varepsilon \operatorname{KL}(\pi|K) \mid \pi \in \Pi(\mu, \nu) \} \quad (3.1)$$

for $\mu \in \mathcal{M}_1(X)$, $\nu \in \mathcal{M}_1(Y)$, $c \in L_+^\infty(X \times Y, \mu \otimes \nu)$, $\varepsilon > 0$, $k := \exp(-c/\varepsilon)$, $K := k \cdot (\mu \otimes \nu)$. By Proposition 1, (3.1) has a unique minimizer π^* that can be represented as $\pi^* = (u^* \otimes v^*) \cdot K$ for corresponding scaling factors u^*, v^* .

Following Benamou [2] (see Sect. 1.1), the strategy of the algorithm is as follows: Divide the space X into two partitions $\mathcal{J}_A$ and $\mathcal{J}_B$ (which imply two partitions of the product space $X \times Y$) that need to ‘overlap’ in a suitable sense. Then, starting from an initial feasible coupling $\pi^0 \in \Pi(\mu, \nu)$, optimize the coupling on each of the cells of $\mathcal{J}_A$ separately, while keeping the marginals on that cell fixed. This can be done independently and in parallel for each cell and the coupling attains a potentially better score while remaining feasible. Then repeat this step on partition $\mathcal{J}_B$, then again on $\mathcal{J}_A$ and so on, continuing to alternate between the two partitions.

To construct $\mathcal{J}_A$ and $\mathcal{J}_B$ we first define a *basic partition* of small cells. The two overlapping partitions are then created by suitable merging of the basic cells. This induces a graph over the basic partition cells as vertices. The algorithm converges when this graph is connected, see Sect. 4 for details.

Definition 4 (*Basic and composite partitions*) A partition of X into measurable sets $\{X_i\}_{i \in I}$, for some finite index set I , is called a basic partition of (X, μ) if the measures $\mu_i := \mu \llcorner X_i$ for $i \in I$ satisfy $\|\mu_i\| > 0$. By construction one has $\sum_{i \in I} \mu_i = \mu$. In addition we will write $K_i := K \llcorner (X_i \times Y)$ for $i \in I$. Often we will refer to a basic partition merely by the index set I .

For a basic partition $\{X_i\}_{i \in I}$ of (X, μ) a composite partition $\mathcal{J}$ is a partition of I . For $J \in \mathcal{J}$ we will use the following notation:

$$X_J := \bigcup_{i \in J} X_i, \quad \mu_J := \sum_{i \in J} \mu_i = \mu \llcorner X_J, \quad K_J := \sum_{i \in J} K_i = K \llcorner (X_J \times Y).$$

Of course, the family $\{X_J\}_{J \in \mathcal{J}}$ is a measurable partition of X and the families $\{X_i \times Y\}_{i \in I}$ and $\{X_J \times Y\}_{J \in \mathcal{J}}$ are measurable partitions of $X \times Y$.

Throughout the article we will use one basic partition I and two corresponding composite partitions $\mathcal{J}_A$ and $\mathcal{J}_B$. A formal statement of the algorithm is given in Algorithm 1.

Example 1 In the three-cell example from Sect. 1.1, the basic partition is given by $\{X_1, X_2, X_3\}$, $I = \{1, 2, 3\}$ and the two composite partitions are given by $\mathcal{J}_A = \{\{1, 2\}, \{3\}\}$ and $\mathcal{J}_B = \{\{1\}, \{2, 3\}\}$. For ℓ odd, i.e. during an $\mathcal{J}_A$ -iteration, the optimization on the ‘single cell’ $\{3\}$ can be skipped since it is made redundant by the optimization on $\{2, 3\}$ in the subsequent $\mathcal{J}_B$ -iteration: One has $P_Y \pi^{(\ell)} \llcorner (X_3 \times Y) = P_Y \pi^{(\ell-1)} \llcorner (X_3 \times Y)$ and thus, when computing the partial marginal for cell $\{2, 3\}$ in the next iteration one finds

$$\begin{aligned} v_{\{2,3\}}^{(\ell+1)} &= P_Y(\pi^{(\ell)} \llcorner (X_{\{2,3\}} \times Y)) = P_Y(\pi^{(\ell)} \llcorner (X_2 \times Y)) + P_Y(\pi^{(\ell)} \llcorner (X_3 \times Y)) \\ &= P_Y(\pi^{(\ell)} \llcorner (X_2 \times Y)) + P_Y(\pi^{(\ell-1)} \llcorner (X_3 \times Y)). \end{aligned}$$

The same holds for the cell $\{1\}$ during the $\mathcal{J}_B$ -iteration. We then obtain the method described in the introduction.

Algorithm 1 Entropic domain decomposition

Input: initial feasible coupling $\pi^0 \in \Pi(\mu, \nu)$

Output: a sequence $(\pi^{(\ell)})_\ell$ of feasible couplings in $\Pi(\mu, \nu)$

```

1:  $\pi^{(0)} \leftarrow \pi^0$ 
2:  $\ell \leftarrow 0$ 
3: loop
4:    $\ell \leftarrow \ell + 1$ 
5:   if ( $\ell$  is odd) then  $\mathcal{J}^{(\ell)} \leftarrow \mathcal{J}_A$  else  $\mathcal{J}^{(\ell)} \leftarrow \mathcal{J}_B$  // select the partition
6:   for all  $J \in \mathcal{J}^{(\ell)}$  do // iterate over each composite cell
7:      $v_J^{(\ell)} \leftarrow P_Y(\pi^{(\ell-1)} \llcorner (X_J \times Y))$  // compute Y-marginal on cell
8:      $\pi_J^{(\ell)} \leftarrow \arg \min \left\{ \varepsilon \text{ KL}(\pi | K) \mid \pi \in \Pi(\mu_J, v_J^{(\ell)}) \right\}$  // minimizer is unique
9:   end for
10:   $\pi^{(\ell)} \leftarrow \sum_{J \in \mathcal{J}^{(\ell)}} \pi_J^{(\ell)}$ 
11: end loop
```

Comment: The theoretical, idealized algorithm runs indefinitely and thus has no stopping criterion. In practice, the stopping criterion may be a fixed number of iterations or a bound on the decrement of the primal score.

Proposition 3 Algorithm 1 is well-defined and for all $\ell \in \mathbb{N}_0$ one has $\pi^{(\ell)} \in \Pi(\mu, \nu)$.

Proof $\pi^{(0)} \in \Pi(\mu, \nu)$ by assumption. Assume now $\pi^{(\ell-1)} \in \Pi(\mu, \nu)$ for some $\ell \geq 1$. Since

$$\mu_J = \mu \llcorner X_J = (P_X \pi^{(\ell-1)}) \llcorner X_J = P_X(\pi^{(\ell-1)} \llcorner (X_J \times Y))$$

we find that $\|v_J^{(\ell)}\| = \|\mu_J\| > 0$ (for $v_J^{(\ell)}$ as introduced in line 7). Also, since $0 \leq \mu_J \leq \mu$ and $0 \leq v_J^{(\ell)} \leq \nu$, the values of $\frac{dv_J^{(\ell)}}{d\nu}$ and $\frac{d\mu_J^{(\ell)}}{d\mu}$ are (up to negligible sets) contained in $[0, 1]$. Therefore, by Proposition 1, the regularized optimal transport problem in line 8 has a unique solution and thus $\pi_J^{(\ell)}$ is well-defined. We verify

$$\begin{aligned} P_X \pi^{(\ell)} &= \sum_{J \in \mathcal{J}^{(\ell)}} P_X \pi_J^{(\ell)} = \sum_{J \in \mathcal{J}^{(\ell)}} \mu_J = \mu, \\ P_Y \pi^{(\ell)} &= \sum_{J \in \mathcal{J}^{(\ell)}} P_Y \pi_J^{(\ell)} = \sum_{J \in \mathcal{J}^{(\ell)}} v_J^{(\ell)} = \sum_{J \in \mathcal{J}^{(\ell)}} P_Y(\pi^{(\ell-1)} \llcorner (X_J \times Y)) = P_Y \pi^{(\ell-1)} = \nu. \end{aligned}$$

Therefore, $\pi^{(\ell)} \in \Pi(\mu, \nu)$ and the claim follows by induction. $\square$

Definition 5 (Notation of iterates) For $\ell \geq 1$ the following sequences will be used throughout the rest of this article:

- Composite partitions $\mathcal{J}^{(\ell)} := \mathcal{J}_A$ if ℓ is odd, $\mathcal{J}_B$ else, as defined in Algorithm 1.
- Primal iterates $\pi^{(\ell)}$, as defined in Algorithm 1.
- Partial primal iterates and corresponding Y -marginals on the composite partition cells,

$$\pi_J^{(\ell)} := \pi^{(\ell)} \llcorner (X_J \times Y), \quad v_J^{(\ell)} := P_Y \pi_J^{(\ell)}$$

for $J \in \mathcal{J}^{(\ell)}$, consistent with their definition in Algorithm 1. Let $u_J^{(\ell)}, v_J^{(\ell)}$ be corresponding scaling factors such that $\pi_J^{(\ell)} = u_J^{(\ell)} \otimes v_J^{(\ell)} \cdot K$. Their existence is provided by Proposition 1.

- Partial primal iterates and corresponding Y -marginals on the basic partition cells,

$$\pi_i^{(\ell)} := \pi^{(\ell)} \llcorner (X_i \times Y), \quad v_i^{(\ell)} := P_Y \pi_i^{(\ell)},$$

for $i \in I$. It is easy to verify that for

$$u_i^{(\ell)}(x) := \begin{cases} u_J^{(\ell)}(x) & \text{if } x \in X_i, \\ 0 & \text{else,} \end{cases}, \quad v_i^{(\ell)}(y) := v_J^{(\ell)}(y) \quad (3.2)$$

where $J \in \mathcal{J}^{(\ell)}$ is uniquely determined by the condition $i \in J$, one finds $\pi_i^{(\ell)} = u_i^{(\ell)} \otimes v_i^{(\ell)} \cdot K$ and therefore, by Propositions 1 and 2, $\pi_i^{(\ell)}$ is the unique optimal entropic coupling between μ_i and $\nu_i^{(\ell)}$. This also follows from a restriction argument as in [26, Theorem 4.6].

Based on Proposition 1, partial iterates satisfy locally some corresponding bounds, which we collect in the following Lemma, the proof of which is postponed to the “Appendix”.

Lemma 1 (Boundedness of scaling factors and density bounds) *Let $\ell > 1$, $i \in I$ and $J \in \mathcal{J}^{(\ell)}$ such that $i \in J$.*

(i) *There exist $\bar{C}, \underline{C} > 0$, possibly depending on ℓ and J , such that*

$$\begin{cases} \underline{C} \leq u_J^{(\ell)}(x) \leq \bar{C} & \text{for } \mu_{J\text{-a.e.}} x \in X_J \\ \underline{C} \cdot \frac{dv_J^{(\ell)}}{dv}(y) \leq v_J^{(\ell)}(y) \leq \bar{C} & \text{for } v\text{-a.e. } y \in Y \end{cases}$$

and, in particular, $\log u_J^{(\ell)} \in L^1(X, \mu_J)$ and $\log v_J^{(\ell)} \in L^1(Y, v_J^{(\ell)})$.

(ii) *The ratios of the $u_J^{(\ell)}$ -scaling factor are bounded,*

$$\frac{u_J^{(\ell)}(x_1)}{u_J^{(\ell)}(x_2)} \leq \exp(\|c\|/\varepsilon) \quad \text{for } \mu_{J\text{-a.e.}} x_1, x_2 \in X_J. \quad (3.3)$$

(iii) *Appropriate versions of (i) and (ii) hold for $u_i^{(\ell)}, v_i^{(\ell)}$ via their definition, (3.2).*

(iv) *The product of the scaling factors $u_i^{(\ell)}, v_i^{(\ell)}$ is locally controlled by*

$$\frac{dv_i^{(\ell)}}{dv}(y) \cdot \frac{\exp(-\|c\|/\varepsilon)}{\|\mu_i\|} \leq u_i^{(\ell)}(x) \cdot v_i^{(\ell)}(y) \leq \frac{\exp(2\|c\|/\varepsilon)}{\|\mu_i\|} \quad (3.4)$$

for μ_i -a.e. $x \in X_i$ and v -a.e. $y \in Y$.

(v) *For any $j \in J$, one has*

$$u_i^{(\ell)}(x) \cdot v_i^{(\ell)}(y) \geq \exp(-3\|c\|/\varepsilon) \|\mu_j\| \cdot u_j^{(\ell-1)}(x') \cdot v_j^{(\ell-1)}(y) \quad (3.5)$$

for $(\mu_i \otimes \mu_j \otimes v)$ -a.e. $(x, x', y) \in X_i \times X_j \times Y$, and

$$\frac{dv_i^{(\ell)}}{dv}(y) \geq \exp(-2\|c\|/\varepsilon) \|\mu_i\| \cdot \frac{dv_j^{(\ell-1)}}{dv}(y) \quad \text{for } v\text{-a.e. } y \in Y. \quad (3.6)$$

To provide some intuition, we now sketch a simple proof that Algorithm 1 converges to the globally optimal solution. For simplicity, we restrict it to the three-cell problem and discrete (and finite) spaces X and Y . While it could be extended to the general setting, we instead refer to Sect. 4.

Proposition 4 *Let X and Y be finite sets. Let $\{X_1, X_2, X_3\}$, $I = \{1, 2, 3\}$ and $\mathcal{J}_A = \{\{1, 2\}, \{3\}\}$ and $\mathcal{J}_B = \{\{1\}, \{2, 3\}\}$ be a basic and two composite partitions of X . Then Algorithm 1 converges to the globally optimal solution of problem (3.1).*

Proof (Sketch) When X and Y are discrete, the spaces of measures and measurable functions all become finite-dimensional real vector spaces. For simplicity, we may assume that μ and v assign strictly positive mass to every point in X and Y . Then K

assigns strictly positive mass to every point in $X \times Y$ and therefore $\pi \ll K$ for any $\pi \in \mathcal{M}_+(X \times Y)$. The function $\text{KL}(\cdot|K)$ is therefore finite and continuous on $\mathcal{M}_+(X \times Y)$. Denote by S the ‘solving’ map $\mathcal{M}_+(X \times Y) \ni \pi \mapsto \arg \min\{\varepsilon \text{KL}(\hat{\pi}|K)|\hat{\pi} \in \Pi(\text{P}_X\pi, \text{P}_Y\pi)\}$ which is well-defined.

We show that S is continuous: Let $(\pi_n)_n$ be a sequence of measures in $\mathcal{M}_+(X \times Y)$ converging to π_∞ (this implies that the entries of π_n are uniformly bounded), then $S(\pi_n) = u_n \otimes v_n \cdot K$ for suitable u_n, v_n by Proposition 1. By re-scaling, with (2.6, 2.7) we may assume that u_n and v_n are uniformly bounded and by compactness in finite dimensions, we may extract some cluster points u_∞ and v_∞ . Since $u_n \otimes v_n \cdot K \in \Pi(\text{P}_X\pi_n, \text{P}_Y\pi_n)$, one finds that $u_\infty \otimes v_\infty \cdot K \in \Pi(\text{P}_X\pi_\infty, \text{P}_Y\pi_\infty)$, which must be equal to $S(\pi_\infty)$ by Proposition 2 (iii). Therefore, $S(\pi_n)$ converges to $S(\pi_\infty)$ and S is continuous.

Let now F_A be the map that takes an iterate $\pi^{(\ell-1)}$ to $\pi^{(\ell)}$ when ℓ is odd, and F_B the map for ℓ even. In the discrete setting, the restriction of a measure to a subset is a continuous map, and since F_A and F_B are built from measure restriction and the solving map S , they are continuous.

Consider now the sequence of odd-numbered iterates $(\pi^{(2\ell+1)})_\ell$. In finite dimensions we may extract a cluster point π^* . By continuity of F_A and F_B , $F_A(F_B(\pi^*))$ must also be a cluster point of $(\pi^{(2\ell+1)})_\ell$. Since the sequence $(\text{KL}(\pi^{(\ell)}|K))_\ell$ is non-increasing (and $\text{KL}(\cdot|K)$ is continuous here), we must have $\text{KL}(\pi^*|K) = \text{KL}(F_A(F_B(\pi^*))|K)$. But $F_A(F_B(\pi^*)) \neq \pi^*$ would imply that the KL divergence has been strictly decreased, so we must have that $F_A(F_B(\pi^*)) = F_B(\pi^*) = \pi^*$, i.e. π^* is partially optimal on all the cells of the composite partitions $\mathcal{J}_A$ and $\mathcal{J}_B$. So there are scaling factors $(u_{\{1,2\}}, v_{\{1,2\}}), (u_{\{3\}}, v_{\{3\}})$ and $(u_{\{1\}}, v_{\{1\}}), (u_{\{2,3\}}, v_{\{2,3\}})$ such that we can write

$$\pi^* = u_{\{1,2\}} \otimes v_{\{1,2\}} \cdot K_{\{1,2\}} + u_{\{3\}} \otimes v_{\{3\}} \cdot K_{\{3\}} = u_{\{1\}} \otimes v_{\{1\}} \cdot K_{\{1\}} + u_{\{2,3\}} \otimes v_{\{2,3\}} \cdot K_{\{2,3\}}.$$

We find that on $X_2 \times Y$ the functions $u_{\{1,2\}} \otimes v_{\{1,2\}}$ and $u_{\{2,3\}} \otimes v_{\{2,3\}}$ must coincide K -a.e., and by re-scaling we may thus assume that $u_{\{1,2\}} = u_{\{2,3\}}$ on X_2 μ -a.e., and $v_{\{1,2\}} = v_{\{2,3\}}$ on Y ν -a.e. and therefore we can ‘glue’ together $u_{\{1,2\}}$ and $u_{\{2,3\}}$ to get some u and set $v = v_{\{1,2\}} = v_{\{2,3\}}$ such that $\pi^* = u \otimes v \cdot K$, which must therefore be globally optimal by Proposition 2 (iii). Since this holds for any cluster point π^* , $(\pi^{(\ell)})_\ell$ converges to this point. $\square$

4 Linear convergence

In this Section we prove linear convergence of Algorithm 1 to the globally optimal solution of (3.1) with respect to the Kullback–Leibler divergence. We start with the proof for the three-cell setup in Sect. 4.1 and finally the proof for general decompositions in Sect. 4.2. Numerical worst-case examples that demonstrate the qualitative accuracy of our convergence rate bounds are presented in Sect. 4.3.

Definition 6 (*Primal suboptimality*) For $\pi \in \Pi(\mu, \nu)$ the (scaled) suboptimality of π for problem (3.1) is denoted by

$$\Delta(\pi) := \text{KL}(\pi|K) - \text{KL}(\pi^*|K) \quad (4.1)$$

where we recall that π^* denotes the unique minimizer.

Remark 2 (*Proof strategy*) The basic strategy for the proofs is inspired by [17] where linear convergence for block coordinate descent methods is shown. Indeed, the domain decomposition algorithm can be interpreted as a block coordinate descent for the entropic optimal transport problem, where in each iteration we add some additional artificial constraints that fix the Y -marginals of the current candidate coupling on each of the composite cells. The fundamental strategy of [17] is to find a bound $\Delta(\pi^{(\ell)}) \leq C (\Delta(\pi^{(\ell-k)}) - \Delta(\pi^{(\ell)}))$ for some $C \in \mathbb{R}_+$ where $k = 1$ in the three-cell setting and $k > 1$ in the general setting. This then quickly implies linear convergence of $\Delta(\pi^{(\ell)})$ to 0. In this article, we exploit the identity $\Delta(\pi^{(\ell-k)}) - \Delta(\pi^{(\ell)}) = \text{KL}(\pi^{(\ell-k)}|\pi^{(\ell)})$ (Lemma 3) and we use a primal-dual gap estimate (Proposition 2 (ii)) to bound $\Delta(\pi^{(\ell)}) \leq \text{KL}(\pi^{(\ell)}|\hat{\pi})$ for a suitable $\hat{\pi}$ of the form $\hat{u} \otimes \hat{v} \cdot K$ (Lemma 4) and then use carefully chosen $\hat{u}$ and $\hat{v}$ to control the latter by the former.

A major challenge is that Algorithm 1 does not produce any ‘full’ dual iterates. We may only use the dual variables for the composite cell problems in line 8, but they are not necessarily consistent across multiple cells. Inspired by the proof of Proposition 4 we propose a way to approximately ‘glue’ them together to obtain a global candidate.

4.1 Three cells

We first state the main result of this section.

Theorem 1 (Linear convergence for three-cell decomposition) *Let (X_1, X_2, X_3) , $I = \{1, 2, 3\}$, be a basic partition of (X, μ) into three cells and set $\mathcal{J}_A = \{\{1, 2\}, \{3\}\}$, $\mathcal{J}_B = \{\{1\}, \{2, 3\}\}$. Then, for iterates of Algorithm 1 one has for $\ell > 1$,*

$$\Delta(\pi^{(\ell)}) \leq \left(1 + \exp\left(-\frac{2\|c\|}{\varepsilon}\right) \frac{\|\mu_2\|}{\|\mu_{[1,3]}\|}\right)^{-1} \cdot \Delta(\pi^{(\ell-1)}). \quad (4.2)$$

In particular, the domain decomposition algorithm converges to the optimal solution.

The proof for Theorem 1 requires some auxiliary Lemmas. Most of the Lemmas are already formulated for more general partitions to allow reusing them in the next Section. The Lemmas strongly rely on algebraic properties of the KL divergence that are used in a similar way throughout the literature, see for example [1, Lemma 2].

Lemma 2 *Let $\pi = (u \otimes v) \cdot K$ for $u \in L_+^\infty(X, \mu)$, $v \in L_+^\infty(Y, \nu)$. Then*

$$\text{KL}(\pi|K) = \langle P_X \pi, \log u \rangle + \langle P_Y \pi, \log v \rangle - \|\pi\| + \|K\|. \quad (4.3)$$

Similarly, given $\pi_i = (u_i \otimes v_i) \cdot K_i$ for $u_i \in L_+^\infty(X, \mu)$, $v_i \in L_+^\infty(Y, \nu)$, $i \in I$, and $\pi = \sum_{i \in I} \pi_i$, then

$$\text{KL}(\pi|K) = \sum_{i \in I} [\langle P_X \pi_i, \log u_i \rangle + \langle P_Y \pi_i, \log v_i \rangle] - \|\pi\| + \|K\|. \quad (4.4)$$

We briefly recall from Sect. 2.1 at this point that $\langle \rho, f \rangle$ denotes integration of the measurable function f against the measure ρ .

Proof We show the second statement. The first then follows by setting $(X_i)_{i \in I}$, I to the trivial partition $(X_1 := X)$, $I := \{1\}$. Arguing as in Proposition 1 (iv) we find $\log u_i \in L^1(X, P_X \pi_i)$, $\log v_i \in L^1(Y, P_Y \pi_i)$ and therefore, all integrals in the following are finite. The proof now follows quickly from direct computation:

$$\begin{aligned} \text{KL}(\pi|K) &= \sum_{i \in I} \text{KL}(\pi_i|K_i) = \sum_{i \in I} \int_{X \times Y} \varphi\left(\frac{d\pi_i}{dK_i}\right) dK_i \\ &= \sum_{i \in I} \left[\int_{X \times Y} \log(u_i(x) v_i(y)) d\pi_i(x, y) - \pi_i(X \times Y) + K_i(X \times Y) \right] \\ &= \sum_{i \in I} \left[\int_{X \times Y} [\log(u_i(x)) + \log(v_i(y))] d\pi_i(x, y) \right] - \|\pi\| + \|K\| \\ &= \sum_{i \in I} [\langle P_X \pi_i, \log u_i \rangle + \langle P_Y \pi_i, \log v_i \rangle] - \|\pi\| + \|K\| \end{aligned}$$

□

Lemma 3 (Expressions for decrement and sub-optimality) *Let $\ell > 1$. Then,*

$$\Delta(\pi^{(\ell-1)}) - \Delta(\pi^{(\ell)}) = \text{KL}(\pi^{(\ell-1)}|K) - \text{KL}(\pi^{(\ell)}|K) = \text{KL}(\pi^{(\ell-1)}|\pi^{(\ell)})$$

and

$$\Delta(\pi^{(\ell)}) = \text{KL}(\pi^{(\ell)}|K) - \text{KL}(\pi^*|K) = \text{KL}(\pi^{(\ell)}|\pi^*).$$

Proof Fix $\ell > 1$. With the notation of Definition 5 we can decompose the current and previous iterate as

$$\pi^{(\ell-1)} = \sum_{i \in I} u_i^{(\ell-1)} \otimes v_i^{(\ell-1)} \cdot K_i, \quad \pi^{(\ell)} = \sum_{i \in I} u_i^{(\ell)} \otimes v_i^{(\ell)} \cdot K_i,$$

with all partial scalings being essentially bounded with respect to μ or ν . So, by using Lemma 2, we compute

$$\text{KL}(\pi^{(\ell-1)}|K) - \text{KL}(\pi^{(\ell)}|K) = \sum_{J \in \mathcal{J}^{(\ell)}} \sum_{i \in J} [\text{KL}(\pi_i^{(\ell-1)}|K_i) - \text{KL}(\pi_i^{(\ell)}|K_i)]$$

$$\begin{aligned}
&= \sum_{J \in \mathcal{J}^{(\ell)}} \sum_{i \in J} \left[\left\langle P_X \pi_i^{(\ell-1)}, \log u_i^{(\ell-1)} \right\rangle + \left\langle P_Y \pi_i^{(\ell-1)}, \log v_i^{(\ell-1)} \right\rangle \right. \\
&\quad \left. - \left\langle P_X \pi_i^{(\ell)}, \log u_i^{(\ell)} \right\rangle - \left\langle P_Y \pi_i^{(\ell)}, \log v_i^{(\ell)} \right\rangle - \|\pi_i^{(\ell-1)}\| + \|\pi_i^{(\ell)}\| \right] \\
&= \sum_{J \in \mathcal{J}^{(\ell)}} \sum_{i \in J} \left[\left\langle \mu_i, \log u_i^{(\ell-1)} \right\rangle + \left\langle v_i^{(\ell-1)}, \log v_i^{(\ell-1)} \right\rangle \right. \\
&\quad \left. - \left\langle \mu_i, \log u_i^{(\ell)} \right\rangle - \left\langle v_i^{(\ell)}, \log v_i^{(\ell)} \right\rangle - \|\pi_i^{(\ell-1)}\| + \|\pi_i^{(\ell)}\| \right]. \quad (4.5)
\end{aligned}$$

Note now that, for any $J \in \mathcal{J}^{(\ell)}$, the Y -marginal on each composite cell is kept fixed during the iteration (see Algorithm 1, lines 7-8), i.e.,

$$\sum_{i \in J} v_i^{(\ell)} = P_Y \pi_J^{(\ell)} = v_J^{(\ell)} = P_Y \pi_J^{(\ell-1)} = \sum_{i \in J} v_i^{(\ell-1)}$$

Hence, since $v_i^{(\ell)} = v_j^{(\ell)} = v_j^{(\ell)}$ for all $i, j \in J$ (see (3.2)), we have

$$\sum_{i \in J} \left\langle v_i^{(\ell)}, \log v_i^{(\ell)} \right\rangle = \sum_{i \in J} \left\langle v_i^{(\ell-1)}, \log v_i^{(\ell)} \right\rangle.$$

Thus, we can continue:

$$\begin{aligned}
(4.5) &= \sum_{J \in \mathcal{J}^{(\ell)}} \sum_{i \in J} \left[\left\langle \mu_i, \log \left(\frac{u_i^{(\ell-1)}}{u_i^{(\ell)}} \right) \right\rangle + \left\langle v_i^{(\ell-1)}, \log \left(\frac{v_i^{(\ell-1)}}{v_i^{(\ell)}} \right) \right\rangle - \|\pi_i^{(\ell-1)}\| + \|\pi_i^{(\ell)}\| \right] \\
&= \sum_{i \in I} \text{KL}(\pi_i^{(\ell-1)} | \pi_i^{(\ell)}) = \text{KL}(\pi^{(\ell-1)} | \pi^{(\ell)}).
\end{aligned}$$

The proof of the second statement follows the same steps. It hinges on the fact that $P_X \pi_i^{(\ell-1)} = \mu_i = P_X \pi_i^*$, all $v_i^* = v^*$ are identical, and $\sum_{i \in I} P_Y \pi_i^{(\ell-1)} = \nu = \sum_{i \in I} P_Y \pi_i^*$ so that one has

$$\sum_{i \in I} \langle P_Y \pi_i^*, \log v_i^* \rangle = \sum_{i \in I} \langle P_Y \pi_i^*, \log v^* \rangle = \sum_{i \in I} \langle P_Y \pi_i^{(\ell-1)}, \log v^* \rangle = \sum_{i \in I} \langle P_Y \pi_i^{(\ell-1)}, \log v_i^* \rangle.$$

□

Lemma 4 (Primal-dual gap) *Let*

$$\tilde{\pi} = \sum_{i \in I} \tilde{\pi}_i, \quad \tilde{\pi}_i = (\tilde{u}_i \otimes \tilde{v}_i) \otimes K_i, \quad \tilde{u}_i \in L_+^\infty(X, \mu), \quad \tilde{v}_i \in L_+^\infty(Y, \nu) \quad \text{for } i \in I$$

and $\hat{\pi} = (\hat{u} \otimes \hat{v}) \cdot K$ with $\hat{u} \in L_+^\infty(X, \mu)$, $\hat{v} \in L_+^\infty(Y, \nu)$. If $\tilde{\pi} \in \Pi(\mu, \nu)$, then

$$\Delta(\tilde{\pi}) \leq \text{KL}(\tilde{\pi} | \hat{\pi}).$$

Proof If $\text{KL}(\tilde{\pi}|\hat{\pi}) = +\infty$ there is nothing to prove. Therefore, assume $\text{KL}(\tilde{\pi}|\hat{\pi}) < +\infty$. Abbreviating $\hat{\pi}_i := \hat{\pi} \llcorner (X_i \times Y)$ for $i \in I$, we find

$$\begin{aligned} \text{KL}(\tilde{\pi}|\hat{\pi}) &= \sum_{i \in I} \text{KL}(\tilde{\pi}_i|\hat{\pi}_i) = \sum_{i \in I} \int_{X \times Y} \varphi\left(\frac{\tilde{u}_i \otimes \tilde{v}_i}{\hat{u} \otimes \hat{v}}\right) d\hat{\pi}_i \\ &= \sum_{i \in I} \left[\int_{X \times Y} \log(\tilde{u}_i \otimes \tilde{v}_i) d\tilde{\pi}_i - \int_{X \times Y} \log(\hat{u} \otimes \hat{v}) d\tilde{\pi}_i - \|\tilde{\pi}_i\| + \|\hat{\pi}_i\| \right] \\ &= \text{KL}(\tilde{\pi}|K) - \int_X \log \hat{u} d\mu - \int_Y \log \hat{v} d\nu + \|\hat{\pi}\| - \|K\| \\ &= \text{KL}(\tilde{\pi}|K) - J(\hat{u}, \hat{v}) \geq \Delta(\tilde{\pi}). \end{aligned}$$

Here the integral in the first line is finite since $\text{KL}(\tilde{\pi}_i|\hat{\pi}_i) \leq \text{KL}(\tilde{\pi}|\hat{\pi}) < \infty$. The first integral in the second line is finite since $\log(\tilde{u}_i \otimes \tilde{v}_i) \in L^1(X \times Y, \tilde{\pi}_i)$, arguing as in Proposition 1 (iv), and therefore so is the second. Since all scaling factors are essentially bounded, we can split the products within the logarithms into separate integrals in the third line where we also used $\tilde{\pi} \in \Pi(\mu, \nu)$. In the fourth line we use the definition of J , (2.9). The last inequality is due to Proposition 2, since for the primal minimizer π^* we find $\text{KL}(\pi^*|K) \geq J(\hat{u}, \hat{v})$ and the primal-dual gap is thus a bound for the suboptimality (4.1). $\square$

Proof (Theorem 1) Assume first that $\ell > 1$ is odd, and so the current composite partition is $\mathcal{J}_A = \{\{1, 2\}, \{3\}\}$. On one hand, using Lemma 3, one finds

$$\Delta(\pi^{(\ell-1)}) - \Delta(\pi^{(\ell)}) = \text{KL}(\pi^{(\ell-1)}|\pi^{(\ell)}) \geq \text{KL}(\pi_2^{(\ell-1)}|\pi_2^{(\ell)}). \quad (4.6)$$

On the other hand, with Lemma 4, one can bound $\Delta(\pi^{(\ell)}) \leq \text{KL}(\pi^{(\ell)}|\hat{\pi})$ for a suitable $\hat{\pi}$ of the form $\hat{\pi} = (\hat{u} \otimes \hat{v}) \cdot K$. Using a piecewise definition on the partition cells for $\hat{u}$, we set here

$$\hat{u} := u_{\{1,2\}}^{(\ell)} + q \cdot u_{\{3\}}^{(\ell-1)} \quad \text{and} \quad \hat{v} := v_{\{1,2\}}^{(\ell)} \quad (4.7)$$

where the factor $q \in \mathbb{R}_{++}$ is to be determined later (observe that $\hat{u}$ and $\hat{v}$ are upper bounded, thanks to Lemma 1). With this choice of $\hat{\pi}$ we find that $\pi_{\{1,2\}}^{(\ell)} = \hat{\pi}_{\{1,2\}}$ and thus

$$\Delta(\pi^{(\ell)}) \leq \text{KL}(\pi^{(\ell)}|\hat{\pi}) = \text{KL}(\pi_3^{(\ell)}|\hat{\pi}_3). \quad (4.8)$$

To complete the proof we must now find a constant $C \in \mathbb{R}_+$ such that

$$\text{KL}(\pi_3^{(\ell)}|\hat{\pi}_3) \leq C \cdot \text{KL}(\pi_2^{(\ell-1)}|\pi_2^{(\ell)}).$$

To this end we write

$$\text{KL}(\pi_2^{(\ell-1)}|\pi_2^{(\ell)}) = \int_{X \times Y} \varphi\left(\frac{d\pi_2^{(\ell-1)}}{d\pi_2^{(\ell)}}\right) d\pi_2^{(\ell)}$$

$$= \int_{X \times Y} \varphi \left(\frac{u_2^{(\ell-1)}(x) v_2^{(\ell-1)}(y)}{u_2^{(\ell)}(x) v_2^{(\ell)}(y)} \right) \frac{u_2^{(\ell)}(x)}{u_2^{(\ell-1)}(x)} u_2^{(\ell-1)}(x) v_2^{(\ell)}(y) k(x, y) d\mu_2(x) dv(y). \quad (4.9)$$

Now apply (3.3) at iterate $(\ell - 1)$ and for $J = \{2, 3\}$ to obtain that for μ -a.e. $x \in X_2, x_3 \in X_3$ and for ν -a.e. $y \in Y$ it holds

$$u_2^{(\ell-1)}(x) k(x, y) \geq u_3^{(\ell-1)}(x_3) k(x_3, y) \cdot \exp(-2\|c\|/\varepsilon). \quad (4.10)$$

Further, note that for $a \geq 0$ the map $\Phi : s \mapsto (\varphi(\frac{a}{s}) \cdot s)$ is convex on $\mathbb{R}_{++}$ and thus by Jensen's inequality one has for $\nu_2^{(\ell)}$ -a.e. $y \in Y$ (such that $v_2^{(\ell)}(y) > 0$ almost surely) that

$$\int_X \varphi \left(\frac{u_2^{(\ell-1)}(x) v_2^{(\ell-1)}(y)}{u_2^{(\ell)}(x) v_2^{(\ell)}(y)} \right) \frac{u_2^{(\ell)}(x)}{u_2^{(\ell-1)}(x)} d\mu_2(x) \geq \|\mu_2\| \cdot \varphi \left(\frac{v_2^{(\ell-1)}(y)}{q v_2^{(\ell)}(y)} \right) q \quad (4.11)$$

where we now fix the value $q := \int_X \frac{u_2^{(\ell)}}{u_2^{(\ell-1)}} d\mu_2$ (which is finite because of the lower and upper boundedness of $u_2^{(\ell)}$ and $u_2^{(\ell-1)}$, cf. Lemma 1 ((i),(iii))). We recall the notation for the normalized integral from (2.1). Plugging first (4.10) and then (4.11) into (4.9) one obtains for μ -almost all $x_3 \in X_3$

$$\begin{aligned} \text{KL}(\pi_2^{(\ell-1)} | \pi_2^{(\ell)}) &\geq e^{-\frac{2\|c\|}{\varepsilon}} \cdot \|\mu_2\| \cdot \int_Y \varphi \left(\frac{v_2^{(\ell-1)}(y)}{q v_2^{(\ell)}(y)} \right) \\ &\quad q u_3^{(\ell-1)}(x_3) v_2^{(\ell)}(y) k(x_3, y) dv(y). \end{aligned} \quad (4.12)$$

Now we work backwards from (4.8). With the specified choices for $\hat{u}$ and $\hat{v}$ made in (4.7), one finds for $\hat{\pi}$ -a.e. $(x, y) \in X_3 \times Y$

$$\frac{d\pi_3^{(\ell)}}{d\hat{\pi}}(x, y) = \frac{u_3^{(\ell)}(x) \cdot v_3^{(\ell)}(y)}{q \cdot u_3^{(\ell-1)}(x) \cdot v_{\{1,2\}}^{(\ell)}(y)} = \frac{u_3^{(\ell-1)}(x) \cdot v_3^{(\ell-1)}(y)}{q \cdot u_3^{(\ell-1)}(x) \cdot v_2^{(\ell)}(y)} = \frac{v_2^{(\ell-1)}(y)}{q \cdot v_2^{(\ell)}(y)}$$

where we have used that $u_3^{(\ell)} \otimes v_3^{(\ell)} = u_3^{(\ell-1)} \otimes v_3^{(\ell-1)}$ (since $\pi_3^{(\ell)} = \pi_3^{(\ell-1)}$ due to $\{3\} \in \mathcal{J}_A$) and $v_3^{(\ell-1)} = v_2^{(\ell-1)}$ (since $\{2, 3\} \in \mathcal{J}_B$). Consequently, with (4.12)

$$\begin{aligned} \text{KL}(\pi_3^{(\ell)} | \hat{\pi}_3) &= \int_{X \times Y} \varphi \left(\frac{v_2^{(\ell-1)}(y)}{q \cdot v_2^{(\ell)}(y)} \right) q u_3^{(\ell-1)}(x) v_2^{(\ell)}(y) k(x, y) d\mu_3(x) dv(y) \\ &\geq e^{-\frac{2\|c\|}{\varepsilon}} \cdot \frac{1}{\|\mu_2\|} \int_X d\mu_3(x) \text{KL}(\pi_2^{(\ell-1)} | \pi_2^{(\ell)}) \end{aligned}$$

and with (4.6) and (4.8)

$$\Delta(\pi^{(\ell)}) \leq \exp\left(\frac{2\|c\|}{\varepsilon}\right) \cdot \frac{\|\mu_3\|}{\|\mu_2\|} \left(\Delta(\pi^{(\ell-1)}) - \Delta(\pi^{(\ell)})\right).$$

For ℓ even one proceed analogously to obtain

$$\Delta(\pi^{(\ell)}) \leq \exp\left(\frac{2\|c\|}{\varepsilon}\right) \cdot \frac{\|\mu_1\|}{\|\mu_2\|} \left(\Delta(\pi^{(\ell-1)}) - \Delta(\pi^{(\ell)})\right).$$

Hence, for any $\ell > 0$, using $\|\mu_{\{1,3\}}\| \geq \max\{\|\mu_1\|, \|\mu_3\|\}$, one obtains

$$\Delta(\pi^{(\ell)}) \leq \exp\left(\frac{2\|c\|}{\varepsilon}\right) \cdot \frac{\|\mu_{\{1,3\}}\|}{\|\mu_2\|} \cdot \left(\Delta(\pi^{(\ell-1)}) - \Delta(\pi^{(\ell)})\right)$$

which implies (4.2). Eventually, the sequence $(\pi^{(\ell)})_\ell$ is bounded and each element lies in $\Pi(\mu, \nu)$. As $\Pi(\mu, \nu)$ is weakly-* closed, $(\pi^{(\ell)})_\ell$ must have a cluster point $\pi^{(\infty)} \in \Pi(\mu, \nu)$. By (4.2) and since Δ is weakly-* lower semi-continuous one has

$$\Delta(\pi^{(\infty)}) \leq \liminf_{\ell \rightarrow \infty} \Delta(\pi^{(\ell)}) = 0.$$

Since $[\Delta(\pi) = \text{KL}(\pi|\pi^*) = 0] \Leftrightarrow [\pi = \pi^*]$ one has $\pi^{(\infty)} = \pi^*$. This also implies that every cluster point of $(\pi^{(\ell)})_\ell$ must equal π^* and that thus the sequence is converging. $\square$

4.2 General decompositions

Now let I be a basic partition of (X, μ) into N basic cells and let $\mathcal{J}_A$ and $\mathcal{J}_B$ be two composite partitions of I . In the special case of three cells discussed in the previous Section, convergence of the algorithm was driven by the overlap of the two partitions on the middle cell. In the general case this overlap structure will be captured by the *partition graph* and related auxiliary objects that we introduce now.

Definition 7 (*Partition graph, cell distance and shortest paths*)

- (i) The partition graph is given by the vertex set I (i.e. each basic cell is represented by one vertex) and the edge set

$$E := \{(i, j) \in I \mid \exists J \in \mathcal{J}_A \cup \mathcal{J}_B \text{ such that } \{i, j\} \subset J\},$$

i.e. there is an edge between two basic cells if they are part of the same composite cell in either of the two composite partitions $\mathcal{J}_A$ or $\mathcal{J}_B$.

- (ii) We assume that no two basic cells are simultaneously contained in a composite cell of $\mathcal{J}_A$ and $\mathcal{J}_B$, i.e. there exist no $i, j \in I$, $J_A \in \mathcal{J}_A$, $J_B \in \mathcal{J}_B$ such that $i, j \in J_A$ and $i, j \in J_B$ (otherwise i and j should be merged into one basic cell). This means that every edge in the partition graph is associated to precisely one of the two composite partitions.

Fig. 2 Illustration of partition graph and cell distance for a typical 2D setup. Vertices, the elements of I , are represented by circles, with cell distance $D(i)$ given as labels. Composite cells of $\mathcal{J}_A$ are indicated by rectangles with solid lines, J_0 is highlighted with grey filling. Composite cells of $\mathcal{J}_B$ are indicated by rectangles with dashed lines. Shortest paths from vertices to J_0 are indicated by black lines (these are often not unique). For the vertex highlighted in grey the indices $(n_{i,k})_{k=0}^{D(i)}$ are given below the corresponding nodes

- (iii) We assume that the partition graph is connected.
- (iv) We denote by $\text{dist} : I \times I \rightarrow \mathbb{N}_0$ the discrete metric on this graph induced by shortest paths, where each edge has length 1.
- (v) Let $J_0 \in \mathcal{J}_A$ be a selected composite cell. We define the *cell distance*

$$D : I \rightarrow \mathbb{N}_0, \quad i \mapsto \text{dist}(i, J_0) = \min\{\text{dist}(i, j) \mid j \in J_0\}$$

and write $M := \max\{D(i) \mid i \in I\}$.

- (vi) For each $i \in I$ we select one shortest path in the graph from J_0 to i , which we represent by a tuple $(n_{i,k})_{k=0}^{D(i)}$ of elements in I with $n_{i,0} \in J_0$, $n_{i,D(i)} = i$, and $D(n_{i,k}) = k$ for $k \in \{0, \dots, D(i)\}$. We refer to Fig. 2 for an illustration of such a construction. One can easily verify that, within any shortest path, the basic cells $n_{i,k}$ and $n_{i,k-1}$ belong to a common composite cell of partition $\mathcal{J}_A$ when k is even, and $\mathcal{J}_B$ when k is odd, or equivalently, of $\mathcal{J}^{(\ell-k)}$ for some odd $\ell > M$, see Fig. 3.

Let us fix for the rest of this Section an odd iterate ℓ so that $\mathcal{J}^{(\ell)} = \mathcal{J}_A$ (the same will then apply for ℓ even by swapping the roles of $\mathcal{J}_A$ and $\mathcal{J}_B$).

Theorem 2 (Linear convergence) *Let $\ell > M$, with M defined in Definition 7 (v). Then, one finds*

$$\Delta(\pi^{(\ell)}) \leq \frac{2MN \exp((6M+7) \|c\|/\varepsilon)}{\mu_{\min}^{2M+1} + 2MN \exp((6M+7) \|c\|/\varepsilon)} \cdot \Delta(\pi^{(\ell-M)}). \quad (4.13)$$

where $\mu_{\min} := \min\{\|\mu_i\| \mid i \in I\}$. In particular, the domain decomposition algorithm converges to the optimal solution.

Remark 3 (*Proof strategy*) For the three-cell case the proof relied on the fact that $\pi_3^{(\ell)}$ remained unchanged in an A -iteration, since $\{3\} \in \mathcal{J}_A$ was a composite cell containing only basic cell 3. This is no longer true for general decompositions. To overcome this, we need to replace parts of the primal iterate $\pi^{(\ell)}$ by partial transport plans from

Fig. 3 Illustration of proof strategies for Lemma 5 and Theorem 2. The horizontal axis represents basic cells along the shortest path $(n_{i,k})_{k=0}^{D(i)}$ from $J_0 \in \mathcal{J}_A$ (left side) to some $i \in I$ (right side), vertical axis represents iterations for some odd $\ell > M$ (i.e. $\mathcal{J}^{(\ell)} = \mathcal{J}_A$). Solid vertical lines represent boundaries of composite cells at given iteration. Grey shading of boxes illustrates the construction of $\tilde{\pi}$ in (4.14). Vertical arrows represent factors $q_{i,k}$, cf. (4.17). Horizontal arrows represent edges in the partition graph and the equalities $v_{n_{i,k}}^{(\ell-k)} = v_{n_{i,k-1}}^{(\ell-k)}$ used in (4.20)

previous iterations (Lemma 5). When bounding the sub-optimality by the decrement (cf. Remark 2) we also need an approximate triangle inequality for the KL-divergence (Lemma 7).

Lemma 5 (Construction of primal upper bound candidate) *For $\ell > M$, ℓ odd, the measure*

$$\tilde{\pi} = \sum_{i \in I} \pi_i^{(\ell-D(i))} \quad (4.14)$$

satisfies $\tilde{\pi} \in \Pi(\mu, \nu)$ and $\text{KL}(\tilde{\pi}|\pi^) \geq \text{KL}(\pi^{(\ell)}|\pi^*)$.*

Proof In the proof we assume that ℓ is odd (i.e. an A-iteration was just completed). The proof for ℓ even works completely analogous by swapping the roles of J_A and J_B .

We start with the claim that $\tilde{\pi} \in \Pi(\mu, \nu)$. Since $\tilde{\pi}$ is a sum of non-negative measures it is non-negative. By construction of the algorithm one always has $P_X \pi_i^{(m)} = \mu_i$ for all iterations $m \in \mathbb{N}_0$ and thus

$$P_X \tilde{\pi} = \sum_{i \in I} P_X \pi_i^{(\ell-D(i))} = \sum_{i \in I} \mu_i = \mu.$$

For $k \in \{0, \dots, M\}$ we introduce the auxiliary measures

$$\tilde{\pi}^k = \sum_{i \in I} \pi_i^{(\ell-\min\{D(i), k\})}.$$

Note that $\tilde{\pi}^0 = \pi^{(\ell)}$ and $\tilde{\pi}^M = \tilde{\pi}$. Further, one obtains for $k \in \{1, \dots, M\}$

$$\begin{aligned}\tilde{\pi}_i^k - \tilde{\pi}_i^{k-1} &= \begin{cases} \pi_i^{(\ell-k)} - \pi_i^{(\ell-k+1)} & \text{if } D(i) \geq k, \\ 0 & \text{else.} \end{cases} \\ \tilde{\pi}^k - \tilde{\pi}^{k-1} &= \sum_{\substack{i \in I: \\ D(i) \geq k}} \left(\pi_i^{(\ell-k)} - \pi_i^{(\ell-k+1)} \right)\end{aligned}$$

Assume now that k is odd. Recalling that ℓ is assumed to be odd, then the set $J^k := \{i \in I : D(i) \geq k\}$ is a union of A -cells and the iteration step from $\pi^{(\ell-k)}$ to $\pi^{(\ell-k+1)}$ is an iteration on A -cells (since $\ell - k + 1$ is odd, recall also Definition 7), i.e., for any $J \in \mathcal{J}_A$ one has

$$P_Y \pi_J^{(\ell-k)} = P_Y \pi_J^{(\ell-k+1)}.$$

Since J^k is a union of A -cells we find therefore

$$P_Y \sum_{i \in J^k} \pi_i^{(\ell-k)} = P_Y \sum_{i \in J^k} \pi_i^{(\ell-k+1)}.$$

Using this, we find that $P_Y \tilde{\pi}^k = P_Y \tilde{\pi}^{k-1}$ for $k \in \{1, \dots, M\}$ and k odd, and we can argue in complete analogy for even k via B -cells. Consequently,

$$P_Y \tilde{\pi} = P_Y \tilde{\pi}^M = P_Y \tilde{\pi}^{M-1} = \dots = P_Y \tilde{\pi}^0 = P_Y \pi^{(\ell)} = P_Y \pi = \nu$$

and thus $\tilde{\pi} \in \Pi(\mu, \nu)$.

Now we establish the KL bound. Again, assume that k is odd, so that J^k is a union of A -cells and the iteration step from $\pi^{(\ell-k)}$ to $\pi^{(\ell-k+1)}$ is an iteration on A -cells. Then for any $J \in \mathcal{J}_A$ one finds

$$\text{KL}(\pi_J^{(\ell-k)} | \pi^*) \geq \text{KL}(\pi_J^{(\ell-k+1)} | \pi^*)$$

and since J^k is a union of A -cells one gets that

$$\begin{aligned}\text{KL}(\tilde{\pi}^k | \pi^*) - \text{KL}(\tilde{\pi}^{k-1} | \pi^*) &= \sum_{i \in I} \left[\text{KL}(\tilde{\pi}_i^k | \pi_i^*) - \text{KL}(\tilde{\pi}_i^{k-1} | \pi_i^*) \right] \\ &= \sum_{i \in J^k} \left[\text{KL}(\pi_i^{(\ell-k)} | \pi_i^*) - \text{KL}(\pi_i^{(\ell-k+1)} | \pi_i^*) \right] \geq 0.\end{aligned}$$

Again, an analogous argument holds when k is even. Consequently one has

$$\text{KL}(\tilde{\pi} | \pi^*) = \text{KL}(\tilde{\pi}^M | \pi^*) \geq \text{KL}(\tilde{\pi}^{M-1} | \pi^*) \geq \dots \geq \text{KL}(\tilde{\pi}^0 | \pi^*) = \text{KL}(\pi^{(\ell)} | \pi^*).$$

□

Lemma 6 (Global density bound) *Let $\ell > M$. Then for ν -a.e. $y \in Y$ and all $i \in I$ one has*

$$\frac{dv_i^{(\ell)}}{dv}(y) \geq \frac{\mu_{\min}^{M+1}}{N} \exp(-2(M+1)\|c\|/\varepsilon). \quad (4.15)$$

Proof Let $J \in \mathcal{J}^{(\ell)}$ such that $i \in J$. We apply (3.6) for each $j \in J$ to obtain

$$\frac{dv_i^{(\ell)}}{dv}(y) \geq \exp(-2\|c\|/\varepsilon) \mu_{\min} \cdot \max_{j \in J} \frac{dv_j^{(\ell-1)}}{dv}(y).$$

for ν -a.e. $y \in Y$. Applying this bound recursively, the max runs over increasingly many cells and after $M+1$ applications of the bound we find:

$$\frac{dv_i^{(\ell)}}{dv}(y) \geq \exp(-2(M+1)\|c\|/\varepsilon) \mu_{\min}^{M+1} \cdot \max_{j \in I} \frac{dv_j^{(\ell-M-1)}}{dv}(y).$$

Observe now that $\sum_{j \in I} \frac{dv_j^{(\ell-M-1)}}{dv}(y) = 1$ and thus the maximum must be bounded from below by $1/N$. $\square$

Proof (Theorem 2) Let us set

$$\hat{\pi} := (\hat{u} \otimes \hat{v}) \cdot K \quad \text{with} \quad \hat{u}_i := \left(\prod_{k=0}^{D(i)-1} q_{i,k} \right) \cdot u_i^{(\ell-D(i))} \quad \text{and} \quad \hat{v} := v_{J_0}^{(\ell)} \quad (4.16)$$

where the factors $q_{i,k}$ are defined as

$$q_{i,k} := \int_X \frac{u_{n_{i,k}}^{(\ell-k)}}{u_{n_{i,k}}^{(\ell-k-1)}} d\mu_{n_{i,k}} \quad (4.17)$$

and are finite thanks to the lower and upper boundedness of u -scalings, cf. Lemma 1, ((i),(iii)). Thus, with Lemmas 3, 4 and $\tilde{\pi}$ from Lemma 5, one obtains

$$\Delta(\pi^{(\ell)}) = \text{KL}(\pi^{(\ell)}|\pi^*) \leq \text{KL}(\tilde{\pi}|\pi^*) \leq \text{KL}(\tilde{\pi}|\hat{\pi}) = \sum_{i \in I} \text{KL}(\tilde{\pi}_i|\hat{\pi}_i) \quad (4.18)$$

where one has

$$\text{KL}(\tilde{\pi}_i|\hat{\pi}_i) = \int_{X_i \times Y} \varphi\left(\frac{d\tilde{\pi}_i}{d\hat{\pi}_i}\right) d\hat{\pi}_i.$$

For $i \in J_0$ one finds $\hat{\pi}_i$ almost everywhere that

$$\frac{d\tilde{\pi}_i}{d\hat{\pi}_i} = \frac{\tilde{u}_i \otimes \tilde{v}_i}{\hat{u}_i \otimes \hat{v}} = \frac{u_i^{(\ell)} \otimes v_i^{(\ell)}}{u_i^{(\ell)} \otimes v_{J_0}^{(\ell)}} = 1 \quad \text{and thus} \quad \text{KL}(\tilde{\pi}_i|\hat{\pi}_i) = 0$$

where we have used that $\hat{u}_i = u_i^{(\ell)}$ and $v_i^{(\ell)} = v_{J_0}^{(\ell)}$ when $i \in J_0$ (which implies $D(i) = 0$).

To control $\text{KL}(\tilde{\pi}_i | \hat{\pi}_i)$ for $i \in I$ with $D(i) > 0$ we will use the following bound for $k \in \{0, \dots, D(i) - 1\}$:

$$\begin{aligned} & \text{KL}(\pi_{n_{i,k}}^{(\ell-k-1)} | \pi_{n_{i,k}}^{(\ell-k)}) \\ &= \int_{X \times Y} \varphi \left(\frac{u_{n_{i,k}}^{(\ell-k-1)}(x) \cdot v_{n_{i,k}}^{(\ell-k-1)}(y)}{u_{n_{i,k}}^{(\ell-k)}(x) \cdot v_{n_{i,k}}^{(\ell-k)}(y)} \right) u_{n_{i,k}}^{(\ell-k)}(x) \cdot v_{n_{i,k}}^{(\ell-k)}(y) k(x, y) d\mu_{n_{i,k}}(x) dv(y) \\ &\geq e^{-\frac{2\|c\|}{\varepsilon}} \cdot \int_{X \times Y} \varphi \left(\frac{u_{n_{i,k}}^{(\ell-k-1)}(x) \cdot v_{n_{i,k}}^{(\ell-k-1)}(y)}{u_{n_{i,k}}^{(\ell-k)}(x) \cdot v_{n_{i,k}}^{(\ell-k)}(y)} \right) \frac{u_{n_{i,k}}^{(\ell-k)}(x)}{u_{n_{i,k}}^{(\ell-k-1)}(x)} u_{n_{i,k}}^{(\ell-k-1)}(\hat{x}) \cdot v_{n_{i,k}}^{(\ell-k)}(y) d\mu_{n_{i,k}}(x) dv(y) \\ &= e^{-\frac{2\|c\|}{\varepsilon}} \cdot \int_Y \left[\int_X \varphi \left(\frac{u_{n_{i,k}}^{(\ell-k-1)}(x) \cdot v_{n_{i,k}}^{(\ell-k-1)}(y)}{u_{n_{i,k}}^{(\ell-k)}(x) \cdot v_{n_{i,k}}^{(\ell-k)}(y)} \right) \frac{u_{n_{i,k}}^{(\ell-k)}(x)}{u_{n_{i,k}}^{(\ell-k-1)}(x)} d\mu_{n_{i,k}}(x) \right] u_{n_{i,k}}^{(\ell-k-1)}(\hat{x}) \cdot v_{n_{i,k}}^{(\ell-k)}(y) dv(y) \\ &\geq e^{-\frac{2\|c\|}{\varepsilon}} \|\mu_{n_{i,k}}\| \int_Y \varphi \left(\frac{v_{n_{i,k}}^{(\ell-k-1)}(y)}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}(y)} \right) q_{i,k} u_{n_{i,k}}^{(\ell-k-1)}(\hat{x}) v_{n_{i,k}}^{(\ell-k)}(y) dv(y) \end{aligned} \quad (4.19)$$

for $\mu_{n_{i,k}}$ -a.e. $\hat{x} \in X_{n_{i,k}}$, where in the first inequality we have used (3.3) and the boundedness of c , and in the second inequality we have used Jensen's inequality as in (4.11) in the proof of Theorem 1.

Now note that one has $v_{n_{i,k}}^{(\ell-k)} = v_{n_{i,k-1}}^{(\ell-k)}$ for $k \in \{1, \dots, D(i)\}$, since $\{n_{i,k}, n_{i,k-1}\} \in J$ for some $J \in \mathcal{J}^{(\ell-k)}$, cf. Definition 7 (vi) and Fig. 3. Therefore, for i with $D(i) > 0$ and recalling $n_{i,D(i)} = i$, one obtains $\hat{\pi}_i$ -almost everywhere that

$$\begin{aligned} \frac{d\tilde{\pi}_i}{d\hat{\pi}_i} &= \frac{\tilde{u}_i \otimes \tilde{v}_i}{\hat{u}_i \otimes \hat{v}_i} = \frac{u_i^{(\ell-D(i))} \otimes v_i^{(\ell-D(i))}}{\left(\prod_{k=0}^{D(i)-1} q_{i,k} \right) \cdot u_i^{(\ell-D(i))} \otimes v_{J_0}^{(\ell)}} = \frac{v_{n_{i,D(i)}}^{(\ell-D(i))}}{\left(\prod_{k=0}^{D(i)-1} q_{i,k} \right) \cdot v_{n_{i,0}}^{(\ell)}} \\ &= \left(\prod_{k=0}^{D(i)-1} q_{i,k} \right)^{-1} \cdot \frac{v_{n_{i,D(i)-1}}^{(\ell-D(i))}}{v_{n_{i,0}}^{(\ell)}} \cdot \underbrace{\left(\prod_{k=1}^{D(i)-1} \frac{v_{n_{i,k-1}}^{(\ell-k)}}{v_{n_{i,k}}^{(\ell-k)}} \right)}_{=1} = \prod_{k=0}^{D(i)-1} \frac{v_{n_{i,k}}^{(\ell-k-1)}}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}}. \end{aligned} \quad (4.20)$$

So continuing the suboptimality bound from (4.18) we obtain

$$\Delta(\pi^{(\ell)}) \leq \sum_{\substack{i \in I: \\ D(i) > 0}} \int_{X \times Y} \varphi \left(\prod_{k=0}^{D(i)-1} \frac{v_{n_{i,k}}^{(\ell-k-1)}}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}} \right) d\hat{\pi}_i.$$

In the following we now transform the ‘ φ of product’ into a ‘sum of φ ’ via Lemma 7. For each term we then use several auxiliary results to bound $\hat{\pi}_i$ from below by the appropriate $\pi_i^{(\ell-k)}$. Using the definition of $q_{i,k}$ in (4.17) and (3.5) we find that

$$\frac{v_{n_{i,k}}^{(\ell-k-1)}}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}} \leq \frac{\exp(3\|c\|/\varepsilon)}{\mu_{\min}} =: L \geq 1$$

and so by using Lemma 7 one obtains for

$$C_1 := 2 D(i) \cdot L^{D(i)-1} \geq D(i) \cdot \max\{2, L^{D(i)-1}\}$$

that

$$\Delta(\pi^{(\ell)}) \leq C_1 \cdot \sum_{\substack{i \in I: \\ D(i) > 0}} \sum_{k=0}^{D(i)-1} \int_{X \times Y} \varphi \left(\frac{v_{n_{i,k}}^{(\ell-k-1)}}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}} \right) d\hat{\pi}_i. \quad (4.21)$$

Further, since $i = n_{i,D(i)}$ and $\{n_{i,D(i)}, n_{i,D(i)-1}\} \subset J$ for some $J \in \mathcal{J}^{(\ell-D(i))}$ (cf. Definition 7 (vi)), by using (3.3) one finds that for μ_i -a.e. $x \in X_i$

$$\begin{aligned} q_{i,D(i)-1} \cdot u_i^{(\ell-D(i))}(x) &= \int_X \frac{u_{n_{i,D(i)}}^{(\ell-D(i))}(x) \cdot u_{n_{i,D(i)-1}}^{(\ell-(D(i)-1))}(\hat{x})}{u_{n_{i,D(i)-1}}^{(\ell-D(i))}(\hat{x})} d\mu_{n_{i,D(i)-1}}(\hat{x}) \\ &\leq \exp(\|c\|/\varepsilon) \int_X u_{n_{i,D(i)-1}}^{(\ell-(D(i)-1))}(\hat{x}) d\mu_{n_{i,D(i)-1}}(\hat{x}) \end{aligned}$$

and therefore, applying iteratively this bound along the chain connecting X_i to $X_{n_{i,0}}$, we obtain for μ_i -a.e. $x \in X_i$ and $\mu_{n_{i,0}}$ -a.e. $x' \in X_{n_{i,0}}$

$$\begin{aligned} \hat{u}_i(x) &= \left[\prod_{k=0}^{D(i)-1} q_{i,k} \right] \cdot u_i^{(\ell-D(i))}(x) \leq e^{\frac{\|c\|}{\varepsilon}} \cdot \int_X \left[\prod_{k=0}^{D(i)-2} q_{i,k} \right] u_{n_{i,D(i)-1}}^{(\ell-D(i)+1)}(\hat{x}) d\mu_{n_{i,D(i)-1}}(\hat{x}) \\ &\leq \cdots \leq \exp(D(i) \|c\|/\varepsilon) \cdot \int_X u_{n_{i,0}}^{(\ell)}(\hat{x}) d\mu_{n_{i,0}}(\hat{x}) \leq \exp((D(i)+1) \|c\|/\varepsilon) \cdot u_{n_{i,0}}^{(\ell)}(x'). \end{aligned} \quad (4.22)$$

Combining successively the upper bound in (3.4), (4.15) and the lower bound in (3.4) one deduces for $(\mu_{n_{i,0}} \otimes \mu_{n_{i,k}})$ -a.e. $(x, x') \in X_{n_{i,0}} \times X_{n_{i,k}}$ and ν -a.e. $y \in Y$ that

$$\begin{aligned} u_{n_{i,0}}^{(\ell)}(x) \cdot v_{n_{i,0}}^{(\ell)}(y) &\leq \frac{\exp(2\|c\|/\varepsilon)}{\mu_{\min}} \leq \frac{N \exp((2M+4)\|c\|/\varepsilon)}{\mu_{\min}^{M+2}} \frac{dv_{n_{i,k}}^{(\ell-k)}}{dv}(y) \\ &\leq \frac{N \exp((2M+5)\|c\|/\varepsilon)}{\mu_{\min}^{M+2}} \cdot \|\mu_{n_{i,k}}\| \cdot u_{n_{i,k}}^{(\ell-k)}(x') \cdot v_{n_{i,k}}^{(\ell-k)}(y) \end{aligned} \quad (4.23)$$

and we set

$$C_2 := \frac{N \exp((2M+5)\|c\|/\varepsilon)}{\mu_{\min}^{M+2}}.$$

Eventually, applying again (3.3), we can easily see that

$$u_{n_{i,k}}^{(\ell-k)}(x) \leq e^{\frac{2\|c\|}{\varepsilon}} \cdot q_{i,k} \cdot u_{n_{i,k}}^{(\ell-k-1)}(x') \quad \text{for } \mu_{n_{i,k}} \text{ -a.e. } x, x' \in X_{n_{i,k}} \quad (4.24)$$

We now have all the ingredients to go back to (4.21). For every fixed $i \in I$ with $\hat{\pi}_i = \hat{u}_i \otimes \hat{v}_i \cdot K_i$ and $0 \leq k \leq D(i) - 1$, and for $\mu_{n_{i,0}}$ -a.e. $x' \in X_{n_{i,0}}$ and $\mu_{n_{i,k}}$ -a.e. $x'' \in X_{n_{i,k}}$, we compute

$$\begin{aligned}
 \int_{X \times Y} \varphi \left(\frac{v_{n_{i,k}}^{(\ell-k-1)}}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}} \right) d\hat{\pi}_i &= \int_{X \times Y} \varphi \left(\frac{v_{n_{i,k}}^{(\ell-k-1)}(y)}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}(y)} \right) \hat{u}_i(x) \hat{v}_i(y) \underbrace{k(x, y)}_{\leq 1} d\mu_i(x) dv(y) \\
 &\stackrel{(4.22)}{\leq} \|\mu_i\| e^{\frac{(M+1)\|c\|}{\varepsilon}} \int_Y \varphi \left(\frac{v_{n_{i,k}}^{(\ell-k-1)}(y)}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}(y)} \right) \cdot u_{n_{i,0}}^{(\ell)}(x') \cdot v_{n_{i,0}}^{(\ell)}(y) dv(y) \\
 &\stackrel{(4.23)}{\leq} C_2 \|\mu_i\| \cdot \|\mu_{n_{i,k}}\| e^{\frac{(M+1)\|c\|}{\varepsilon}} \int_Y \varphi \left(\frac{v_{n_{i,k}}^{(\ell-k-1)}(y)}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}(y)} \right) u_{n_{i,k}}^{(\ell-k)}(x'') \cdot v_{n_{i,k}}^{(\ell-k)}(y) dv(y) \\
 &\stackrel{(4.24)}{\leq} C_2 \|\mu_i\| \cdot \|\mu_{n_{i,k}}\| e^{\frac{(M+3)\|c\|}{\varepsilon}} \int_Y \varphi \left(\frac{v_{n_{i,k}}^{(\ell-k-1)}(y)}{q_{i,k} v_{n_{i,k}}^{(\ell-k)}(y)} \right) q_{i,k} \cdot u_{n_{i,k}}^{(\ell-k-1)}(x'') \cdot v_{n_{i,k}}^{(\ell-k)}(y) dv(y) \\
 &\stackrel{(4.19)}{\leq} C_2 \|\mu_i\| e^{\frac{(M+5)\|c\|}{\varepsilon}} \text{KL}(\pi_{n_{i,k}}^{(\ell-k-1)} | \pi_{n_{i,k}}^{(\ell-k)})
 \end{aligned}$$

This eventually yields

$$\begin{aligned}
 \Delta(\pi^{(\ell)}) &\leq C_1 \cdot C_2 \cdot e^{\frac{(M+5)\|c\|}{\varepsilon}} \sum_{\substack{i \in I: \\ D(i) > 0}} \sum_{k=0}^{D(i)-1} \|\mu_i\| \text{KL} \left(\pi_{n_{i,k}}^{(\ell-k-1)} | \pi_{n_{i,k}}^{(\ell-k)} \right) \\
 &\leq C_1 \cdot C_2 \cdot e^{\frac{(M+5)\|c\|}{\varepsilon}} \sum_{\substack{j \in I: \\ D(j) < M}} \sum_{i \in I} \|\mu_i\| \text{KL} \left(\pi_j^{(\ell-D(j)-1)} | \pi_j^{(\ell-D(j))} \right) \\
 &\leq C_1 \cdot C_2 \cdot e^{\frac{(M+5)\|c\|}{\varepsilon}} \sum_{j \in I} \sum_{k=0}^{M-1} \text{KL} \left(\pi_j^{(\ell-k-1)} | \pi_j^{(\ell-k)} \right) \\
 &= C_1 \cdot C_2 \cdot e^{\frac{(M+5)\|c\|}{\varepsilon}} \cdot \left(\Delta(\pi^{(\ell-M)}) - \Delta(\pi^{(\ell)}) \right)
 \end{aligned}$$

where we used in the second line that $D(j) = k$ if $j = n_{i,k}$ for $j \in I$ (since $n_{i,k}$ is at the k -th step of the shortest path to cell i). This yields the sought-after bound:

$$\Delta(\pi^{(\ell)}) \leq \frac{2MN e^{\frac{(6M+7)\|c\|}{\varepsilon}}}{\mu_{\min}^{2M+1}} \cdot \left(\Delta(\pi^{(\ell-M)}) - \Delta(\pi^{(\ell)}) \right).$$

Convergence of the iterates to the optimal solution follows as in Theorem 1. □

4.3 Numerical worst-case examples

In the preceding Sections we derived bounds on the convergence rate based on various worst-case estimates. But convergence might actually be much faster. In this Section we provide numerical examples that demonstrate that, while the pre-factors in our bounds might not be tight, in the three-cell-case the qualitative dependency of the

Fig. 4 Setup for worst-case examples. Top row: three cells, Example 2, cost c , initial coupling $\pi^{(0)}$ and embedding of $X = Y = \{1, 2, 3\}$ into $\mathbb{R}^2$ such that c becomes the distance function. The domain decomposition algorithm must exchange the masses of the first and third column of $\pi^{(0)}$. Bottom row: cost c and initial coupling $\pi^{(0)}$ for chain setup, Example 4, for $N = 5$

Fig. 5 Results for ε -dependency in three-cell worst-case example, see Example 2. Left: Sub-optimality $\Delta(\pi^{(\ell)})$ over iterations for various values of ε . The values converge to zero linearly. Right: We extract the linear contraction factor $\lambda(\varepsilon)$ from the left plot and compare it to the theoretical bound from Theorem 1. This is visualized best in the form $-\log(\lambda(\varepsilon)^{-1} - 1)$ for which the bound yields $\frac{2\|c\|}{\varepsilon} - \log(q/(1-q))$, see (4.25). In particular, the $1/\varepsilon$ -behaviour is accurate and the pre-factor appears to be off only by a factor of ≈ 2

convergence rate on the parameters ε and the mass in the middle cell $\mu(X_2)$ are accurate; and that convergence slows down in the general case as the maximum distance M from the root cell J_0 in the partition graph increases.

Example 2 (Three cells, dependency on ε) Let $X = Y = \{1, 2, 3\}$ with cost c as given in Fig. 4. This configuration can be obtained with embeddings of X and Y into $\mathbb{R}^2$. Set $\mu = \nu = p \cdot \delta_1 + q \cdot \delta_2 + p \cdot \delta_3$ for $q \in (0, 1)$, $p = \frac{1-q}{2}$. We choose $q = 0.3$. Basic and composite partitions are chosen as in the previous three-cell examples (Example 1): $X_i = \{i\}$ for $i \in I = \{1, 2, 3\}$, $\mathcal{J}_A = \{\{1, 2\}, \{3\}\}$, $\mathcal{J}_B = \{\{1\}, \{2, 3\}\}$.

The unique optimal coupling in the unregularized case is given by $\pi^* = p \cdot \delta_{(1,1)} + q \cdot \delta_{(2,2)} + p \cdot \delta_{(3,3)}$. As initial coupling for the algorithm we pick $\pi^{(0)} := p \cdot \delta_{(1,3)} + q \cdot \delta_{(2,2)} + p \cdot \delta_{(3,1)}$. So the mass in cell X_1 must be moved into cell X_3 and vice versa.

Fig. 6 Results for q -dependency in three-cell worst-case example, see Example 3. Left: Sub-optimality $\Delta(\pi^{(\ell)})$ over iterations for various values of q . The values converge to zero linearly. Right: We extract the linear contraction factor $\lambda(q)$ from the left plot and compare it to the theoretical bound from Theorem 1. This is visualized best in the form $\lambda(q)^{-1} - 1$ for which the bound yields $\exp(-2\|c\|/\varepsilon) \cdot \frac{q}{1-q}$, see (4.26). In particular, the proportionality to q (for small values) is accurate. The pre-factor appears to be off by a factor of ≈ 50

In the unregularized case $\pi^{(0)}$ would be a fixed-point of the iterations, since $\pi^{(0)}$ is partially optimal on $X_{\{1,2\}}$ and $X_{\{2,3\}}$. This provides a simple example that the unregularized domain decomposition algorithm requires additional assumptions for convergence to the global solution.

In the regularized setting, some mass will also be put onto the points $(1, 2)$, $(2, 1)$, $(2, 3)$, and $(3, 2)$ despite the high cost, thus leading to the possibility to move mass between X_1 to X_3 via X_2 and thus eventually solving the problem. But this mass tends to zero exponentially as $\varepsilon \rightarrow 0$ thus resulting in an exponential number of iterations.

Numerical results are summarized in Figure 5. To estimate $\Delta(\pi^{(\ell)})$ we obtain a high precision approximate solution π^* by applying a single Sinkhorn algorithm to the problem. Since this can directly exchange mass between cells 1 and 3, for the single solver this problem is not particularly difficult (the same applies to the other examples in this Section). Let $\lambda(\varepsilon)$ be the empirical contraction factor for the sub-optimality. Theorem 1 yields the bound

$$\lambda(\varepsilon) \leq \left(1 + \exp(-2\|c\|/\varepsilon) \cdot \frac{q}{1-q}\right)^{-1} \Leftrightarrow -\log(\lambda(\varepsilon)^{-1} - 1) \leq \frac{2\|c\|}{\varepsilon} - \log(q/(1-q)). \quad (4.25)$$

As indicated in Fig. 5, the $1/\varepsilon$ -dependency accurately describes the convergence behaviour.

Example 3 (Three cells, dependency on q) We revisit the previous Example 2, but now we fix the regularization parameter $\varepsilon = 10$ and vary the mass q in the middle cell. Intuitively, far from the optimal solution, the mass exchanged between the cells in each iteration should be proportional to q and thus for small q the number of required iterations should be approximately proportional to q^{-1} . This would imply that for small q the contraction ratio $\lambda(q)$ scales as $\exp(-Cq) \approx 1 - Cq$ for some constant C . For small q this matches the behaviour of the bound of Theorem 1, which implies, cf. (4.25),

Fig. 7 Results for N -dependency in chain graph worst-case example, see Example 4. Left: Sub-optimality $\Delta(\pi^{(\ell)})$ over iterations for various values of N . The values converge to zero linearly. Right: Linear contraction factors $\lambda(N)$ extracted from the left plot. Since they approach 1 very quickly, they are visualized best in the form $1 - \lambda(N)$ in log-scale

$$\lambda(q)^{-1} - 1 \geq \exp(-2\|c\|/\varepsilon) \cdot \frac{q}{1-q}. \quad (4.26)$$

Numerical results are summarized in Figure 6. In our setting the proportionality between $\lambda(q)^{-1} - 1$ and q is confirmed numerically. Note that we chose $\varepsilon = 10$ relatively high. For smaller values we observed that the mass exchanged between the cells in each iterations was no longer proportional to q and thus the interplay between the parameters ε and q became more complicated.

Example 4 (*Chain graph, dependency on graph size*) The dependency of the contraction ratio bound of Corollary 2 on the graph structure is more complex. The graph structure determines the numbers M , N and implicitly $\mu_{\min}$ (since mass must be distributed over increasingly many cells). Due to the variety of worst-case estimates that entered into the proof of (4.13) we do not expect it to be particularly tight. Nevertheless, it is easy to confirm numerically that convergence slows down when the graph size increases.

We extend Example 2. For $N \geq 3$ let $X := Y := I := \{1, \dots, N\}$, $X_i := \{i\}$ for $i \in I$. Let $\mu := \nu := \frac{1}{N} \sum_{i \in I} \delta_i$. Set

$$c(i, j) := \begin{cases} 0 & \text{if } i = j, \\ 1 & \text{if } (i, j) \in \{(1, N), (N, 1)\}, \\ 10 & \text{else.} \end{cases}$$

The unique optimal coupling is given by $\pi^* = \frac{1}{N} \sum_{i=1}^N \delta_{(i,i)}$, we choose $\pi^{(0)} = \frac{1}{N} (\sum_{i=2}^{N-1} \delta_{(i,i)} + \delta_{(1,N)} + \delta_{(N,1)})$. c and $\pi^{(0)}$ are illustrated in Fig. 4. For N even, we set

$$\mathcal{J}_A := \{\{1, 2\}, \{3, 4\}, \dots, \{N-1, N\}\}, \quad \mathcal{J}_B := \{\{1\}, \{2, 3\}, \dots, \{N-2, N-1\}, \{N\}\},$$

for N odd we adapt this in the obvious fashion. Finally, let $\varepsilon = 1.4$ (but the qualitative results do not depend on ε). Similar to Example 2, mass from the first and last cell must be exchanged, but this time via an increasing number of intermediate cells. Numerical results for this setting are shown in Fig. 7. We find that $\lambda(N)$ approaches 1 as N increases.

5 Relation to parallel sorting

The linear rates obtained in Theorems 1 and 2 become very slow for small ε . The number of required iterations to achieve prescribed sub-optimality scales like $\exp(C/\varepsilon)$ as ε approaches zero. This rate is comparable to the result by Franklin and Lorenz [10] on the convergence of the discrete Sinkhorn algorithm. They show that the intermediate marginal $\mu^{(\ell)} = P_X \pi^{(\ell)}$ in the Sinkhorn algorithm satisfies $d(\mu^{(\ell)}, \mu) \leq q^\ell \cdot d(\mu^{(0)}, \mu)$ where d denotes Hilbert's projective metric and $q \approx 1 - 4 \exp(-\|c\|/\varepsilon)$ for $\varepsilon \ll \|c\|$.

While we do not claim that the bounds in Theorems 1 and 2 are tight, they at least qualitatively capture the right scaling behaviour with respect to ε , the mass contained in the basic cells, and the (approximate) graph diameter M , as demonstrated with some worst-case examples in Sect. 4.3.

We emphasize however that in Sect. 4 we make virtually no assumptions on the cost function (beyond bounded) and on the structure of the basic and composite partitions (beyond basic cells carrying non-zero mass and a connected partition graph). This freedom is essential for the construction of the examples in Sect. 4.3. Conversely, in practice we are usually not dealing with such restrictive problems, but often need to solve transport problems on low-dimensional domains, with highly structured cost functions, and it is possible to choose highly structured basic and composite partitions. Based on the intuition from Fig. 1 we expect that the algorithm will perform much better in such settings, which is confirmed by numerical experiments in Sect. 7.

Unfortunately, a convergence analysis that leverages the additional geometric structure to obtain better rates is much more challenging than the worst-case bounds from Sect. 4 and thus beyond the scope of this article.

To provide at least some insight we discuss in this Section an example of discrete, one-dimensional, unregularized optimal transport where the domain decomposition algorithm becomes equivalent to the parallel odd-even transposition sorting algorithm. Let

$$X = \{x_1, \dots, x_N\}, \quad Y = \{y_1, \dots, y_N\},$$

for some $N \in \mathbb{N}$. Let μ and ν be the normalized counting measures on X and Y , which can be written as

$$\mu = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}, \quad \nu = \frac{1}{N} \sum_{i=1}^N \delta_{y_i}.$$

and let $c : X \times Y \rightarrow \mathbb{R}_+$ satisfy the strict Monge property [5], i.e.

$$c(x_i, y_j) + c(x_r, y_s) < c(x_i, y_s) + c(x_r, y_j) \quad \text{for} \quad 1 \leq i < r \leq N, \quad 1 \leq j < s \leq N. \quad (5.1)$$

A simple example for such a setting would be when X and Y are (sorted) sets of points in $\mathbb{R}$, $x_i < x_j$, $y_i < y_j$ for $1 \leq i < j \leq N$, and $c(x_i, y_j) = h(x_i - y_j)$ for a strictly convex function h . This covers the Wasserstein distances for $p > 1$.

For this setting we are now interested in the unregularized optimal transport problem

$$\min \left\{ \int_{X \times Y} c \, d\pi \mid \pi \in \Pi(\mu, \nu) \right\}. \quad (5.2)$$

It is easy to see that this is equivalent to solving the linear assignment problem between the sets X and Y with respect to the cost c and that the unique optimal solution is given by $\pi = \frac{1}{N} \sum_{i=1}^N \delta_{(x_i, y_{\sigma(i)})}$ (uniqueness is implied by strict inequality in (5.1)).

Let now a basic partition of X be given by the partition into singletons, i.e. $X_i = \{x_i\}$ for $i \in I = \{1, \dots, N\}$ and the two composite partitions $\mathcal{J}_A$ and $\mathcal{J}_B$ will be set to

$$\mathcal{J}_A = (\{1, 2\}, \{3, 4\}, \dots, \{N-1, N\}), \quad \mathcal{J}_B = (\{1\}, \{2, 3\}, \dots, \{N-2, N-1\}, \{N\}) \quad (5.3)$$

where for simplicity of exposition we assume that N is even. Further, let $\pi^{(0)} = \frac{1}{N} \sum_{i=1}^N \delta_{(x_i, y_{\sigma(i)})}$ where $\sigma : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$ is some permutation. We can now conceive running an unregularized version of Algorithm 1 where we replace the regularized transport solver in line 8 by the unregularized counterpart. It can quickly be seen that (due to the strict Monge property) the solution of the (now unregularized) sub-problems in line 8 is always unique and at each iteration $\pi^{(\ell)}$ corresponds to a permutation $\sigma^{(\ell)}$. Each composite cell (with the exception of the cells $\{1\}$ and $\{N\}$ in $\mathcal{J}_B$) consists of two consecutive elements of X and in line 8 it is tested whether the two corresponding assigned points in Y should be kept in the current order or flipped. Therefore, in this special case the domain decomposition algorithm reduces to the odd-even transposition parallel sorting algorithm [14, Exercise 37] which is known to converge in $O(N)$ iterations. It is not hard to see that the analysis for odd-even transposition sort could be adapted to the continuous one-dimensional strictly-Monge case to establish convergence in $O(1/\mu_{\min})$ steps.

Unfortunately, the techniques do not generalize to multiple dimensions and an accurate analysis of the convergence speed of the domain decomposition algorithm in that setting is therefore still an open problem. Nevertheless, we are optimistic that the expected performance for ‘structured problems’ in higher dimensions will be much better than the worst-case bounds from Sect. 4, which is confirmed numerically in Sect. 7.

6 Practical large-scale implementation

Algorithm 2 Practical implementation of entropic domain decomposition

Input:

- initial feasible basic Y -marginals $(v_i^0)_{i \in I}$ (i.e. satisfying $\|v_i^0\| = \|\mu_i\|$ for $i \in I$ and $\sum_{i \in I} v_i^0 = \nu$)
- scaling factor $u^0 \in L_+^\infty(X, \mu)$

Output:

- final feasible basic Y -marginals $(v_i)_{i \in I}$
- final partial scaling factors $(u_J)_{J \in \mathcal{J}_A}, (u_J)_{J \in \mathcal{J}_B}$ on both composite partitions

```

1: for all  $i \in I$  do:  $v_i \leftarrow v_i^0$ 
2: for all  $J \in \mathcal{J}_A$  do:  $u_{A,J} \leftarrow u_{\perp} X_J$ 
3: for all  $J \in \mathcal{J}_B$  do:  $u_{B,J} \leftarrow u_{\perp} X_J$ 
4:  $\ell \leftarrow 0$ 
5: repeat
6:    $\ell \leftarrow \ell + 1$ 
7:   if ( $\ell$  is odd) then  $C \leftarrow A$  else  $C \leftarrow B$                                 // select the partition
8:   for all  $J \in \mathcal{J}_C$  do                                                            // solve over each composite cell
9:      $v_{\text{cell}} \leftarrow \sum_{i \in J} v_i$                                               // compute  $Y$ -marginal on cell
10:     $\pi_{\text{cell}}, u_{C,J} \leftarrow \text{SINKHORN}(\mu_J, v_{\text{cell}}, u_{C,J})$                   // approx. sol. w. Sinkhorn algorithm
11:    for all  $i \in J$  do:  $v_i \leftarrow \text{PY}(\pi_{\text{cell} \llcorner (X_i \times Y))$                 // extract new basic partial marginals
12:     $\text{BALANCEMEASURES}((v_i)_{i \in J})$ 
13:    for all  $i \in J$  do:  $\text{TRUNCATE}(v_i)$ 
14:  end for
15: until stopping criterion
16: return  $(v_i)_{i \in I}, (u_{A,J})_{J \in \mathcal{J}_A}, (u_{B,J})_{J \in \mathcal{J}_B}$ 

```

Comment: The functions SINKHORN and BALANCEMEASURES are described in Section 6.2. The function TRUNCATE refers to truncation up to a small threshold to obtain sparse vectors, as described in Section 6.1, item (ii).

The variable π_{cell} is a local variable that is not stored beyond the for-loop over $\mathcal{J}_C$.

6.1 Computational adaptations

Our strongest motivation for studying the entropic domain decomposition algorithm is to develop an efficient numerical method. The statement given in Algorithm 1 is mathematically abstract and served as preparation for theoretical analysis. In this Section we describe a more concrete computational implementation, formalized in Algorithm 2. We summarize the main modifications:

- Instead of storing the full coupling $\pi^{(\ell)}$, we only store the Y -marginal $v_i^{(\ell)} := \text{PY}\pi_i^{(\ell)}$ of each basic cell $i \in I$. This suffices to compute the Y -marginals of $\pi^{(\ell)}$ on each composite cell, as done in Algorithm 1, line 7. The partial couplings $\pi_i^{(\ell)}$ can quickly be recovered with the scaling factors u , see (iv) below.
- To further reduce the memory footprint, we truncate $v_i^{(\ell)}$ after each iteration, i.e. we only keep a sparse representation of the entries above a small threshold, which we set to 10^{-15} . In theory this invalidates the global density bound (4.15) in Lemma 6 and thus jeopardizes convergence of the method. In practice our threshold

- is low enough such that $v_i^{(\ell)}$ on neighbouring basic cells have strong overlap and thus, based on the intuition from Sect. 5, on geometric problems we still expect convergence to an approximate solution. This can be confirmed numerically via the primal-dual gap (Sect. 7.3 and Table 1).
- (iii) We replace ‘abstract exact solution’ of the entropic transport problems on the composite cells by the Sinkhorn algorithm (compare Algorithm 1, line 8 and Algorithm 2, line 10). This means, we need to be able to handle approximate solutions for the cell problems. We describe how this is done in Sect. 6.2.
 - (iv) We explicitly keep track of the partial scaling factors u on the composite cells of $\mathcal{J}_A$ and $\mathcal{J}_B$. This allows accurate initialization of the Sinkhorn algorithm on each composite cell, in particular after the first few iterations. Further, it enables us to quickly reconstruct the partial couplings on composite and basic cells. Finally, we can use them to estimate the sub-optimality via the primal-dual gap, see Sect. 6.3.
 - (v) We combine Algorithm 2 with a multi-scale scheme and the ε -scaling heuristic described in [23]. This is critical for computational efficiency, as it allows to obtain a good (and sparse) initial coupling π^0 (or its partial Y -marginals on the basic cells) and drastically reduces the number of required iterations. Some more details are given in Sect. 6.4.

6.2 Sinkhorn solver and handling approximate cell solutions

The Sinkhorn solver at line 10 in Algorithm 2 only returns an approximate solution. This requires some additional post-processing. Note that for efficiency, we initialize the solver with the scaling factor $u_{C,J}$ from the previous iteration, i.e. we start the solver with a Y -iteration (Remark 1).

As stopping criterion of the Sinkhorn solver we use the L^1 -marginal error of the X -marginal after the Y -iteration. We require that this error is smaller than $\|\mu_J\| \cdot \text{Err}$ where Err is some global error threshold. In this way, the summed L^1 -error of all X -marginals over all composite cells is bounded by Err after each iteration.

We terminate the algorithm after a Y -iteration, i.e. the Y -marginals are satisfied exactly, $P_Y \pi_{\text{cell}} = v_{\text{cell}}$, and after the for-loop in line 14 one has $\sum_{i \in I} v_i = v$, which is crucial for the validity of the domain decomposition scheme. Since $\mu_J = \mu \llcorner X_J$ is constant throughout the algorithm it is easier to handle a deviation between $P_X \pi_{\text{cell}}$ and μ_J .

However, when π_{cell} is not contained in $\Pi(\mu_J, v_{\text{cell}})$, after computing the basic cell marginals in line 11 it may occur that $\|v_i\| \neq \|\mu_i\|$. This means, that while $\sum_{i \in J} v_i = v_{\text{cell}}$, the masses of the v_i may not exactly match the masses in their basic cells. This would cause problems in subsequent iterations. Therefore, the function BALANCEMEASURES determines for each basic cell $i \in J$ the deviation $\|v_i\| - \|\mu_i\|$ and then moves mass between the different $(v_i)_{i \in J}$ to compensate the deviations, while preserving their non-negativity and their sum. This can be done efficiently directly in the sparse data structure and without adding new support points.

Table 1 Summary of average performance (with standard deviation) of the domain decomposition algorithm and a single sparse Sinkhorn algorithm [23] for comparison

Image size	64 × 64	128 × 128	256 × 256	512 × 512	1024 × 1024	2048 × 2048
Domain decomposition (stopping criterion: $L^1 - \text{error} \leq 10^{-4}$)						
<i>Total runtime in seconds</i>						
Time (1 worker)	(3.7 ± 0.0)E0	(1.5 ± 0.0)E1	(5.9 ± 0.0)E1	(2.4 ± 0.0)E2	(1.1 ± 0.0)E3	(4.8 ± 0.0)E3
Time (2 workers)	(2.1 ± 0.0)E0	(7.9 ± 0.1)E0	(3.1 ± 0.0)E1	(1.2 ± 0.0)E2	(5.7 ± 0.1)E2	(2.5 ± 0.0)E3
Time (3 workers)	(1.7 ± 0.0)E0	(5.7 ± 0.1)E0	(2.1 ± 0.0)E1	(8.5 ± 0.1)E1	(3.8 ± 0.0)E2	(1.7 ± 0.0)E3
Time (4 workers)	(1.5 ± 0.0)E0	(4.6 ± 0.1)E0	(1.7 ± 0.0)E1	(6.5 ± 0.0)E1	(2.9 ± 0.0)E2	(1.3 ± 0.0)E3
Time (5 workers)	(1.5 ± 0.0)E0	(4.0 ± 0.1)E0	(1.4 ± 0.0)E1	(5.4 ± 0.0)E1	(2.4 ± 0.0)E2	(1.1 ± 0.0)E3
Time final layer (5 w.)	(8.7 ± 0.4)E-1	(2.6 ± 0.0)E0	(1.0 ± 0.0)E1	(4.1 ± 0.0)E1	(2.0 ± 0.0)E2	(8.5 ± 0.0)E2
<i>Times for specific functions in seconds (5 workers)</i>						
SINKHORN	(1.3 ± 0.0)E0	(3.6 ± 0.1)E0	(1.2 ± 0.0)E1	(4.9 ± 0.0)E1	(2.3 ± 0.0)E2	(1.0 ± 0.0)E3
BALANCEMEASURES	(1.3 ± 0.0)E-1	(2.6 ± 0.1)E-1	(7.3 ± 0.3)E-1	(2.2 ± 0.1)E0	(7.0 ± 0.1)E0	(2.6 ± 0.0)E1
TRUNCATE	(1.4 ± 0.1)E-2	(5.7 ± 0.5)E-2	(2.1 ± 0.1)E-1	(8.1 ± 0.2)E-1	(3.2 ± 0.1)E0	(1.3 ± 0.0)E1
Refinement	(3.3 ± 0.2)E-2	(1.0 ± 0.1)E-1	(3.2 ± 0.1)E-1	(1.2 ± 0.0)E0	(5.0 ± 0.0)E0	(2.1 ± 0.0)E1
total number of entries in all $(v_i)_{i \in I}$						
Maximum	(8.4 ± 0.2)E4	(3.5 ± 0.0)E5	(1.4 ± 0.0)E6	(5.5 ± 0.0)E6	(2.1 ± 0.0)E7	(8.3 ± 0.0)E7
Final	(1.7 ± 0.0)E4	(6.7 ± 0.1)E4	(2.6 ± 0.0)E5	(1.0 ± 0.0)E6	(3.9 ± 0.0)E6	(1.6 ± 0.0)E7
<i>Solution quality</i>						
Relative PD gap	(1.9 ± 1.6)E-4	(1.0 ± 0.6)E-4	(1.0 ± 0.7)E-4	(4.3 ± 3.2)E-5	(1.0 ± 1.3)E-4	(6.2 ± 3.0)E-5
L^1 marginal error X	(5.1 ± 0.3)E-5	(5.0 ± 0.2)E-5	(4.9 ± 0.1)E-5	(4.7 ± 0.1)E-5	(4.8 ± 0.1)E-5	(4.4 ± 0.1)E-5
L^1 marginal error Y	(7.8 ± 25.8)E-7	(1.2 ± 1.3)E-8	(3.7 ± 0.2)E-8	(1.5 ± 0.1)E-7	(4.4 ± 0.1)E-7	(1.1 ± 0.0)E-6
<i>Single sparse multi-scale Sinkhorn (stopping criterion: $L^\infty - \text{error} \leq 10^{-7}$)</i>						
Time	(3.6 ± 0.6)E0	(2.7 ± 0.7)E1	(1.3 ± 0.6)E2	(2.8 ± 1.2)E2	-	-
# kernel entries, max	(6.1 ± 0.4)E5	(2.8 ± 0.1)E6	(1.2 ± 0.0)E7	(5.1 ± 0.0)E7	-	-
# kernel entries, final	(9.3 ± 0.3)E4	(4.0 ± 0.1)E5	(1.7 ± 0.0)E6	(6.9 ± 0.1)E6	-	-
L^1 marginal error X	(1.1 ± 0.3)E-4	(3.1 ± 0.9)E-4	(7.5 ± 2.1)E-4	(2.6 ± 0.7)E-3	-	-
Relative dual score	(-1.8 ± 2.3)E-4	(-1.4 ± 0.5)E-3	(-4.9 ± 1.7)E-3	(-6.9 ± 2.9)E-3	-	-

For each image size results are averaged over 45 problem instances (Sect. 7.1). The relative PD gap and the relative dual score are defined in (7.1) and (7.2)

6.3 Estimating the primal-dual gap

Explicitly evaluating the primal-dual gap between (2.4) and (2.8) (see Proposition 2 (ii)) is a useful tool to verify that a high-quality approximate solution was found. As primal candidate we can use the current primal iterate. We can sum up the contribution of each π_{cell} to (2.4) in the for-loop starting at line 8. The partial replacement with older iterates, as in Lemma 5, was merely required for the proof strategy, for numerical evaluation it is not necessary.

As a dual candidate we could use $\hat{u}$ and $\hat{v}$ as constructed in (4.16, 4.17). It is not hard to see that upon perfect convergence of the domain decomposition algorithm, the functions $u_{J_A}^{(\ell)}$ and $u_{J_B}^{(\ell-1)}$ on two overlapping composite cells $J_A \in \mathcal{J}_A$, $J_B \in \mathcal{J}_B$ would only differ by a constant factor on their overlap $J_A \cap J_B \neq \emptyset$. Consequently, this ‘glued’ together $(\hat{u}, \hat{v})$ would be a dual maximizer (cf. proof strategy of Proposition 4). In practice, before perfect convergence, we find however that this construction yields unnecessarily bad dual estimates. We now sketch a slightly more sophisticated construction (along with an explanation why the above construction does not work well numerically).

Inspired by the partition graph, Definition 7, and the construction (4.16, 4.17) we now define another weighted graph. The vertex set is given by the composite cells $\mathcal{J}_A$. Between two vertices $J_1, J_2 \in \mathcal{J}_A$ there will be an edge if there exists some $J_B \in \mathcal{J}_B$ such that $J_1 \cap J_B \neq \emptyset$ and $J_2 \cap J_B \neq \emptyset$. The corresponding edge-weight will be set to

$$q_{(J_1, J_2)} := \int_{J_1 \cap J_B} \frac{u_{A, J_1}}{u_{B, J_B}} d\mu \cdot \int_{J_2 \cap J_B} \frac{u_{B, J_B}}{u_{A, J_2}} d\mu.$$

Note that the weight depends on the orientation of the edge and that there may be multiple edges between the same two vertices. Similar to (4.16, 4.17) we could now fix a root cell $J_0 \in \mathcal{J}_A$ and for any $J \in \mathcal{J}_A$ we fix some path $(J_0, J_1, \dots, J_n = J)$ and set

$$\hat{u}_J := \prod_{k=0}^{n-1} q_{(J_k, J_{k+1})} \cdot u_{A, J}.$$

Again, after perfect convergence, the $\hat{u}$ obtained by gluing the $\hat{u}_J$ would be a dual maximal X -scaling. The value of $\hat{u}_J$ would not depend on the choice of the path from J_0 to J and when going around a cycle in the graph, the product of the corresponding weight factors will be equal to 1. This means that $\log q_{(J_1, J_2)}$ is a discrete gradient on the graph, i.e. summing $\log q_{(J_1, J_2)}$ around a cycle yields zero.

But before perfect convergence, the weights are inconsistent in the sense that $\log q_{(J_1, J_2)}$ is not a gradient and by picking some arbitrary path from J_0 to J we distribute this inconsistency in an arbitrary way over $\hat{u}_J$. In practice this leads to poor candidates $\hat{u}$, in particular far from the root cell.

We weaken this effect by doing a discrete Helmholtz decomposition of $\log q_{(J_1, J_2)}$. More concretely, we look for a minimizer of

$$\min \left\{ \sum_{\substack{J_1, J_2 \in \mathcal{J}_A: \\ (J_1, J_2) \text{ is graph edge}}} ((V(J_1) - V(J_2)) - \log q_{(J_1, J_2)})^2 \mid V : \mathcal{J}_A \rightarrow \mathbb{R} \right\}.$$

The gradient of the minimizing potential V will be an optimal approximation of the edge weights in a mean-square sense. This problem can be solved quickly via a linear system, also when there are multiple edges between two vertices. We then set $\hat{u}_J := \exp(V_J) \cdot u_{A, J}$. We observe that this distributes the inconsistencies in the weights more evenly over the graph and yields better dual estimates.

Since we truncate the partial Y -marginals (see Sect. 6.1 (ii)), unlike in the theoretical analysis the factors v_J are not defined on all of Y . But once the factors V are available, the v_J can be ‘glued’ together similar to $\hat{u}$.

In practice this method yields small primal-dual gaps (cf. Sect. 7.3), confirming that the entropic domain decomposition works well.

6.4 Multi-scale scheme and ε -scaling

In [23] it is described in detail how a combination of a multi-scale scheme and ε -scaling increases the computational efficiency of the Sinkhorn algorithm and how to implement them numerically. Similarly, the two techniques are crucial for computational efficiency of the domain decomposition scheme. For instance, to initialize Algorithms 1 and 2 one requires a feasible coupling $\pi^0 \in \Pi(\mu, \nu)$ (or its partial Y -marginals on basic cells). Without further prior knowledge, the only canonical candidate would be the product measure $\mu \otimes \nu$. But this is almost certainly far from the optimal coupling, hence requiring many iterations, and it is dense, i.e. requiring a lot of memory in the initial iterations. Instead, as usual, we start by solving the problem on a coarse scale and then refine the (approximate) solutions to subsequent finer scales for initialization.

We now provide a few more details on how this is done. Assume X and Y are discrete, finite sets. Let $\hat{X}$ and $\hat{Y}$ be coarse approximations of X and Y , let $\text{pa} : X \sqcup Y \rightarrow \hat{X} \sqcup \hat{Y}$ be a function that assigns to every element of X and Y their *parent* element in $\hat{X}$ and $\hat{Y}$. Coarse approximations of μ and ν are then given by $\hat{\mu} := \text{pa}_\# \mu$ and $\hat{\nu} := \text{pa}_\# \nu$ where $\text{pa}_\#$ denotes the push-forward of measures. Let $\{\hat{X}_i\}_{i \in \hat{I}}$ be a basic partition of $(\hat{X}, \hat{\mu})$ and let $\{X_i\}_{i \in I}$ be a basic partition of (X, μ) . We choose them in such a way that the fine basic partition can be interpreted as a refinement of the coarse basic partition, i.e. for each $i \in I$ there is some unique $j \in \hat{I}$ such that $\text{pa}(x) \in \hat{X}_j$ for all $x \in X_i$. We denote the function $I \rightarrow \hat{I}$, taking i to j by Pa .

Let now $(\hat{\nu}_j)_{j \in \hat{I}}$ be the tuple of partial marginals as returned by Algorithm 2 on the coarse scale. We use as initialization on the subsequent finer scale the measures

defined via

$$v_i^0(y) := v(y) \cdot \frac{\hat{v}_{\text{Pa}(i)}(\text{pa}(y))}{\hat{v}(\text{pa}(y))} \cdot \frac{\mu(X_i)}{\hat{\mu}(\hat{X}_{\text{Pa}(i)})} \quad (6.1)$$

where by a slight abuse of notation, we write $v(y)$ to mean $v(\{y\})$ for singletons.

Proposition 5 (6.1) provides a valid initialization for Algorithm 2.

Proof We need to show that $\sum_{i \in I} v_i^0 = v$ and $\|v_i^0\| = \|\mu_i\|$ for all $i \in I$. For the former, for $y \in Y$ we find

$$\sum_{i \in I} v_i^0(y) = \sum_{j \in \hat{I}} \sum_{\substack{i \in I: \\ \text{Pa}(i)=j}} v_i^0(y) = v(y) \sum_{j \in \hat{I}} \frac{\hat{v}_j(\text{pa}(y))}{\hat{v}(\text{pa}(y))} \sum_{\substack{i \in I: \\ \text{Pa}(i)=j}} \frac{\mu(X_i)}{\hat{\mu}(\hat{X}_j)} = v(y)$$

where in the third expression we first see that the third factor equals 1, since $\text{pa}^{-1}(\hat{X}_j) = \cup_{i \in I: \text{Pa}(i)=j} X_i$, and subsequently the second factor equals 1, since $\sum_{j \in \hat{I}} \hat{v}_j = \hat{v}$, which is implied by the domain decomposition algorithm on the coarse scale.

Similarly, for the latter, for $i \in I$,

$$\begin{aligned} \|v_i^0\| &= \sum_{\hat{y} \in \hat{Y}} \sum_{\substack{y \in Y: \\ \text{pa}(y)=\hat{y}}} v_i^0(y) = \frac{\mu(X_i)}{\hat{\mu}(\hat{X}_{\text{Pa}(i)})} \sum_{\hat{y} \in \hat{Y}} \frac{\hat{v}_{\text{Pa}(i)}(\hat{y})}{\hat{v}(\hat{y})} \underbrace{\sum_{\substack{y \in Y: \\ \text{pa}(y)=\hat{y}}} v(y)}_{=\hat{v}(\hat{y})} \\ &= \frac{\mu(X_i)}{\hat{\mu}(\hat{X}_{\text{Pa}(i)})} \sum_{\hat{y} \in \hat{Y}} \hat{v}_{\text{Pa}(i)}(\hat{y}) = \frac{\mu(X_i)}{\hat{\mu}(\hat{X}_{\text{Pa}(i)})} \hat{v}_{\text{Pa}(i)}(\hat{Y}) = \mu(X_i) \end{aligned}$$

where we have used the consistency condition $\|\hat{v}_j\| = \|\hat{\mu}_j\| = \hat{\mu}(\hat{X}_j)$ for $j \in \hat{I}$ in the coarse domain decomposition algorithm. $\square$

Of course this scheme can be repeated over multiple successive levels of approximations.

In addition, we refine the scaling factors u from the coarse to the fine level to obtain good initializations. For this, we first ‘glue’ together the final u at the coarse level, as described in Sect. 6.3. Then we do a linear interpolation onto the fine points (which we perform in log-space).

Finally, we remark that when changing the parameter ε during ε -scaling from ε_{old} to ε_{new} , the scalings (u, v) should not be left constant, but instead $(\varepsilon \cdot \log(u), \varepsilon \cdot \log(v))$ should be kept constant, so we set $u_{\text{new}} := \exp(\frac{\varepsilon_{\text{old}}}{\varepsilon_{\text{new}}} \cdot \log u_{\text{old}})$, and similarly for v_{new} , cf. [23, Section 3.2].

Fig. 8 Two example images for size 256×256 , at different layers, $l = 8$ being the original. Yellow represents high, blue low mass density. The combination of strong local mass concentrations, smooth regions and areas with very little mass leads to challenging test problems with non-trivial transport maps (color figure online)

7 Numerical geometric large scale examples

We now demonstrate the efficiency of Algorithm 2 on large-scale geometric problems and show that it compares favorably to the single Sinkhorn algorithm.

7.1 Preliminaries

Hardware and implementation All numerical experiments were performed on a standard desktop computer with an Intel Core i7-8700 CPU with 6 physical cores and 32 GB of RAM, running a standard Ubuntu-based distribution. We developed an experimental implementation for Python 3.7. Parallelization was implemented with MPI, via the module mpi4py.¹ We used a simple master/worker setup with one thread (master) keeping track of the overall problem and submitting the sub-tasks to the other threads (workers). Parallelization was used for solving of the composite cell problems, the BALANCEMEASURES function and data refinement from coarse to fine layers.

Stopping criterion and measure truncation We set the stopping criterion for the internal Sinkhorn solver to $\text{Err} = 10^{-4}$ (see Sect. 6.2). This bounds the L^1 -marginal error of the primal iterates. Partial marginals $(\nu_i)_{i \in I}$ are truncated at 10^{-15} and stored as sparse vectors.

Test data We focus on solving the Wasserstein-2 optimal transport problem, i.e. we set $c(x, y) = \|x - y\|^2$ but the scheme can be applied to arbitrary costs and we expect efficient performance for any cost of the form $c(x, y) = h(x - y)$ for strictly convex h , such that (in the continuous, unregularized setting) the unique optimal plan is concentrated on the graph of a Monge map, see [11].

¹ <https://mpi4py.readthedocs.io/>.

As test data we use 2D images with dimensions $2^n \times 2^n$ for $n = 6$ to $n = 11$, i.e. from 64×64 up to 2048×2048 , the number of pixels per image ranging from 4.1E3 to 4.2E6. Similar to [22,23] the images were randomly generated, consisting of mixtures of Gaussians with random variances and magnitudes. This represents challenging problem data, leading to non-trivial optimal couplings, involving strong compression, expansion and anisotropies, as visualized in Figures 8 and 11. For each problem size we generated 10 test images, i.e. 45 pairwise non-trivial transport problems, to get some estimate of the average performance.

Multi-scale and ε -scaling Each image is then represented by a hierarchy of increasingly finer images of size $2^l \times 2^l$ where l ranges from 3 to n . We refer to l as the *layer*. This induces a simple quad-tree structure over the images where each pixel at layer l is the parent of four pixels at layer $l + 1$, see Sect. 6.4 and [22,23] for more details.

We embed the pixels into $\mathbb{R}^2$ as a Cartesian grid of edge length $\Delta x_n := 1$ between two neighbouring pixels at layer $l = n$ and with edge length $\Delta x_l := 2^{n-l}$ on coarser layers. At each layer l we start solving with $\varepsilon = 2 \cdot \Delta x_l^2$, for which we apply four iterations (two on $\mathcal{J}_A$ and two on $\mathcal{J}_B$). Then we decrease ε by a factor of two and apply two more iterations (one on $\mathcal{J}_A$ and one on $\mathcal{J}_B$). We repeat this (so that $\varepsilon = 0.5 \cdot \Delta x_l^2$) and then refine the image. This leaves the scale of $c(\cdot, \cdot)/\varepsilon$ approximately invariant throughout the layers (cf. [23]). At the finest layer, we perform two additional iterations at $\varepsilon = 0.25$, which implies that the influence of entropic regularization in our final solution is rather weak (cf. Fig. 10).

This means, that for two images of size $2^n \times 2^n$ we only perform $(n - 2) \cdot 8 + 2$ iterations of the domain decomposition algorithm, even though there are 2^{2n-4} basic cells in the X -image (see below) and the graph diameter (Definition 7) is on the order $O(2^{n-2})$. We observe that this logarithmic number of iterations is fully sufficient (cf. Table 1). This is only possible due to the geometric problem structure (Sect. 5) and the multi-scale scheme with ε scaling.

Basic and composite partitions As basic partition we divide each image (at a given layer l) into blocks of $s \times s$ pixels where the cell size s is a suitable divisor of 2^l . The composite partition $\mathcal{J}_A$ is then generated by summarizing the basic cells into 2×2 blocks, the partition $\mathcal{J}_B$ is generated in the same way but with an offset of one basic cell in each direction, i.e. with single basic cells in the corners and 2×1 blocks along the image boundaries. This is visualized in Fig. 2. Each basic cell at layer l is therefore split into four basic cells at layer $l + 1$, thus satisfying the assumption from Sect. 6.4.

The cell size s is an important parameter. For small s the composite cell problems are small (and thus easier), but one must solve a larger number of composite problems, implying more communication overhead in parallelization and requiring more domain decomposition iterations. For large s the number of composite problems is smaller, thus inducing less communication overhead and requiring less iterations, but solving the composite problems becomes more challenging. We found that $s = 4$ yielded the best results for our code. Studying the influence of this parameter in more detail and reducing the computational overhead in the implementation are a relevant direction for future work.

Fig. 9 Primal iterates for the images from Fig. 8, visualized as described in Sect. 7.2, on layer $l = 8$, $\varepsilon = 2$, for $\ell = 0, 1, 2$. The initial coupling is blurry due to the refinement construction via (6.1). For iterates after $\ell = 2$ it is hard to spot further changes

Fig. 10 Primal iterates for the images from Fig. 8, visualized as described in Sect. 7.2, on layer $l = 8$, for $\varepsilon = 2$ ($\ell = 4$ on that layer) and $\varepsilon = 0.25$ ($\ell = 10$), see Sect. 7.1 for a description of ε -scaling and iteration schedule. For $\varepsilon = 2$ some blur between neighbouring cells is visible, for $\varepsilon = 0.25$ it is virtually gone (up to some inevitable overlap due to the discretization)

7.2 Visualization of primal iterates

To get an impression of the algorithm's behaviour we visualize the evolution of the primal iterates $\pi^{(\ell)}$. This is difficult since in our examples the iterates are measures on $\mathbb{R}^4$. But the partition structure provides some assistance. We assign colors to the basic cells in an alternating pattern (e.g. a checkerboard tiling). Denote by col_i the color assigned to basic cell $i \in I$ (col_i could just be a RGB vector). We can then approximately visualize a coupling $\pi \in \Pi(\mu, \nu)$ with partial marginals $\nu_i := P_Y \pi_i$ via the color image on Y generated by $\sum_{i \in I} \text{col}_i \cdot \frac{d\nu_i}{d\nu}$. When π is a deterministic coupling, e.g. introduced by a Monge map $T : X \rightarrow Y$, then region $T(X_i)$ will be colored with col_i , i.e. the image would look like a Cartesian grid that was morphed by the map T . When π is non-deterministic, different ν_i can overlap and in these regions the resulting color will be a linear interpolation of the colors col_i , weighted by the ν_i .

Consider the optimal transport problem between the two images from Fig. 8 with size 256×256 . Various aspects of this problem are illustrated in Figs. 9, 10 and 11.

Figure 9 shows the effect of the first two iterations on the finest layer $l = 8$. The initial image is essentially uniform, since the initial marginals ν_i^0 , as generated with (6.1) on basic cells $i \in I$ that were generated by sub-dividing the same coarse basic cell $j \in \hat{I}$ are, up to a constant factor, identical. After the first iteration (on $\mathcal{J}_A$) a grid pattern reemerges, after the second iteration (on $\mathcal{J}_B$), it is hard to spot further changes (cf. Fig. 10, where for $\varepsilon = 2$ the iterate after four iterations is shown).

The effect of ε -scaling within a given layer is visualized in Fig. 10. For $\varepsilon = 2$ the checkerboard pattern is blurred between neighbouring cells. For $\varepsilon = 0.25$, neighbouring cells overlap by at most one pixel, which would also be true in the discrete

Fig. 11 Primal iterates for the images from Fig. 8, visualized as described in Sect. 7.2, on layers $l = 3$ to $l = 8$, for $\varepsilon = 0.5 \cdot \Delta x_l$, where $\Delta x_l = 2^{l-8}$ is the distance between adjacent grid points on a given layer, and $\ell = 8$ on each layer, except on the finest layer, where $\varepsilon = 0.25$ and $\ell = 10$. See Sect. 7.1 for a description of ε -scaling and iteration schedule. During each refinement a basic cell is split into four smaller basic cells and additional details in the optimal coupling become discernible

unregularized setting (recall that our pixels do not all carry the same mass, so the discrete unregularized optimal coupling is usually not deterministic). This demonstrates that the algorithm can be run until very low regularization parameters.

Finally, Fig. 11 demonstrates the evolution of the multi-scale scheme. It shows the final coupling on each layer for the smallest ε used on that layer, just before refinement (except for the finest layer). One can trace how increasingly finer structures of the optimal coupling emerge as the resolution improves.

7.3 Quantitative evaluation

Having convinced ourselves by means of visualizations that the entropic domain decomposition algorithm seems to work as intended, we now turn to a more quantitative analysis of its performance. A comparison to the performance of a single sparse multi-scale Sinkhorn algorithm [23] is provided in Sect. 7.4.

Runtime We measure the runtime of the algorithm on different problem sizes, with the number of worker threads ranging from 1 to 5. For simplicity, for a single worker thread, the implementation still uses two threads (one master, one worker), but the performance cannot be distinguished from a true single-thread implementation, since the master thread is essentially idle in this case. In addition we are interested in the

Fig. 12 Average runtimes of domain decomposition algorithm. The runtime increases approximately linearly (or slightly super-linearly) with the size of the marginal images, allowing the solution of large problems. The runtime of a single Sinkhorn algorithm (with different stopping criterion) is shown for comparison (see Sect. 7.4 for details). Runtime decreases approximately in inverse proportionality with the number of worker threads (up to five workers), indicating efficient parallelization (except for small images where overhead is more significant and thus the effect of parallelization saturates earlier)

Fig. 13 Number of non-zero entries to store the partial marginal list $(v_i)_{i \in I}$ at the finest layer. Maximal number (for largest value of ε) and final number are shown. For comparison the number of non-zero entries in the truncated (stabilized) kernel matrix of the single Sinkhorn algorithm is shown (again, maximal and final number on finest layer). All numbers scale linearly in the image size, indicating that the truncation schemes efficiently reduce memory demand. For domain decomposition the final number of entries per pixel is approximately four, i.e. the coupling is rather concentrated and the effect of entropic blurring is weak. Maximal and final entry numbers are significantly lower with domain decomposition, the maximal number (which represents the memory bottleneck) is almost reduced by a factor 10

time required for solving the final (i.e. finest) layer, and for the times required by the main sub-routines, i.e. the Sinkhorn solver, the measure balancing, truncation and refinement. These are listed in Table 1 and the total runtimes are visualized in Figure 12.

We find that the runtime scales approximately linearly (or at most slightly super-linear) with the size of the marginals which is crucial for solving large problems. The time required for solving the final layer comprises approximately 80% of the total time, indicating that even with a good initialization, solving the finest layer is a challenging problem. The time spent on the SINKHORN sub-solver for the composite cell problems accounts for approximately 96% of the runtime for large problems, which means that other functions cause relatively little overhead. On small problems these ratios are a bit lower, since the overhead of the multi-scale scheme is more significant.

For large problems the runtime is approximately proportional to the inverse number of workers (up to five workers), indicating that parallelization is computationally efficient. Again, for small problems, computational overhead becomes more significant, such that the benefit of parallelization saturates earlier. For large problems we estimate that MPI communication accounts for approximately 3% of time in a domain decomposition iteration.

Sparsity In addition to runtime the memory footprint of the algorithm is an important criterion for practicability. It is dominated by storing the list $(v_i)_{i \in I}$ (after truncation at 10^{-15} , see Sect. 6.1), since only a few small sub-problems are solved simultaneously at any given time. We are therefore interested in the number of non-zero entries in this list throughout the algorithm. This number is usually maximal on the finest layer (most pixels) and for the highest value of ε on that layer (with strongest entropic blur), which therefore constitutes the memory bottleneck of the algorithm.

This maximal number and the final number of entries are listed in Table 1 and visualized in Fig. 13. We find that both scale approximately linear with marginal size. Again, this is a prerequisite for solving large problems. Maximally the algorithm requires approximately 20 entries per pixel in the first image, and approximately 4 per pixel upon convergence. The latter indicates that entropic blur is indeed strongly reduced in the final iterate.

Solution quality The data reported in Table 1 also confirms that the algorithm provides high-quality approximate solutions of the problem. The relative primal-dual (PD) gap is defined as (see (2.9) for the definition of J)

$$\frac{\text{KL}(\pi|K) - J(u, v)}{\text{KL}(\pi|K) - \|K\|}, \quad (7.1)$$

i.e. we divide the primal-dual gap by the primal score (where, for simplicity we subtract the constant term $\|K\|$). The dual variables (u, v) are obtained from the composite duals by gluing, as discussed in Sect. 6.3. We find that this is on the order of 10^{-4} , i.e. a high relative accuracy.

The L^1 -error of the X -marginal of the primal iterate is consistently even smaller than the prescribed bound $\text{Err} = 10^{-4}$ (cf. Sect. 6.2). This happens since we do not

check the stopping criterion for the SINKHORN sub-solver after every iteration. The L^1 -error of the Y -marginal is not (numerically) zero, since a small error accumulates over time during the several additions and subtractions of the partial marginals and in the BALANCEMEASURES function. But since it is still substantially below the X -error we do not consider this to be a problem. Additional stabilization-routines could be added to the algorithm, should it become necessary.

7.4 Comparison to single Sinkhorn algorithm

As a reference we compare the performance of the domain decomposition algorithm to that of a single Sinkhorn algorithm (with adaptive kernel truncation, coarse-to-fine scheme and ε -scaling as in [23]). The results are also summarized in Table 1, and visualized in Figs. 12 and 13. The single Sinkhorn algorithm was limited to sizes up to 512×512 due to numerical stability and memory issues (see below for details).

Both algorithms have several parameters that influence the trade-off between accuracy and runtime. One important parameter is the stopping criterion for the sub-solver Sinkhorn algorithm in the former, and for the global single Sinkhorn algorithm in the latter. Another one is the truncation threshold for partial marginals (cf. Sect. 6.1) in the former, and for the (stabilized) kernel matrix in the latter [23, Section 3.3].

Stopping criterion The stopping criterion for the domain decomposition solver was picked such that the global L^1 -error of the marginals is below $\text{Err} = 10^{-4}$ (cf. Sections 6.2 and 7.1). We observe that if we pick the same criterion for the single Sinkhorn solver, it takes much longer than the domain decomposition method (even without parallelization) and the runtime scales significantly super-linearly with marginal size.

A more detailed analysis reveals that the single Sinkhorn algorithm spends a major part of time on the largest ε on the finest layer, i.e. immediately after the last refinement. We conjecture that this is due to the fact that the single Sinkhorn algorithm represents the primal iterates only implicitly through the dual iterates (u, v) via $(u \otimes v) \cdot K$ (cf. Sect. 2.2). During layer-refinement only the dual variables are refined explicitly. This can introduce quite non-local marginal errors in the primal iterate (i.e. mass must be moved far to correct them) that take many iterations to correct. We tried various heuristics to reduce this effect (e.g. more sophisticated interpolation of the dual variables during refinement) but unfortunately without success.

In contrast, the domain decomposition method keeps explicitly track of primal and dual (partial) iterates, both of which are refined separately during layer transitions. In particular the total mass of the refined primal iterate within each refined basic cell is correct after refinement (cf. (6.1) and Proposition 5) and thus fewer global mass movements are required for correcting these errors.

Therefore, as in [23], numerically we use the L^∞ error as stopping criterion for the single Sinkhorn algorithm. We set the stopping accuracy to 10^{-7} , such that (without parallelization) both algorithms have approximately the same runtime.

Sparsity To perform its iterations the single Sinkhorn algorithm must keep track of a truncated (stabilized) kernel matrix, i.e. one column of non-zero entries for each pixel

of the first marginal. Following [23, Section 3.3, Equation (3.8)] we keep entries of the (stabilized) kernel matrix where $(u \otimes v) \cdot k \geq \theta$ with the choice $\theta = 10^{-10}$.

The domain decomposition method (in the form of Algorithm 2) only needs to store a column of non-zero entries for each basic cell (of size $s \times s$) in the first marginal. Thus, to represent the primal iterate with comparable accuracy, the number of non-zero entries required in the domain decomposition method is reduced by a factor s^{-2} .

In this sense, for the parameter choices in our experiments, the domain decomposition algorithm represents the primal iterates slightly more accurately. At the same time, the number of entries to store the final iterate is reduced to approximately 15% compared to the single Sinkhorn algorithm. The maximum number of entries is even reduced to approximately 11%. On problem size 2048×2048 the single Sinkhorn algorithm ran out of memory on our test machine.

Solution quality As discussed above, to obtain similar runtimes for both algorithms, a less strict stopping criterion had to be chosen for the single Sinkhorn algorithm. Therefore, it produces higher L^1 -marginal errors for the X -marginal that increase with problem size. The Y -marginal error is (numerically) zero, since the algorithm was terminated after a Y -iteration (cf. Remark 1). The relative dual score reported in Table 1 is defined by

$$\frac{J(u_{\text{single}}, v_{\text{single}}) - J(u_{\text{domdec}}, v_{\text{domdec}})}{J(u_{\text{domdec}}, v_{\text{domdec}}) - \|K\|}, \quad (7.2)$$

cf. (7.1), where $(u_{\text{domdec}}, v_{\text{domdec}})$ refers to the (glued) final dual variables of the domain decomposition method (Sect. 6.3) and $(u_{\text{single}}, v_{\text{single}})$ to the final dual iterates of the single Sinkhorn algorithm. The fact that this number is negative indicates that the domain decomposition provides better dual candidates than the single Sinkhorn.

Numerical stability For small regularization the standard Sinkhorn algorithm becomes numerically unstable due to the finite precision of discretized numbers on computers. This can be remedied to a large extent by the absorption technique [23]. But refinement between layers remains a delicate point since the interpolated dual variables from the coarser layer may not provide a perfect initial guess at the finer layer. This can lead to numerically zero or empty columns and rows in the (stabilized) kernel matrix and thus will lead to invalid entries in the scalings (u, v) . While it is possible to fix these cases, on large problems this is relatively cumbersome and computationally inefficient. We therefore decided to abort the algorithm in these cases instead. In our experiments this concerned 3 of the 45 examples of size 512×512 and the majority of the 1024×1024 examples (which we have therefore removed from the comparison).

These issues may also arise in the domain decomposition algorithm but there they are localized to small problems on composite cells which are easy to fix (e.g. by temporarily increasing ε). We have implemented a corresponding safe-guard. In all our examples it was invoked on two 1024×1024 instances.

Summary Clearly the domain decomposition method is superior to the single Sinkhorn algorithm in various aspects. In comparison, the domain decomposition method (without parallelization) produces more accurate primal and dual iterates in comparable time, with approximately 11% of the memory. It runs more reliably on large problems, since numerical hickups of the Sinkhorn algorithm can be fixed locally. On top of that the domain decomposition method allows for real non-local parallelization, whereas the Sinkhorn algorithm can only be parallelized very locally over the matrix-vector multiplications.

8 Conclusion and outlook

In this article we studied Benamou's domain decomposition algorithm for optimal transport in the entropy regularized setting. The key observation is that the regularized version converges to the unique globally optimal solution under very mild assumptions. We proved linear convergence of the algorithm w.r.t the KL-divergence and illustrated the (potentially very slow) rates with numerical examples.

We then argued that on 'geometric problems' the algorithm should converge much faster. To confirm this experimentally, we discussed a practical, efficient version of the algorithm with reduced memory footprint, numerical safeguards for approximation errors, primal-dual certificates and a coarse-to-fine scheme. With this we were able to compute optimal transport between images with ≈ 4 megapixels on a standard desktop computer, matching two megapixel images in approximately four minutes. Even without parallelization, the algorithm compared favourably to a single Sinkhorn algorithm in terms of runtime, memory, solution quality and numerical reliability. With parallelization the runtime was efficiently reduced with little communication overhead.

A practical open question for future research is the development of more sophisticated and efficient implementations, e.g. using GPUs and/or multiple machines, with more sophisticated parallelization structure (e.g. no single unit needs to keep track of the full problem). This should then be tested on 3D problems and with more general cost functions. In the course of this, parameters such as the basic cell size should be studied in more detail.

On the theoretical side the convergence analysis on 'geometric problems' is an interesting challenge, to provide a more thorough understanding for the low number of required iterations. We conjecture that a 'discrete Helmholtz decomposition' as discussed in Sect. 6.3 may become relevant for this. The relation between optimal transport and the Helmholtz decomposition was already noted in [4] and exploited numerically in [12].

Acknowledgements This work was supported by the Emmy Noether programme and the SFB TRR 109 of the DFG.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give

appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Appendix

Proof (Proposition 1) (i) is provided by [8, Theorem 2.1]. Existence of (u^*, v^*) in (ii) is established by [20, Theorem 3]. Since π^* is unique, $u^* \otimes v^*$ is unique K -almost everywhere. Since k is strictly bounded away from zero, this translates to $(\mu \otimes \nu)$ -a.e. uniqueness and finally to μ -a.e. and ν -a.e. uniqueness of u^* and v^* up to re-scaling.

The integral equations in (iii) follow from the marginal constraints $\pi^* \in \Pi(\hat{\mu}, \hat{\nu})$ and the definition of K .

Using now that $\exp(-\|c\|_\infty/\varepsilon) \leq k(x, y) \leq 1$ for all $(x, y) \in X \times Y$, one easily deduces that for all $x \in X$,

$$\exp(-\|c\|_\infty/\varepsilon) \int_Y v^*(y) \, d\nu(y) \leq \int_Y k(x, y) v^*(y) \, d\nu(y) \leq \int_Y v^*(y) \, d\nu(y),$$

which gives $v^* \in L^1(Y, \nu)$ and $\|v^*\|_{L^1(Y, \nu)} > 0$ (otherwise there would exist no u^* to satisfy the left equation of (2.5), recall that we assumed $\|\hat{\mu}\| > 0$). Analogously we see that $u^* \in L^1(X, \mu)$ and $\|u^*\|_{L^1(X, \mu)} > 0$, so that (2.6) and (2.7) follow from (2.5). Since $\frac{d\hat{\mu}}{d\mu}$ and $\frac{d\hat{\nu}}{d\nu}$ were assumed to be (essentially) bounded this implies $u^* \in L_+^\infty(X, \mu)$ and $v^* \in L_+^\infty(Y, \nu)$. Finally, by the first inequality of (2.6),

$$\int_X |\log u^*| \, d\hat{\mu} = \int_X |\log u^*| \frac{d\hat{\mu}}{d\mu} \, d\mu \leq \|v^*\|_{L^1(Y, \nu)} \int_X |\log u^*| u^* \, d\mu < \infty$$

where the last inequality is due to the μ -essential upper bound on u^* and the lower-boundedness of $\mathbb{R}_+ \ni s \mapsto s \log s$ (with limit 0 for $s = 0$). By this, and an analogous argument for v^* , we find $\log u^* \in L^1(X, \hat{\mu})$, $\log v^* \in L^1(Y, \hat{\nu})$. $\square$

Proof (Proposition 2)

(i). Since $u \in L_+^\infty(X, \mu)$, $\log u$ is also μ -essentially bounded from above, thus the first integral is either finite or $-\infty$. The same holds for the second term. The third and fourth terms of J are always finite. Hence, the value of J is unambiguously defined and contained in $[-\infty, \infty)$.

(ii). First observe that $\text{KL}(\pi|K) > -\infty$ and $J(u, v) < +\infty$ by definition. If $J(u, v) = -\infty$ or $\text{KL}(\pi|K) = +\infty$ there is nothing to prove. Thus, let us assume both values are finite and directly compute

$$\begin{aligned} \text{KL}(\pi|K) - J(u, v) &= \int_{X \times Y} \left[\varphi\left(\frac{d\pi}{dK}\right) + (u \otimes v - 1) \right] dK - \underbrace{\left(\int_X \log u \, d\hat{\mu} + \int_Y \log v \, d\hat{\nu} \right)}_{= \int_{X \times Y} (\log u \oplus \log v) \, d\pi} \\ &= \int_{X \times Y} \left[\varphi\left(\frac{d\pi}{dK}\right) + (u \otimes v - 1) - (\log u \oplus \log v) \right] dK \end{aligned}$$

$$\begin{aligned}
&= \int_{X \times Y: \substack{u \otimes v > 0}} \left[\varphi \left(\frac{d\pi}{dK} \right) + \underbrace{\varphi^* (\log u \oplus \log v)}_{=u \otimes v - 1} - (\log u \oplus \log v) \cdot \frac{d\pi}{dK} \right] dK \\
&\quad + \int_{X \times Y: \substack{u \otimes v = 0}} \underbrace{[\varphi(0) - 1]}_{=0} dK \geq 0
\end{aligned}$$

Here we used that $J(u, v) > -\infty$ implies $u > 0$ $\hat{\mu}$ -a.e. and $v > 0$ $\hat{\nu}$ -a.e., so $\pi \in \Pi(\hat{\mu}, \hat{\nu})$ has no mass in the region of the second integral (where $u \otimes v = 0$). The integrand in the first integral is non-negative by the Fenchel–Young inequality and we added the second (vanishing) integral to emphasize that all contributions were considered.

(iii). Let π^* be of the described form with scalings (u^*, v^*) . Arguing as in Proposition 1 we find that $\log u^* \in L^1(X, \hat{\mu})$, $\log v^* \in L^1(Y, \hat{\nu})$ and consequently

$$\text{KL}(\pi^* | K) = \int_X \log u^* d\hat{\mu} + \int_Y \log v^* d\hat{\nu} - \|\pi^*\| + \|K\| = J(u^*, v^*)$$

and so by (ii), π^* and (u^*, v^*) are optimal for (2.4) and (2.8). Existence of such a π^* is provided by Proposition 1, and therefore, maximizers for (2.8) exist. $\square$

Proof (Lemma 1) (i). The result follows from Proposition 1 ((iii)) applied to $\hat{\mu} = \mu_J$ and $\hat{\nu} = \nu_J^{(\ell)}$. Hence, $\log u_J^{(\ell)} \in L^1(X, \mu_J)$ and $\log v_J^{(\ell)} \in L^1(Y, \nu_J^{(\ell)})$, and in particular $\|u_J^{(\ell)}\|_{L^1(X, \mu)} > 0$ and $\|v_J^{(\ell)}\|_{L^1(Y, \nu)} > 0$. The bounds on $u_J^{(\ell)}$ and $v_J^{(\ell)}$ follow from (2.6) and (2.7) in the form

$$\begin{aligned}
\frac{1}{\|v_J^{(\ell)}\|_1} \frac{d\mu_J}{d\mu}(x) &\leq u_J^{(\ell)}(x) \leq \frac{\exp(\|c\|_\infty/\varepsilon)}{\|v_J^{(\ell)}\|_1} \frac{d\mu_J}{d\mu}(x) \quad \text{for } \mu_J\text{-a.e. } x \in X_J, \\
\frac{1}{\|u_J^{(\ell)}\|_1} \frac{d\nu_J^{(\ell)}}{d\nu}(y) &\leq v_J^{(\ell)}(y) \leq \frac{\exp(\|c\|_\infty/\varepsilon)}{\|u_J^{(\ell)}\|_1} \frac{d\nu_J^{(\ell)}}{d\nu}(y) \quad \text{for } \nu\text{-a.e. } y \in Y,
\end{aligned}$$

combined with $(d\mu_J/d\mu)(x) = 1$ for μ_J -a.e. $x \in X_J$ and $(d\nu_J^{(\ell)}/d\nu)(y) \leq 1$ for ν -a.e. $y \in Y$. Observe that the bounding constants, in this case $\underline{C} = \min(1/\|v_J^{(\ell)}\|_1, 1/\|u_J^{(\ell)}\|_1)$ and $\bar{C} = \max(\exp(\|c\|_\infty/\varepsilon)/\|v_J^{(\ell)}\|_1, \exp(\|c\|_\infty/\varepsilon)/\|u_J^{(\ell)}\|_1)$, generally depend on the scaling choice for the duals (Proposition 1 ((ii))) and are not universal over cells and iterations.

(ii). First observe that the scaling factors $u_J^{(\ell)}, v_J^{(\ell)}$ satisfy

$$\begin{cases} u_J^{(\ell)}(x) \int_Y k(x, y) v_J^{(\ell)}(y) d\nu(y) = \frac{d\mu_J}{d\mu}(x) & \text{for } \mu\text{-a.e. } x \in X, \\ v_J^{(\ell)}(y) \int_X k(x, y) u_J^{(\ell)}(x) d\mu(x) = \frac{d\nu_J^{(\ell)}}{d\nu}(y) & \text{for } \nu\text{-a.e. } y \in Y, \end{cases} \quad (8.1)$$

thanks to Proposition 1 (iii) with $\hat{\mu} = \mu_J$ and $\hat{\nu} = \nu_J^{(\ell)}$. The first equation in (8.1) yields

$$\frac{u_J^{(\ell)}(x_1)}{u_J^{(\ell)}(x_2)} = \frac{\int_Y \exp(-c(x_2, y)/\varepsilon) v_J^{(\ell)}(y) d\nu(y)}{\int_Y \exp(-c(x_1, y)/\varepsilon) v_J^{(\ell)}(y) d\nu(y)}$$

for μ_J -a.e. $x_1, x_2 \in X_J$, again using that $(d\mu_J/d\mu)(x_1) = (d\mu_J/d\mu)(x_2) = 1$. Thus, (ii) follows by noting that $c(x_2, y) \geq c(x_1, y) - \|c\|$.

(iv). For the upper bound we combine (8.1) and (3.3) on the basic cell i , to obtain

$$\begin{aligned} u_i^{(\ell)}(x) \cdot v_i^{(\ell)}(y) &= u_i^{(\ell)}(x) \cdot \frac{dv_i^{(\ell)}}{dv}(y) \left(\int_X k(x', y) u_i^{(\ell)}(x') d\mu_i(x') \right)^{-1} \\ &\leq u_i^{(\ell)}(x) \cdot \frac{dv_i^{(\ell)}}{dv}(y) \left(\exp(-2\|c\|/\varepsilon) \cdot u_i^{(\ell)}(x) \cdot \int_X d\mu_i(x') \right)^{-1} \leq \frac{\exp(2\|c\|/\varepsilon)}{\|\mu_i\|} \end{aligned}$$

where we have used $k(x', y) \geq \exp(-\|c\|/\varepsilon)$ in the first inequality and $(dv_i^{(\ell)}/dv)(y) \leq 1$ for ν -a.e. $y \in Y$ in the second. The lower bound works almost entirely analogous by using $k(x', y) \leq 1$.

(v). With $j \in J \in \mathcal{J}^{(\ell)}$, we recall

$$\nu_J^{(\ell)} = \sum_{k \in J} P_Y(\pi^{(\ell-1)} \llcorner (X_k \times Y)) = \sum_{k \in J} \nu_k^{(\ell-1)} \geq \nu_j^{(\ell-1)}. \quad (8.2)$$

Since $i, j \in J \in \mathcal{J}^{(\ell)}$ our convention is that $v_i^{(\ell)} = v_j^{(\ell)} = \nu_J^{(\ell)}$. Combining (3.3) and (8.2), one obtains for $(\mu_i \otimes \mu_j \otimes \nu)$ -a.e. $(x, x', y) \in X_i \times X_j \times Y$

$$\begin{aligned} u_i^{(\ell)}(x) \cdot v_i^{(\ell)}(y) &= \int_X u_i^{(\ell)}(x) \cdot v_i^{(\ell)}(y) d\mu_J(\hat{x}) \geq \exp(-\|c\|/\varepsilon) \cdot \nu_J^{(\ell)}(y) \cdot \int_X u_J^{(\ell)}(\hat{x}) k(\hat{x}, y) d\mu_J(\hat{x}) \\ &= \frac{\exp(-\|c\|/\varepsilon)}{\|\mu_J\|} \cdot \frac{dv_J^{(\ell)}}{dv}(y) \geq \exp(-\|c\|/\varepsilon) \cdot \frac{dv_j^{(\ell-1)}}{dv}(y) \\ &= \exp(-\|c\|/\varepsilon) \cdot \nu_j^{(\ell-1)}(y) \int_X u_j^{(\ell-1)}(\hat{x}) k(\hat{x}, y) d\mu_j(\hat{x}) \\ &\geq \exp(-3\|c\|/\varepsilon) \|\mu_j\| \cdot u_j^{(\ell-1)}(x') \cdot v_j^{(\ell-1)}(y). \end{aligned}$$

Eventually, (3.6) follows with similar arguments: we use the first two lines in the derivation here above and $\exp(-\|c\|/\varepsilon) \leq k(x, y) \leq 1$ to obtain, for ν -a.e. $y \in Y$

$$\frac{dv_i^{(\ell)}}{dv}(y) = \int_X u_i^{(\ell)}(x) v_i^{(\ell)}(y) k(x, y) d\mu_i(x) \geq \exp(-2\|c\|/\varepsilon) \|\mu_i\| \cdot \frac{dv_j^{(\ell-1)}}{dv}(y).$$

□

Lemma 7 (Approximate triangle inequality for KL) *Let $L \in (0, \infty)$, $n \in \mathbb{N}$. Then for all $(s_1, \dots, s_n) \in [0, L]^n$ one finds*

$$\varphi\left(\prod_{k=1}^n s_k\right) \leq C \cdot \sum_{k=1}^n \varphi(s_k) \quad \text{for } C = n \cdot \max\{2, L^{n-1}\}. \quad (8.3)$$

Proof For convenience introduce the functions $S, A, B : \mathbb{R}_+^n \rightarrow \mathbb{R}$,

$$S(s) = \prod_{k=1}^n s_k, \quad A(s) = \varphi(S(s)), \quad B(s) = \sum_{k=1}^n \varphi(s_k).$$

Note that A and B are zero if and only if $s = (1, \dots, 1)$. In this case (8.3) holds for any $C \in \mathbb{R}$. When one of the s_k equals zero, the bound holds for $C \geq 1$. So, to obtain a feasible constant C we obtain an upper bound on

$$\sup \left\{ \frac{A(s)}{B(s)} \mid s \in (0, L]^n, s \neq (1, \dots, 1) \right\}. \quad (8.4)$$

Consider now the case where $S(s) \geq 1$ and $s_k < 1$ for some k . In this case we find that

$$\frac{\partial A}{\partial s_k}(s) = \varphi'(S(s)) \frac{S(s)}{s_k} = \log(S(s)) \frac{S(s)}{s_k} \geq 0, \quad \frac{\partial B}{\partial s_k}(s) = \varphi'(s_k) = \log(s_k) < 0.$$

In this case we can increase $\frac{A(s)}{B(s)}$ by increasing s_k .

Likewise, assume $0 < S(s) < 1$, $s_k > 1$. Then we find

$$\frac{\partial A}{\partial s_k}(s) < 0, \quad \frac{\partial B}{\partial s_k}(s) > 0.$$

So we can increase $\frac{A(s)}{B(s)}$ by decreasing s_k . In summary, we can restrict the supremum (8.4) to the cases $[S(s) \geq 1, s_k \geq 1]$ and $[S(s) < 1, 0 < s_k \leq 1]$.

Supremum over $[S(s) \geq 1, s_k \geq 1]$: with the additional constraint we can rewrite (8.4) as

$$\sup \left\{ \sup \left\{ \frac{A(s)}{B(s)} \mid s \in [1, L]^n, S(s) = \tilde{S} \right\} \mid \tilde{S} \in (1, L^n] \right\}$$

where we can ignore the case $\tilde{S} = 1$ since then $s = (1, \dots, 1)$. By reparametrizing $s_k = \exp(r_k)$ we find:

$$\sup \left\{ \sup \left\{ \frac{\varphi(\tilde{S})}{\sum_{k=1}^n \varphi(\exp(r_k))} \mid r \in [0, \log L]^n, \sum_{k=1}^n r_k = \log \tilde{S} \right\} \mid \tilde{S} \in (1, L^n] \right\}$$

Since $r \mapsto \sum_{k=1}^n \varphi(\exp(r_k))$ is convex on $\mathbb{R}_+^n$ and invariant under permuting the order of the arguments r_k , the inner supremum is attained when setting $r_k = \frac{1}{n} \log \tilde{S}$ and thus one finds (by reparametrizing $\tilde{S} = \tilde{s}^n$)

$$\sup \left\{ \frac{\varphi(\tilde{s}^n)}{n \cdot \varphi(\tilde{s})} \mid \tilde{s} \in (1, L] \right\}$$

Observe now that

$$g(\tilde{s}) := \varphi(\tilde{s}^n) - n^2 \tilde{s}^{n-1} \varphi(\tilde{s}) \leq 0 \quad \text{for all } \tilde{s} > 1$$

because $g(1) = 0$ and $g'(\tilde{s}) = -n^2(n-1)\tilde{s}^{n-2}\varphi(\tilde{s}) \leq 0$ for all $\tilde{s} > 1$. In particular,

$$\frac{\varphi(\tilde{s}^n)}{n\varphi(\tilde{s})} \leq n\tilde{s}^{n-1} \leq nL^{n-1} \quad \text{for all } \tilde{s} \in (1, L]$$

which yields the conclusion.

Supremum over $[S(s) < 1, 0 < s_k \leq 1]$: for $s \in [0, 1]^n$ let $r_k := 1 - s_k$ and set $R = \sum_{k=1}^n r_k$. Then one finds $S(s) \geq \max\{0, 1 - R\}$ and, since φ is decreasing on $[0, 1]$, $A(s) = \varphi(S(s)) \leq \varphi(\max\{0, 1 - R\})$. Similarly for B , with convexity of φ and Jensen's inequality, one has

$$B(s) = \sum_{k=1}^n \varphi(s_k) \geq n \cdot \varphi\left(\frac{1}{n} \sum_{k=1}^n s_k\right) = n \cdot \varphi\left(1 - \frac{1}{n} R\right).$$

This allows us to rewrite (8.4) for the second case as

$$\sup \left\{ \frac{A(s)}{B(s)} \mid s \in [0, 1]^n, S(s) < 1 \right\} \leq \sup \left\{ \frac{\varphi(\max\{0, 1 - R\})}{n \cdot \varphi(1 - \frac{1}{n} R)} \mid R \in (0, n] \right\}$$

Since the denominator is increasing in R and the numerator is constant on $R \in [1, n]$, we can restrict to $R \in (0, 1]$ and prove that the supremum is bounded by $2n$. To do so, we introduce

$$g(R) := \varphi(1 - R) - 2n^2 \varphi\left(1 - \frac{1}{n} R\right) \quad \text{for } R \in [0, 1].$$

To conclude it suffices to prove that $g(R) \leq 0$ for all $R \in [0, 1]$ and any $n \geq 2$. Indeed, we have $g(0) = 0$, $g(1) < 0$ and the derivative $g'(R) = -\log(1 - R) + 2n \log\left(1 - \frac{1}{n} R\right)$ is such that $g'(0) = 0$, $g'(R) \rightarrow +\infty$ as $R \rightarrow 1^-$ and g' has a unique zero on $(0, 1)$ at $R_n = \frac{n}{2n-1}$. We conclude that in $(0, 1)$ the function g has a unique critical point (a minimum) attained at R_n , thus $g(R) \leq 0$ for all $R \in [0, 1]$. $\square$

References

1. Altschuler, J., Weed, J., Rigollet, P.: Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In: *Advances in Neural Information Processing Systems*, vol. 30 (NIPS 2017) (2017)
2. Benamou, J.D.: A domain decomposition method for the polar factorization of vector-valued mappings. *SIAM J. Numer. Anal.* **32**(6), 1808–1838 (1994)
3. Benamou, J.D., Froese, B.D., Oberman, A.M.: Numerical solution of the optimal transportation problem using the Monge–Ampère equation. *J. Comput. Phys.* **260**(1), 107–126 (2014)
4. Brenier, Y.: Polar factorization and monotone rearrangement of vector-valued functions. *Commun. Pure Appl. Math.* **44**(4), 375–417 (1991)
5. Burkhard, R.E., Klinz, B., Rudolf, R.: Perspectives of Monge properties in optimization. *Discrete Appl. Math.* **70**(2), 95–161 (1996)
6. Carlier, G., Duval, V., Peyré, G., Schmitzer, B.: Convergence of entropic schemes for optimal transport and gradient flows. *SIAM J. Math. Anal.* **49**(2), 1385–1418 (2017)
7. Cominetti, R., San Martín, J.: Asymptotic analysis of the exponential penalty trajectory in linear programming. *Math. Program.* **67**, 169–187 (1992)
8. Csiszar, I.: I-divergence geometry of probability distributions and minimization problems. *Ann. Probab.* **3**(1), 146–158 (1975). <https://doi.org/10.1214/aop/1176996454>
9. Cuturi, M.: Sinkhorn distances: lightspeed computation of optimal transportation distances. In: *Advances in Neural Information Processing Systems*, vol. 26 (NIPS 2013), pp. 2292–2300 (2013)
10. Franklin, J., Lorenz, J.: On the scaling of multidimensional matrices. *Linear Algebra Appl.* **114–115**, 717–735 (1989)
11. Gangbo, W., McCann, R.J.: The geometry of optimal transportation. *Acta Math.* **177**(2), 113–161 (1996)
12. Haker, S., Zhu, L., Tannenbaum, A., Angenent, S.: Optimal mass transport for registration and warping. *Int. J. Comput. Vis.* **60**, 225–240 (2004)
13. Kitagawa, J., Mérigot, Q., Thibert, B.: Convergence of a Newton algorithm for semi-discrete optimal transport. *J. Eur. Math. Soc.* (2019). <https://doi.org/10.4171/JEMS/889>
14. Knuth, D.E.: *The Art of Computer Programming, Volume 3: Sorting and Searching*, 2nd edn. Addison-Wesley Professional, Boston (1998)
15. Léonard, C.: From the Schrödinger problem to the Monge–Kantorovich problem. *J. Funct. Anal.* **262**(4), 1879–1920 (2012)
16. Lévy, B.: A numerical algorithm for L2 semi-discrete optimal transport in 3D. *ESAIM Math. Model. Numer. Anal.* **49**(6), 1693–1715 (2015)
17. Luo, Z.Q., Tseng, P.: On the convergence of the coordinate descent method for convex differentiable minimization. *J. Optim. Theory Appl.* **72**(1), 7–35 (1992)
18. Mérigot, Q.: A multiscale approach to optimal transport. *Comput. Graph. Forum* **30**(5), 1583–1592 (2011)
19. Peyré, G., Cuturi, M.: Computational optimal transport. *Found. Trends Mach. Learn.* **11**(5–6), 355–607 (2019). <https://doi.org/10.1561/22000000073>
20. Rüschendorf, L., Thomsen, W.: Note on the Schrödinger equation and I-projections. *Stat. Probab. Lett.* **17**, 369–375 (1993)
21. Santambrogio, F.: *Optimal Transport for Applied Mathematicians, Progress in Nonlinear Differential Equations and Their Applications*, vol. 87. Birkhäuser, Boston (2015)
22. Schmitzer, B.: A sparse multi-scale algorithm for dense optimal transport. *J. Math. Imaging Vis.* **56**(2), 238–259 (2016)
23. Schmitzer, B.: Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM J. Sci. Comput.* **41**(3), A1443–A1481 (2019). <https://doi.org/10.1137/16M1106018>
24. Schmitzer, B., Schnörr, C.: A hierarchical approach to optimal transport. In: *Scale Space and Variational Methods (SSVM 2013)*, pp. 452–464 (2013)
25. Solomon, J., de Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., Guibas, L.: Convolutional Wasserstein distances: efficient optimal transportation on geometric domains. *ACM Trans. Graph. (Proc. SIGGRAPH 2015)* **34**(4), 1–11 (2015)

26. Villani, C.: Optimal Transport: Old and New, Grundlehren der mathematischen Wissenschaften, vol. 338. Springer, Berlin (2009)

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