Locally Correct Interleavings Between Merge Trees

Thijs Beurskens
TU Eindhoven, The Netherlands

Tim Ophelders
TU Eindhoven, The Netherlands

Bettina Speckmann
TU Eindhoven, The Netherlands

Kevin Verbeek
TU Eindhoven, The Netherlands

Abstract

Merge trees are typically used as a topological summary of scalar fields. To analyze e.g. a time-varying scalar field via its merge tree representation, one needs a suitable method to match and compare two merge trees. We consider the interleaving distance as a method to match and compare merge trees. An *interleaving* between two merge trees consists of two maps, one in each direction. These maps must satisfy ancestor relations and hence introduce a "shift" between points and their image. An optimal interleaving minimizes the maximum shift; the interleaving distance is the value of this shift. However, to study the evolution of merge trees, we need not only a number but also a meaningful matching between the two trees. The two maps of an optimal interleaving induce a matching, but due to the bottleneck nature of the interleaving distance, this matching fails to capture local similarities between the trees. In this paper we hence propose a notion of local optimality for interleavings. To do so, we define the *residual interleaving distance*, a generalization of the interleaving distance that allows additional constraints on the maps. This allows us to define *locally correct interleavings*, which use a range of shifts across the two merge trees that reflect the local similarity well. We give a constructive proof that a locally correct interleaving always exists.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Interleaving distance, merge trees, local correctness, matchings, topological data analysis

Digital Object Identifier 10.4230/LIPIcs.SoCG.2026.11

Related Version *Full Version*: `https://arxiv.org/abs/2512.16474`

Funding *Thijs Beurskens*: Supported by the Dutch Research Council (NWO) under project no. OCENW.M20.089.
Tim Ophelders: Supported by the Dutch Research Council (NWO) under project no. VI.Veni.212.260.

1 Introduction

Terrains that vary over time are a common data type in various application areas, such as scientific computing or geographic information science. Typically such data sets are very large; to analyze them efficiently one often needs a compact abstraction that captures salient features. Topological data analysis (TDA) provides several *topological descriptors*, such as persistence diagrams, Reeb graphs, and Morse-Smale complexes, that can serve as abstractions for 2D, 3D, or higher-dimensional terrains. In this paper we focus in particular on *merge trees*, which are graph-based descriptors that encode the evolution of connected components of sublevel or superlevel sets of a terrain (see Figure 1). In the context of time-varying terrains, they have been used for feature tracking [7], cluster detection [16], and various other tasks in scientific visualization [17].

42nd International Symposium on Computational Geometry (SoCG 2026).
Editors: Hee-Kap Ahn, Michael Hoffmann, and Amir Nayyeri; Article No. 11; pp. 11:1–11:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

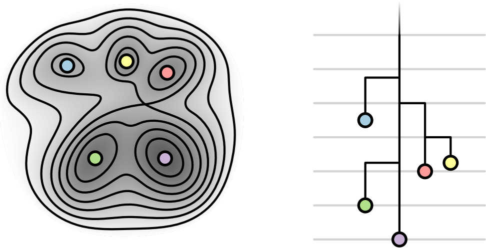

Figure 1 A terrain and its merge tree; leaves and internal vertices represent minima and saddles.

To analyze temporal sequences of terrains one does not only need a compact representation, but also efficient ways to compare these representations over time. For merge trees in particular, several distance measures have been proposed, including the (local) edit distance [13, 14], the Wasserstein distance [11], and the interleaving distance [8]. However, a distance measure returns only a numeric value that quantifies the similarity of two merge trees T_1 and T_2. In many application scenarios, we would also like to track the merge trees over time and hence need to understand which parts of T_1 most closely correspond to parts of T_2. That is, we are interested in a matching between T_1 and T_2.

The interleaving distance in fact relies on a type of matching between the two trees: the so-called *interleaving*. An interleaving is a pair of ancestor-preserving maps, one from T_1 to T_2 and one from T_2 to T_1, whose compositions send points to ancestors: if a point x of T_1 is mapped to a point y of T_2, then y must be mapped to an ancestor of x, and vice versa. The *interleaving distance* is the maximum *shift* – the distance between a point and its image – minimized over all interleavings between the trees. We give exact definitions of the interleaving distance and interleavings in Sections 2 and 3.

Two merge trees typically admit many optimal interleavings. The interleaving distance is a bottleneck measure and hence a given optimal interleaving might be locally quite "loose": the shift might exceed what is needed and thereby fail to capture local similarities between the trees; see Figure 2. In the following we are hence developing a notion of locally optimal interleavings which allows us to distinguish between "tight" and "loose" optimal interleavings.

Contributions. Our work is inspired by the concept of locally correct Fréchet matchings [3]; the *Fréchet distance* is a bottleneck measure for curves based on a matching between points on the curves. A locally correct Fréchet matching is then an optimal matching that remains optimal even when considering any two matched subcurves. This is a desirable feature, however, it cannot be directly reproduced for the interleaving distance. First of all, there might not even be a one-to-one correspondence between the points of T_1 and T_2. Furthermore,

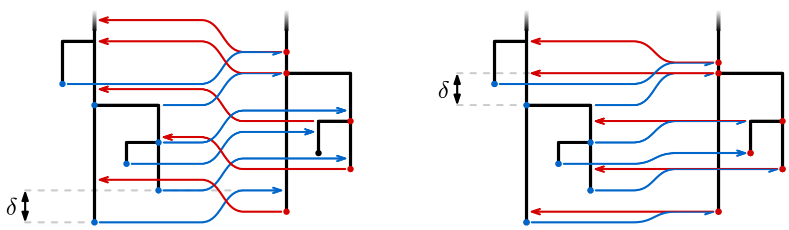

Figure 2 Two optimal interleavings between two merge trees (both with shift δ). Intuitively, the right interleaving induces a tighter and more meaningful matching than the left interleaving.

two optimal interleavings between different pairs of subtrees do not necessarily combine into an optimal interleaving between the two complete trees. As a first step we hence introduce a novel generalization of the interleaving distance that can handle additional constraints on the maps. Specifically, we aim to use constraints that force a point x to be mapped to an ancestor of a particular point y within the other tree. Over interleavings that satisfy such constraints, we aim to minimize the maximum shift over points x that are either unconstrained or map to a strict ancestor of their corresponding point y. We refer to the resulting distance as the *residual interleaving distance*. In Section 3, we provide formal definitions and show that an optimal interleaving exists for any constraints.

In Section 4 we then propose a definition for *locally correct interleavings* that builds upon the residual interleaving distance. Specifically, we define a restriction of an interleaving $\mathcal{I} = (\alpha, \beta)$ to two subsets S_1 of T_1 and S_2 of T_2 as the restriction of α to S_1 together with the restriction of β to S_2. We call $\mathcal{I}$ locally correct if it minimizes the residual interleaving distance with respect to any restriction of $\mathcal{I}$. We constructively prove that a locally correct interleaving always exists. All omitted proofs can be found in the full version.

Related work. The interleaving distance between merge trees was defined by Morozov, Beketayev, and Weber [8] as a tool to study the stability of merge trees. The authors of [1] establish a connection between the interleaving distance and the Gromov-Hausdorff distance, and they argue that approximating the interleaving distance within a factor of 3 is NP-hard. They also describe an $\mathcal{O}(\min(n, \sqrt{rn}))$-approximation algorithm, where n is the total size of the trees and r is the ratio of the longest to the shortest edge. Touli and Wang [15] provide an equivalent definition of the interleaving distance in terms of a single map and use this characterization to design an FPT-algorithm for computing the interleaving distance exactly. More recently, the authors of [6] show that the interleaving distance can be formulated in terms of a related metric [9] for labeled merge trees, also known as phylogenetic trees.

Several attempts have been made to develop versions of the interleaving distance that can be used in practical applications. In particular, the results in [6] gave rise to heuristic algorithms for computing the interleaving distance [4] and for comparing terrains [16, 18]. Pegoraro [10] uses yet another reformulation of the interleaving distance, in terms of couplings, as the basis of a heuristic for computing the distance via linear integer programming. Lastly, the authors of [2] introduce a variant of the interleaving distance that imposes total orders on the merge trees, which can be computed efficiently via a connection to the Fréchet distance between one-dimensional curves.

2 Preliminaries

Let T be a rooted tree, that is, a tree with one vertex identified as the *root*. In the remainder of this paper, we identify T with a *topological realization*: we represent each edge of T by a unit segment $[0, 1]$ and connect these segments according to the adjacencies of their corresponding edges in T. We distinguish elements of the topological realization T, which we refer to as *points* of T, from vertices of the combinatorial tree, which we denote by $V(T)$. A *merge tree* is a pair (T, f), where T is a finite rooted tree, and $f\colon T \to \mathbb{R} \cup \{\infty\}$ is a continuous *height function* on T that is strictly increasing towards the root and such that the root has height ∞. When clear from context, we use T to refer to the pair (T, f).

For two points x_1 and x_2 of T, we say x_1 is a *descendant* of x_2, written $x_1 \preceq x_2$, if there exists an f-monotonically increasing path in T from x_1 to x_2. If furthermore $x_1 \neq x_2$ then x_1 is a *strict descendant* of x_2. We say x_2 is a *(strict) ancestor* of x_1 if x_1 is a (strict) descendant of x_2. Lastly, for a point x and a height value $h \geq f(x)$, we use $x|^h$ to denote the unique ancestor of x at height $f(x|^h) = h$ (see Figure 3a).

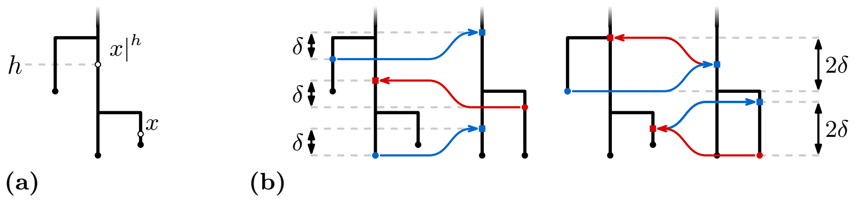

Figure 3 **(a)** The point $x|^h$ is the ancestor of x at height h. **(b)** A pair of δ-compatible maps. Left: **C1** and **C2** ensure that each point is mapped to a point of the other tree δ higher. Right: **C3** and **C4** ensure that both compositions map points to their ancestor 2δ higher.

Compatible maps. The original definition of the interleaving distance between two merge trees (T_1, f_1) and (T_2, f_2) was given by Morozov, Beketayev, and Weber [8].

▶ **Definition 1** ([8]). *Fix a value $\delta \geq 0$. Two continuous maps $\alpha\colon T_1 \to T_2$ and $\beta\colon T_2 \to T_1$ are* δ-compatible *between T_1 and T_2 if for all $x \in T_1$ and all $y \in T_2$:*

(C1) $f_2(\alpha(x)) = f_1(x) + \delta$,
(C2) $f_1(\beta(y)) = f_2(y) + \delta$,
(C3) $\beta(\alpha(x)) = x|^{f_1(x)+2\delta}$, *and*
(C4) $\alpha(\beta(y)) = y|^{f_2(y)+2\delta}$.

The interleaving distance *$d(T_1, T_2)$ between T_1 and T_2 is the infimum δ for which there exist δ-compatible maps (see Figure 3b).*

Critical values. It has been shown [1, 15] that the interleaving distance between any two merge trees attains a value in a set of *critical values*. Specifically, let $\Delta := \Delta(T_1, T_2)$ be the union of three sets Δ_1, Δ_2, and Δ_3, with

$$\begin{aligned}
\Delta_1 &:= \{|f_1(v) - f_2(w)| : v \in V(T_1), w \in V(T_2)\},\\
\Delta_2 &:= \{\tfrac{1}{2}|f_1(v_1) - f_1(v_2)| : v_1, v_2 \in V(T_1)\},\\
\Delta_3 &:= \{\tfrac{1}{2}|f_2(w_1) - f_2(w_2)| : w_1, w_2 \in V(T_2)\}.
\end{aligned} \tag{1}$$

▶ **Lemma 2** ([1, 15]). *For any two merge trees T_1 and T_2, it holds that $d(T_1, T_2) \in \Delta(T_1, T_2)$.*

3 Generalizing the interleaving distance

The interleaving distance is based on a pair of maps. As noted in [1], we can relax the requirements on these maps to allow varying height differences between points and their images. Building on this relaxed definition, we introduce a generalized version of the interleaving distance that accommodates constraints on the maps we consider. Specifically, each constraint specifies that a point x must be mapped to an ancestor of a particular point y within the other tree. Our goal is then to find the optimal pair of maps that "extends" a given set of constraints. We call the corresponding distance the *residual interleaving distance*. To define it formally, we first introduce *partial up-maps*, which represent the allowed maps between the two merge trees. In Section 3.1 we relax Definition 1 using partial up-maps. In Section 3.2 we then give the formal definition of the residual interleaving distance.

Partial up-maps. Let (T_1, f_1) and (T_2, f_2) be two merge trees. An *arrow* from T_1 to T_2 is any pair of points $(x, y) \in T_1 \times T_2$ such that $f_1(x) \leq f_2(y)$. The *shift* of an arrow (x, y) is the height difference between x and y, denoted $\sigma(x, y) := f_2(y) - f_1(x)$. Consider a subset of points $S \subseteq T_1$. A map $\varphi\colon S \to T_2$ is a *partial up-map* from S to T_2 if (i) $f_2(\varphi(x)) \geq f_1(x)$

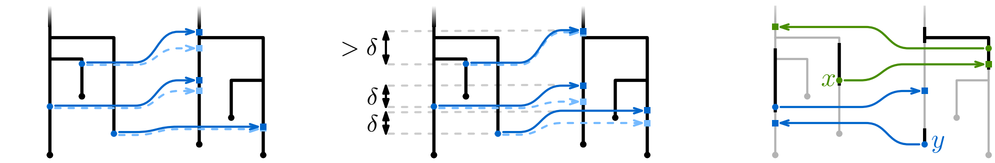

Figure 4 **(a)** A partial up-map (non-dashed) that extends another partial up-map (dashed). **(b)** A partial up-map φ (dashed) and the map $\varphi^{\uparrow}[\delta]$ (non-dashed). **(c)** A partial interleaving.

for all points x of S and (ii) it *preserves ancestors*, that is, $x_1 \preceq x_2$ implies $\varphi(x_1) \preceq \varphi(x_2)$ for all points x_1, x_2 of S. We use $\mathrm{dom}(\varphi) \coloneqq S$ to denote the *domain* of φ. For each point x of S, the pair $(x, \varphi(x))$ is an arrow. The *graph* of φ is the set of arrows of φ, denoted $A_\varphi \coloneqq \{(x, \varphi(x)) : x \in S\}$. We define the *shift* of the partial up-map φ as the supremum shift over all arrows of φ:

$$\sigma(\varphi) \coloneqq \sup\{\sigma(a) : a \in A_\varphi\}. \tag{2}$$

Given an arrow (x, y) from S to T_2, we say φ *extends* (x, y) if $\varphi(x) \succeq y$, and we say φ *uses* (x, y) if $\varphi(x) = y$. Similarly, for a subset $S' \subseteq S$ and a partial up-map φ' from S' to T_2, we say φ *extends* or *uses* φ' if, respectively, φ extends or uses all arrows of φ' (Figure 4a).

For any $\delta \geq 0$, we can define a partial up-map $\varphi^{\uparrow}[\delta]$ from S to T_2 that extends φ and such that the shift of each arrow is at least δ. Specifically, for each point x of S, let $h_x \coloneqq \max(f_1(x) + \delta, f_2(\varphi(x)))$ and define $\varphi^{\uparrow}[\delta](x) \coloneqq \varphi(x)|^{h_x}$ (see Figure 4b).

▶ **Lemma 3.** *For any $\delta \geq 0$, the map $\varphi^{\uparrow}[\delta]$ is a partial up-map that extends φ.*

Proof. Let $\varphi' \coloneqq \varphi^{\uparrow}[\delta]$. By construction, we have $f_2(\varphi'(x)) \geq f_1(x)$. To see that φ' preserves ancestors, take two points x_1, x_2 of S such that x_2 is an ancestor of x_1. Since φ preserves ancestors, $\varphi(x_2)$ is an ancestor of $\varphi(x_1)$. If $\varphi'(x_1) = \varphi(x_1)$, we get $\varphi'(x_1) \preceq \varphi(x_2) \preceq \varphi'(x_2)$. Otherwise $\varphi'(x_1)$ is a strict ancestor of $\varphi(x_1)$. By construction, $f_2(\varphi'(x_1)) = f_1(x_1) + \delta \leq f_1(x_2) + \delta \leq f_2(\varphi'(x_2))$. Both $\varphi'(x_1)$ and $\varphi'(x_2)$ are ancestors of $\varphi(x_1)$, so $\varphi'(x_2)$ is an ancestor of $\varphi'(x_1)$. So φ' is a partial up-map, and by construction also extends φ. ◀

We define arrows and up-maps from a subset S of T_2 to T_1 symmetrically.

Partial interleavings. Fix two subsets $S_1 \subseteq T_1$ and $S_2 \subseteq T_2$. A *partial interleaving* $\mathcal{P}$ between S_1 and S_2 is a pair of partial up-maps $\varphi \colon S_1 \to T_2$ and $\psi \colon S_2 \to T_1$ such that for all points x of S_1 and all points y of S_2 it holds that if $\varphi(x) \preceq y$ then $\psi(y) \succeq x$, and if $\psi(y) \preceq x$ then $\varphi(x) \succeq y$ (Figure 4c). The *shift* of $\mathcal{P}$, denoted $\sigma(\mathcal{P})$, is the maximum of the shift of φ and the shift of ψ. The *graph* of $\mathcal{P}$ is the union of the graph of φ and the graph of ψ, denoted $A_{\mathcal{P}} \coloneqq A_\varphi \cup A_\psi$. Given another partial interleaving $\mathcal{P}' = (\varphi', \psi')$, we say $\mathcal{P}$ *extends* $\mathcal{P}'$ if φ extends φ' and ψ extends ψ'. Similarly, $\mathcal{P}$ *uses* $\mathcal{P}'$ if φ uses φ' and ψ uses ψ'.

▶ **Lemma 4.** *For any $\delta \geq 0$, the pair $(\varphi^{\uparrow}[\delta], \psi^{\uparrow}[\delta])$ is a partial interleaving that extends (φ, ψ).*

Proof. Denote $\varphi' \coloneqq \varphi^{\uparrow}[\delta]$ and $\psi' \coloneqq \psi^{\uparrow}[\delta]$. By Lemma 3, the maps φ' and ψ' are partial up-maps that extend φ and ψ, respectively. Therefore, it suffices to show that the pair (φ', ψ') is a partial interleaving. Let x be a point of S_1 and y be a point of S_2 such that $\varphi'(x) \preceq y$. By construction, the point $\varphi'(x)$ is an ancestor of $\varphi(x)$, so y is also an ancestor of $\varphi(x)$. Since (φ, ψ) is a partial interleaving, the point $\psi(y)$ must be an ancestor of x. As $\psi'(y)$ is an ancestor of $\psi(y)$, we get $\psi'(y) \succeq x$. Symmetrically, if $\psi'(y) \preceq x$ then $\varphi'(x) \succeq y$. ◀

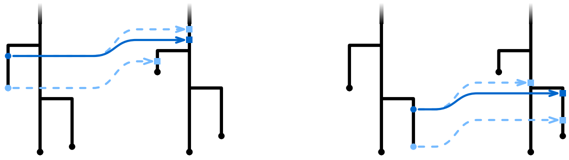

Figure 5 The shift of the dashed interleaving is not equal to a critical value, so we can obtain another interleaving (non-dashed) with strictly smaller shift by "pushing down" the image of each point. If we push below a vertex (right) we carefully need to choose the correct subtree.

3.1 Complete Interleavings

If $S_1 = T_1$, we call a partial up-map $\alpha\colon S_1 \to T_2$ a *complete up-map*, or simply an *up-map*, from T_1 to T_2. Symmetrically, if $S_2 = T_2$, we call a partial up-map $\beta\colon S_2 \to T_1$ a (complete) up-map from T_2 to T_1. Correspondingly, if $S_1 = T_1$ and $S_2 = T_2$, we call the partial interleaving (α, β) a *(complete) interleaving*. The defining property for a complete interleaving then becomes that for all $x \in T_1$ it holds that $\beta(\alpha(x)) \succeq x$ and for all $y \in T_2$ it holds that $\alpha(\beta(y)) \succeq y$.

▶ **Lemma 5.** *$d(T_1, T_2) = \inf\{\delta : T_1$ and T_2 admit a complete interleaving with shift $\delta\}$.*

Optimal interleavings. The interleaving distance is defined as an infimum, but in both [5, 10] the authors prove that the infimum can be replaced by a minimum: there always exists an interleaving with shift $d(T_1, T_2)$ (Theorem 6). We call an interleaving that realizes the interleaving distance an *optimal interleaving*. Both proofs rely on alternative definitions of the interleaving distance, namely in terms of labellings [5] and in terms of couplings [10]. We can show the same statement more directly: if the shift δ of an interleaving $\mathcal{I}$ is not equal to a critical value, we can obtain an interleaving whose shift $\delta' < \delta$ is equal to a critical value. Intuitively, we "push down" each arrow of $\mathcal{I}$ so that its shift becomes δ' (see Figure 5). Later, in Section 3.2, we describe a formal construction to prove a generalization of Theorem 6.

▶ **Theorem 6.** *Any two merge trees T_1 and T_2 admit an optimal interleaving.*

3.2 Residual Interleaving Distance

Fix a partial interleaving $\mathcal{P} = (\varphi, \psi)$. We refer to a (partial) interleaving that extends $\mathcal{P}$ as a (*partial*) *$\mathcal{P}$-extension*. We first observe that there always exists a complete $\mathcal{P}$-extension. Indeed, the maps that take every point of T_1 to the root at ∞ of T_2, and every point of T_2 to the root at ∞ of T_1 form a complete interleaving that extends $\mathcal{P}$.

Our goal is to find a "tightest possible" complete $\mathcal{P}$-extension, that is, an interleaving that extends $\mathcal{P}$ and minimizes the shift of the remaining arrows. For some intuition, consider the example in Figure 6. The map φ consists of a single (dashed) arrow (x, y). Intuitively, the map α_2 describes a "tighter" extension of φ than the map α_1. However, if we only disregard the shift of (x, y), the shift of the remaining arrows of α_1 is equal to the (supremum) shift of the remaining arrows of α_2. The arrow (x, y) implicitly induces a "fan" of arrows that all valid up-maps need to extend. Specifically, any ancestor of x needs to be mapped to an ancestor of y. To distinguish between the maps α_1 and α_2, we therefore not only disregard the shift of (x, y), but of any arrow that is "contained" within such a fan.

Formally, we define the *fan* of an arrow (x, y) as the set of arrows (x', y) for all ancestors x' of x with height at most that of y, written,

$$F[(x, y)] \coloneqq \{(x', y) : x' \in T_1 \text{ such that } x' \succeq x \text{ and } f_1(x') \le f_2(y)\}.$$

The *fan* of a collection of arrows A, denoted $F[A]$, is the union of all fans $F[a]$ for $a \in A$. We refer to the fan $F[A_\varphi]$ as the *fan* of φ.

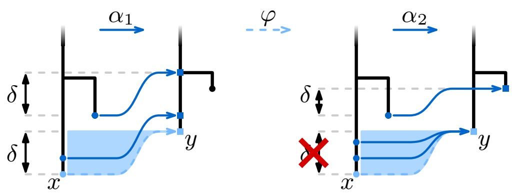

Figure 6 The φ-residual shift of arrows within the fan $F[(x, y)]$ (illustrated with the shaded area) is equal to 0, so the φ-residual shift of α_2 is smaller than the φ-residual shift of α_1.

We use fans to define the *φ-residual shift*, or *φ-shift* for short. Recall that we use A_φ to denote the set of all arrows $(x, \varphi(x))$. For an arrow a from T_1 to T_2, the φ-residual shift of a, denoted $\sigma_\varphi(a)$, is zero if a lies within $F[A_\varphi]$, and is $\sigma(a)$ otherwise. Similar to the definition of shift in Eq. 2, we define the φ-residual shift of an up-map φ', denoted $\sigma_\varphi(\varphi')$, as the supremum φ-residual shift over all arrows of φ'. That is,

$$\sigma_\varphi(x, y) \coloneqq \begin{cases} 0 & \text{if } (x, y) \in F[A_\varphi], \\ \sigma(x, y) & \text{else,} \end{cases} \qquad \sigma_\varphi(\varphi') \coloneqq \sup\{\sigma_\varphi(a) : a \in A_{\varphi'}\}.$$

Consider again the example in Figure 6. We disregard the shifts of all arrows within the fan $F[(x, y)]$. As a result, the φ-shift of α_2 is δ', whereas the φ-shift of α_1 is δ.

We define the *ψ-(residual) shift* of a partial up-map ψ' symmetrically. Finally, for a partial interleaving $\mathcal{P}' = (\varphi', \psi')$, we define the *$\mathcal{P}$-residual* shift, or *$\mathcal{P}$-shift*, of $\mathcal{P}'$ as the maximum of the φ-shift of φ' and the ψ-shift of ψ', denoted $\sigma_\mathcal{P}(\mathcal{P}')$.

▶ **Definition 7.** *The $\mathcal{P}$-residual interleaving distance $d_\mathcal{P}(T_1, T_2)$ between T_1 and T_2 is the infimum δ for which there exists a complete $\mathcal{P}$-extension whose $\mathcal{P}$-residual shift is δ.*

When they are clear from context, we omit the arguments T_1 and T_2 and simply write $d_\mathcal{P}$ to denote the $\mathcal{P}$-residual interleaving distance between T_1 and T_2. We say an extension is *optimal* if its residual shift is equal to the residual interleaving distance.

Critical values and pairs. For the remainder of this section, we consider *finite* partial interleavings $\mathcal{P}$: both $\mathrm{dom}(\varphi)$ and $\mathrm{dom}(\psi)$ consist of finitely many points. Recall from Section 2 that the non-residual interleaving distance always attains a value in the set of critical values $\Delta \coloneqq \Delta(T_1, T_2)$. For the residual interleaving distance, this set Δ is not always sufficient (see Figure 7). Therefore, we extend Δ to a set of $\mathcal{P}$-critical values. Specifically, we "mark" all vertices in T_1 and T_2, and all points x of T_1 and y of T_2 that appear in an arrow (x, y) or (y, x) of $\mathcal{P}$. We call these the *$\mathcal{P}$-critical points* of T_1 and T_2, denoted $C_1[\mathcal{P}]$ and $C_2[\mathcal{P}]$ respectively. Since $\mathcal{P}$ is finite, there are only finitely many critical points. We

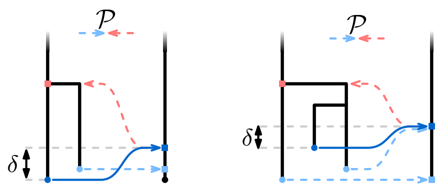

Figure 7 In both examples, the non-dashed arrow is (part of) an optimal interleaving; the value δ is not equal to a critical value. It is equal to a $\mathcal{P}$-critical value.

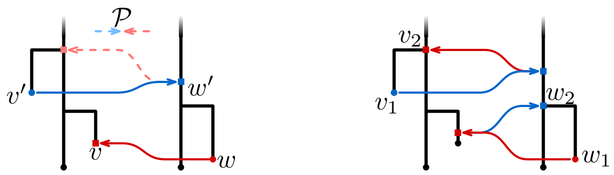

Figure 8 There are two types of critical pairs: arrow critical pairs (w, v) and (v', w') (left) and zigzag critical pairs (v_1, v_2) and (w_1, w_2) (right). The pair (v', w') correspond to a $\mathcal{P}$-critical value.

define a set $\Delta_1[\mathcal{P}]$ that contains all height differences between critical points:

$$\Delta_1[\mathcal{P}] \coloneqq \{|f_1(v) - f_2(w)| : v \in C_1[\mathcal{P}], w \in C_2[\mathcal{P}]\}.$$

We define the set of $\mathcal{P}$-*critical values* as the set $\Delta[\mathcal{P}](T_1, T_2) \coloneqq \Delta_1[\mathcal{P}] \cup \Delta_2 \cup \Delta_3$, where Δ_2 and Δ_3 are as defined in Eq. (1). Each critical value corresponds to a pair of critical points; we refer to such a pair b as a $\mathcal{P}$-*critical pair*. We distinguish two types; see Figure 8.

- If $b = (v, w)$ for a point $v \in C_1[\mathcal{P}]$ and a point $w \in C_2[\mathcal{P}]$, then b is an *arrow critical pair* that corresponds to the critical value $\sigma(v, w)$. We say a partial interleaving (φ, ψ) *uses* b if $\varphi(v) = w$. The case that $b = (w, v)$ is symmetric.
- If $b = (v_1, v_2)$ for two vertices v_1, v_2 in $V(T_1)$, then b is a *zigzag critical pair* that corresponds to the critical value $\frac{1}{2}(f_1(v_2) - f_1(v_1))$. We say a partial interleaving (φ, ψ) *uses* b if there is a point $y \in T_2$ with $f_2(y) = f_1(v_1) + \frac{1}{2}(f_1(v_2) - f_1(v_1))$ such that $\varphi(v_1) = y$ and $\psi(y) = v_2$. The case that $b = (w_1, w_2)$ for two vertices w_1, w_2 in $V(T_2)$ is symmetric.

Similar to Lemma 2, we show that the $\mathcal{P}$-residual interleaving distance always attains a value in the set of $\mathcal{P}$-critical values. A $\mathcal{P}$-critical pair is a *realizing* critical pair if it corresponds to the $\mathcal{P}$-residual interleaving distance. We can additionally show that there always exists an optimal $\mathcal{P}$-extension and that any such extension uses a realizing critical pair.

▶ **Lemma 8.** *Let $\mathcal{P}$ be a finite interleaving. For any complete $\mathcal{P}$-extension with $\mathcal{P}$-shift δ that does not use a critical pair that corresponds to δ, there exists another $\mathcal{P}$-extension whose $\mathcal{P}$-shift is equal to a $\mathcal{P}$-critical value that is strictly less than δ.*

Proof sketch. Consider a $\mathcal{P}$-extension $\mathcal{I} = (\alpha, \beta)$ with $\mathcal{P}$-shift δ, and assume $\mathcal{P}$ does not use a critical pair that corresponds to δ. We modify $\mathcal{I}$ into an extension $\mathcal{I}' = (\alpha', \beta')$ whose $\mathcal{P}$-shift is equal to a $\mathcal{P}$-critical value that is strictly less than δ. Let $\delta' \in \Delta[\mathcal{P}]$ be the greatest value that satisfies $\delta' < \delta$. Intuitively, we modify $\mathcal{I}$ by "pushing down" each arrow (recall Figure 5). We describe the construction of α'; the construction of β' is symmetric.

We assume, without loss of generality, that each arrow of α has shift at least δ'. Otherwise, we replace α by $\alpha^{\uparrow}[\delta']$. Since the residual shift of an arrow is bounded by its shift, and since $\delta' < \delta$, the φ-residual shift of $\alpha^{\uparrow}[\delta']$ is bounded by δ. For $x \in T_1$, let $v_x \in C_1[\mathcal{P}]$ be the highest descendant critical point of x; if $x \in C_1[\mathcal{P}]$, then $v_x = x$. Moreover, let $w_x \in C_2[\mathcal{P}]$ be the highest descendant critical point of $\alpha(v_x)$. Let $h_x \coloneqq \max(f_1(x) + \delta', f_2(w_x))$. We define $\alpha'(x)$ as the ancestor of w_x at height h_x, that is, $\alpha'(x) \coloneqq w_x|^{h_x}$.

It is not hard to prove that α' is an up-map that extends φ. To show that the φ-shift of α' is bounded by δ', we argue that any arrow (x, y) of α' with shift strictly greater than δ' must be inside the fan of φ. Specifically, we show that then also the shift δ^* of the arrow (v_x, w_x) is strictly greater than δ'. The pair (v_x, w_x) is a critical pair; since $\mathcal{I}$ does not use a critical pair that corresponds to δ, it moreover follows that δ^* cannot be equal to δ. So, by choice of δ', we get $\delta^* > \delta$. As a result, (v_x, w_x), and therefore also (x, w_x) must lie within the fan of φ: their residual shift is 0. Symmetrically, the map β' is an up-map that extends ψ and has ψ-shift at most δ'. We can then argue that the resulting pair $\mathcal{I}'$ is an interleaving. ◀

Lemma 8 directly implies Corollaries 9 and 10.

▶ **Corollary 9.** *If $\mathcal{P}$ is finite, then any optimal $\mathcal{P}$-extension uses a realizing critical pair.*

▶ **Corollary 10.** *If $\mathcal{P}$ is finite, then $d_{\mathcal{P}}(T_1, T_2) \in \Delta[\mathcal{P}](T_1, T_2)$.*

Let $\delta_1 \in \Delta[\mathcal{P}]$ be the greatest value that satisfies $\delta_1 \leq d_{\mathcal{P}}$ and let $\delta_2 \in \Delta[\mathcal{P}]$ be the smallest value that satisfies $\delta_2 > d_{\mathcal{P}}$. By definition of $d_{\mathcal{P}}$, for all $\varepsilon > 0$ there is a $\mathcal{P}$-extension with $\mathcal{P}$-shift $d_{\mathcal{P}} + \varepsilon$. In particular, for ε small enough, there is a $\mathcal{P}$-extension $\mathcal{I}$ with $\mathcal{P}$-shift $\delta_1 < \sigma_{\mathcal{P}}(\mathcal{I}) < \delta_2$. In other words, there are no critical pairs that correspond to $\sigma_{\mathcal{P}}(\mathcal{I})$. This means that $\mathcal{I}$ cannot use a critical pair that corresponds to the $\mathcal{P}$-shift of $\mathcal{I}$. By Lemma 8, it then follows that there exists a $\mathcal{P}$-extension with $\mathcal{P}$-shift $\delta_1 = d_{\mathcal{P}}$.

▶ **Theorem 11.** *Any finite partial interleaving $\mathcal{P}$ between any two merge trees T_1 and T_2 admits an optimal $\mathcal{P}$-extension.*

Isolated extensions. To reason about "greatest" shifts and "highest" ancestors, it is useful to have optimal extensions that satisfy two additional finiteness properties. Specifically, assume the $\mathcal{P}$-residual interleaving distance is strictly positive, that is, $d_{\mathcal{P}} > 0$. We say an optimal $\mathcal{P}$-extension $\mathcal{I}$ is *isolated* if (1) the $\mathcal{P}$-shift of only finitely many arrows of $\mathcal{I}$ is equal to $d_{\mathcal{P}}$, and (2) there exists some $\varepsilon > 0$ such that for each arrow (x, y) of $\mathcal{I}$ with $\mathcal{P}$-shift at least $d_{\mathcal{P}} - \varepsilon$, there exists a descendant x' of x, such that (x', y) is an arrow in $\mathcal{I}$ with the same target and $\mathcal{P}$-shift $d_{\mathcal{P}}$.

▶ **Lemma 12.** *If $\mathcal{P}$ is finite and $d_{\mathcal{P}} > 0$, there exists an isolated $\mathcal{P}$-extension.*

4 Locally correct interleavings

The interleaving distance is a bottleneck distance: it yields a single value that quantifies the worst discrepancy between the two merge trees. While this global value is useful for comparison, we are interested in the actual interleavings that realize the distance. Importantly, we seek interleavings that are not only globally optimal, but also locally meaningful. Specifically, if a part of the interleaving is fixed, we want the remainder to be optimal relative to this fixed part. To formalize this notion of local optimality, we build upon the residual interleaving distance introduced in the previous section.

Let (T_1, f_1) and (T_2, f_2) be two merge trees, and fix a partial interleaving $\mathcal{P} = (\varphi, \psi)$. For a subset $S_1 \subseteq \mathrm{dom}(\varphi)$, we define the *restriction* of φ to S_1 as the map $\varphi|_{S_1} : S_1 \to T_2$ given by $\varphi|_{S_1}(x) := \varphi(x)$ for all points x of S_1. Any restriction of φ is a partial up-map, and φ extends all of its restrictions. The *restriction* of ψ to a subset $S_2 \subseteq T_2$ is defined symmetrically. Lastly, the pair $\mathcal{R} = (\varphi|_{S_1}, \psi|_{S_2})$ is a partial interleaving such that $\mathcal{P}$ extends $\mathcal{R}$. We call $\mathcal{R}$ a *restriction* of $\mathcal{P}$. Let $\mathcal{I}$ be a complete interleaving. We make the following observation:

▶ **Observation 13.** *For any restriction $\mathcal{R}$ of $\mathcal{I}$, it holds that $d_{\mathcal{R}}(T_1, T_2) \leq \sigma_{\mathcal{R}}(\mathcal{I})$.*

A complete interleaving $\mathcal{I}$ is *locally correct* if for all restrictions $\mathcal{R}$ of $\mathcal{I}$, the $\mathcal{R}$-residual shift of $\mathcal{I}$ is equal to the $\mathcal{R}$-residual interleaving distance.

▶ **Definition 14.** *A complete interleaving $\mathcal{I} = (\alpha, \beta)$ is* locally correct *if and only if for all $S_1 \subseteq T_1$ and $S_2 \subseteq T_2$, the restriction $\mathcal{R} = (\alpha|_{S_1}, \beta|_{S_2})$ satisfies*

$$d_{\mathcal{R}}(T_1, T_2) = \sigma_{\mathcal{R}}(\mathcal{I}). \tag{3}$$

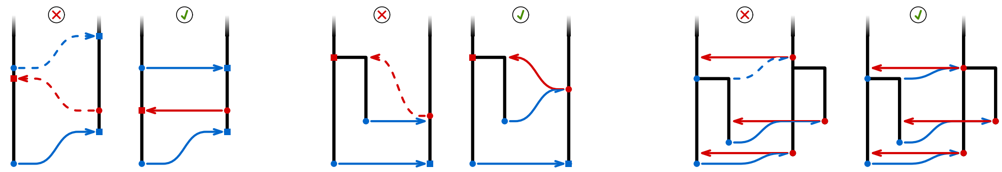

Figure 9 Three pairs of merge trees; for each pair we show an interleaving that is locally correct (✓), and one that is not (⊗). If we fix the non-dashed arrows in the ⊗ cases, the dashed arrows violate Eq. (3); e.g. for the middle example, if we fix the blue arrows, then the image of the red point needs to be an ancestor of both blue points, requiring a shift greater than the interleaving distance.

See Figure 9 for some examples of interleavings that are (not) locally correct.

Suppose $\mathcal{I}$ is locally correct. Taking $\mathcal{R}$ as the empty restriction, we see that $\mathcal{I}$ satisfies $d(T_1, T_2) = \sigma(\mathcal{I})$. In other words, any locally correct interleaving is an optimal interleaving. The reverse is not necessarily true; not every optimal interleaving is locally correct (see Figure 2). That raises the question: does a locally correct interleaving always exist? In the remainder of this section, we answer this question affirmatively.

▶ **Theorem 15.** *Any two merge trees T_1 and T_2 admit a locally correct interleaving.*

Overview of the proof. To prove Theorem 15, we give an explicit construction of a locally correct interleaving. Specifically, we incrementally build a partial interleaving $\mathcal{P}$ by augmenting it with "bottleneck" arrows. We show that with every augmentation, the $\mathcal{P}$-residual interleaving distance strictly decreases, until, after a finite number of iterations, it becomes zero. To conclude the proof, we then show that any complete interleaving that extends the resulting partial interleaving $\mathcal{P}$ is locally correct.

4.1 Bottlenecks

The interleaving distance is determined by the greatest shift among all arrows of an optimal interleaving. Locally, we might improve such an interleaving by "pushing down" some of the arrows, thereby decreasing their shifts. However, we cannot do this for all arrows; otherwise, we would obtain an interleaving whose shift is strictly smaller than the interleaving distance. We are interested in a *bottleneck*: a set of arrows that we cannot "push down".

We first define the *relative difference* of two, not necessarily related, partial up-maps φ and φ'. Specifically, consider the set of points for which either φ is defined and φ' is not, or for which both are defined but where they do not agree:

$$S := \mathrm{dom}(\varphi) \setminus \{x \in \mathrm{dom}(\varphi) \cap \mathrm{dom}(\varphi') : \varphi(x) = \varphi'(x)\}.$$

We define the relative difference of φ and φ' as the restriction of φ to S, denoted $\varphi \setminus \varphi' := \varphi|_S$. Similarly, for two partial interleavings $\mathcal{P} = (\varphi, \psi)$ and $\mathcal{P}' = (\varphi', \psi')$, we define the *relative difference* $\mathcal{P} \setminus \mathcal{P}'$ as the pair of relative differences $(\varphi \setminus \varphi', \psi \setminus \psi')$. The resulting pair is a restriction of $\mathcal{P}$, so it directly follows that any relative difference is a partial interleaving.

Augmentations. Fix a partial interleaving $\mathcal{P} = (\varphi, \psi)$. In the remainder of this section, we assume that $\mathcal{P}$ is finite, and that the $\mathcal{P}$-residual interleaving distance is strictly positive, that is, $d_{\mathcal{P}} > 0$. A $\mathcal{P}$-extension $\mathcal{Q}$ is a $\mathcal{P}$-*augmentation* if (i) the $\mathcal{P}$-residual shift of $\mathcal{Q}$ is at most $d_{\mathcal{P}}$, that is, $\sigma_{\mathcal{P}}(\mathcal{Q}) \leq d_{\mathcal{P}}$, and (ii) the $\mathcal{Q}$-residual interleaving distance is strictly less than $d_{\mathcal{P}}$, that is, $d_{\mathcal{Q}} < d_{\mathcal{P}}$. We say that $\mathcal{Q}$ is *minimal* if it does not extend any other $\mathcal{P}$-augmentation.

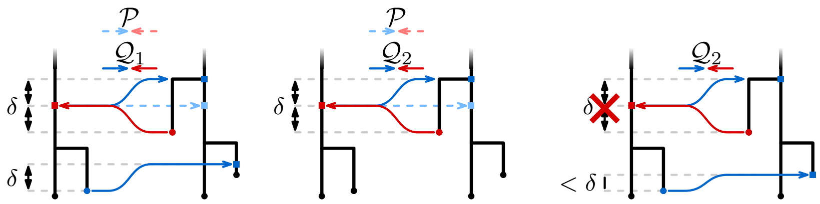

Figure 10 A partial interleaving $\mathcal{P}$ and two augmentations $\mathcal{Q}_1$ and $\mathcal{Q}_2$. $\mathcal{Q}_1$ is not a minimal augmentation, as $\mathcal{Q}_2$ is strictly smaller. The $\mathcal{Q}_2$-residual shift is strictly smaller than δ (right).

▶ **Lemma 16.** *Any finite interleaving $\mathcal{P}$ with $d_{\mathcal{P}} > 0$ admits a finite minimal $\mathcal{P}$-augmentation.*

Let $\mathcal{Q}$ be a minimal $\mathcal{P}$-augmentation. We refer to the relative difference $\mathcal{Q} \setminus \mathcal{P}$ as a $\mathcal{P}$-*bottleneck*. Intuitively, a bottleneck captures a minimal set of arrows that we need to "add" to decrease the residual interleaving distance. See Figure 10 for an illustration.

Recall (the proof of) Lemma 8. If we apply the lemma to a given $\mathcal{P}$-extension $\mathcal{I}$ that does not use a realizing critical pair, we can construct another $\mathcal{P}$-extension $\mathcal{I}'$ that has strictly smaller shift. The exact same proof still works for a slightly stronger statement: it suffices to assume that the relative difference $\mathcal{I} \setminus \mathcal{P}$ does not use a realizing critical pair.

▶ **Observation 17.** *Let $\mathcal{P}$ be a finite interleaving and let $\mathcal{I}$ be a $\mathcal{P}$-extension. If $\mathcal{I} \setminus \mathcal{P}$ does not use a realizing critical pair, there exists a $\mathcal{P}$-extension $\mathcal{I}'$ with $\sigma_{\mathcal{P}}(\mathcal{I}') < \sigma_{\mathcal{P}}(\mathcal{I})$.*

We use this observation to prove the following lemma.

▶ **Lemma 18.** *If $\mathcal{P}$ is finite and $d_{\mathcal{P}} > 0$, any $\mathcal{P}$-bottleneck $\mathcal{Q} \setminus \mathcal{P}$ uses a realizing critical pair.*

Proof. Let $\mathcal{I}$ be an optimal $\mathcal{Q}$-extension. Since $\mathcal{Q}$ extends $\mathcal{P}$, so does $\mathcal{I}$. Moreover, as $\mathcal{I}$ is optimal, the $\mathcal{Q}$-shift of $\mathcal{I}$ is equal to the $\mathcal{Q}$-residual interleaving distance, that is, $\sigma_{\mathcal{Q}}(\mathcal{I}) = d_{\mathcal{Q}}$. Since $\mathcal{Q}$ is an augmentation, the shift of all arrows of $\mathcal{Q} \setminus \mathcal{P}$ is at most $d_{\mathcal{P}}$, and $d_{\mathcal{Q}} < d_{\mathcal{P}}$. Combining, we obtain $\sigma_{\mathcal{P}}(\mathcal{I}) \leq d_{\mathcal{P}}$. In other words, $\mathcal{I}$ is an optimal $\mathcal{P}$-extension. By Observation 17, it follows that $\mathcal{I} \setminus \mathcal{P}$ uses a realizing critical pair b. As $d_{\mathcal{Q}} < d_{\mathcal{P}}$, we get that the $\mathcal{Q}$-shift of the arrows that correspond to b must be 0. In other words, $\mathcal{Q} \setminus \mathcal{P}$ uses b. ◀

A direct consequence of Lemma 18 is that if, for a given $\mathcal{P}$-extension $\mathcal{Q}$, the relative difference $\mathcal{Q} \setminus \mathcal{P}$ does not use a realizing critical pair, then $\mathcal{Q}$ cannot be a $\mathcal{P}$-augmentation. In other words, we have the following corollary:

▶ **Corollary 19.** *Let $\mathcal{P}$ be a finite interleaving with $d_{\mathcal{P}} > 0$, and let $\mathcal{Q}$ be a $\mathcal{P}$-extension. If $\sigma_{\mathcal{P}}(a) < d_{\mathcal{P}}(T_1, T_2)$ for all arrows $a \in A_{\mathcal{Q}}$, then $d_{\mathcal{Q}}(T_1, T_2) \geq d_{\mathcal{P}}(T_1, T_2)$.*

For a partial interleaving $\mathcal{P} = (\varphi, \psi)$, we use $\mathcal{P}.\varphi$ and $\mathcal{P}.\psi$ to denote φ and ψ, respectively.

4.2 Constructing a Locally Correct Interleaving

We show that a locally correct interleaving always exists by incrementally constructing one. Specifically, let $\mathcal{P}$ be an empty interleaving. We iteratively replace $\mathcal{P}$ with a finite minimal $\mathcal{P}$-augmentation, and we show that after a finite number of iterations the $\mathcal{P}$-residual interleaving distance becomes zero, at which point we stop iterating. We conclude by showing that for the resulting partial interleaving $\mathcal{P}$, any optimal $\mathcal{P}$-extension is locally correct. For this, we maintain three invariants throughout the entire construction. First, we argue that $\mathcal{P}$ remains *dominant*: the shift of each arrow of $\mathcal{P}$ is strictly greater than $d_{\mathcal{P}}$. Secondly, we show that for all restrictions of $\mathcal{P}$ it holds that $d_{\mathcal{R}} \geq \sigma_{\mathcal{R}}(\mathcal{P})$; we say $\mathcal{P}$ is partially locally correct. Lastly, we show that after each step of the construction, $\mathcal{P}$ "specifies" at least one additional vertex of the input merge trees. Since each vertex can be specified at most once, it follows that the construction terminates.

Invariants. Formally, let $\mathcal{P}$ be a partial interleaving and assume $d_{\mathcal{P}} > 0$. Recall that another partial interleaving $\mathcal{Q}$ uses $\mathcal{P}$ if they completely agree on the domains of $\mathcal{P}.\varphi$ and $\mathcal{P}.\psi$.

▶ **Lemma 20.** *Let $\mathcal{P}$ be a dominant, finite interleaving with $d_{\mathcal{P}} > 0$. For any finite minimal $\mathcal{P}$-augmentation $\mathcal{Q}$, (i) $\mathcal{Q}$ uses $\mathcal{P}$, and (ii) $\mathcal{Q} \setminus \mathcal{P}$ uses only realizing critical pairs.*

We use this lemma to argue that $\mathcal{P}$ remains locally correct throughout the construction.

▶ **Lemma 21.** *Let $\mathcal{P}$ be a finite interleaving and suppose $d_{\mathcal{P}} > 0$. Let $\mathcal{Q}$ be a finite minimal $\mathcal{P}$-augmentation. If $\mathcal{P}$ is dominant and locally correct, then so is $\mathcal{Q}$.*

Proof. We first argue that $\mathcal{Q}$ is dominant. Let a be an arrow of $\mathcal{Q}$. If a is also an arrow of $\mathcal{P}$, then from the fact that $\mathcal{P}$ is dominant we know that $\sigma(a) > d_{\mathcal{P}} > d_{\mathcal{Q}}$. Otherwise, a must be an arrow of the bottleneck $\mathcal{Q} \setminus \mathcal{P}$. By Lemma 20, it then directly follows that $\sigma(a) = d_{\mathcal{P}} > d_{\mathcal{Q}}$.

Next, we argue that $\mathcal{Q}$ is locally correct. Consider a restriction $\mathcal{R}$ of $\mathcal{Q}$; we need to show that $d_{\mathcal{R}} \geq \sigma_{\mathcal{R}}(\mathcal{Q})$. If $\mathcal{R} = \mathcal{Q}$, we immediately obtain $\sigma_{\mathcal{R}}(\mathcal{Q}) = 0 \leq d_{\mathcal{R}}$. Otherwise, let $\mathcal{R}'$ be the restriction of $\mathcal{R}$ to the domains of $\mathcal{P}.\varphi$ and $\mathcal{P}.\psi$. Since $\mathcal{Q}$ uses $\mathcal{P}$, it follows that $\mathcal{R}'$ uses $\mathcal{P}$ and is hence a restriction of $\mathcal{P}$. If $\mathcal{R}' = \mathcal{P}$, then by definition of an augmentation we know that the $\mathcal{R}'$-shift of any arrow of $\mathcal{Q}$ is $d_{\mathcal{P}}$. We obtain $\sigma_{\mathcal{R}'}(\mathcal{Q}) = d_{\mathcal{P}}$. Moreover, since $\mathcal{R} \neq \mathcal{Q}$, the $\mathcal{R}$-shift of at least one arrow of $\mathcal{Q} \setminus \mathcal{P}$ is non-zero. We obtain $\sigma_{\mathcal{R}}(\mathcal{Q}) = \sigma_{\mathcal{R}'}(\mathcal{Q}) = d_{\mathcal{P}}$. Since $\mathcal{Q}$ is a minimal augmentation and extends $\mathcal{R}$, we know that $\mathcal{R}$ cannot be an augmentation of $\mathcal{P}$. It follows that $d_{\mathcal{R}} \geq d_{\mathcal{P}}$. Putting it together, we obtain $d_{\mathcal{R}} = \sigma_{\mathcal{R}}(\mathcal{Q})$.

Lastly, if $\mathcal{R}' \neq \mathcal{P}$, then since $\mathcal{P}$ is locally correct and $\mathcal{R}'$ is a restriction of $\mathcal{P}$, we have

$$d_{\mathcal{R}'} \geq \sigma_{\mathcal{R}'}(\mathcal{P}). \tag{4}$$

We first show that $\sigma_{\mathcal{R}'}(\mathcal{P}) \geq \sigma_{\mathcal{R}}(\mathcal{Q})$. First, none of the arrows of $\mathcal{R} \setminus \mathcal{R}'$ are arrows of $\mathcal{P}$, so $\sigma_{\mathcal{R}'}(\mathcal{P}) = \sigma_{\mathcal{R}}(\mathcal{P})$. Next, the $\mathcal{R}$-residual shift of $\mathcal{Q}$ is the maximum of the $\mathcal{R}$-residual shifts of $\mathcal{P}$ and $\mathcal{Q} \setminus \mathcal{P}$. Since $\mathcal{P}$ is dominant, we know that the shift of any arrow of $\mathcal{P}$ is strictly greater than $d_{\mathcal{P}}$. As $\mathcal{P} \setminus \mathcal{R}$ is not empty, it follows that $\sigma_{\mathcal{R}}(\mathcal{P}) > d_{\mathcal{P}}$. On the other hand, by definition of an augmentation, we know that the shift of any arrow of $\mathcal{Q} \setminus \mathcal{P}$ is at most $d_{\mathcal{P}}$. So, we get $\sigma_{\mathcal{R}}(\mathcal{Q}) = \max(\sigma_{\mathcal{R}}(\mathcal{P}), \sigma_{\mathcal{R}}(\mathcal{Q} \setminus \mathcal{P})) = \sigma_{\mathcal{R}}(\mathcal{P})$. Combining, we obtain $\sigma_{\mathcal{R}'}(\mathcal{P}) \geq \sigma_{\mathcal{R}}(\mathcal{Q})$.

Next, we argue that $d_{\mathcal{R}} \geq d_{\mathcal{R}'}$. By construction, $\mathcal{R}$ extends $\mathcal{R}'$. As observed before, we have $\sigma_{\mathcal{R}'}(\mathcal{P}) = \sigma_{\mathcal{R}}(\mathcal{P}) > d_{\mathcal{P}}$. Combining with Equation (4), we obtain $d_{\mathcal{R}'} \geq \sigma_{\mathcal{R}'}(\mathcal{P}) > d_{\mathcal{P}}$. Let a be an arrow of $\mathcal{R} \setminus \mathcal{R}'$. Then a is an arrow of the bottleneck $\mathcal{Q} \setminus \mathcal{P}$, so the shift of a is exactly $d_{\mathcal{P}}$. In other words, $\sigma_{\mathcal{R}'}(a) = d_{\mathcal{P}} < d_{\mathcal{R}'}$. This means that we can apply Corollary 19, for the finite interleaving $\mathcal{R}'$ and the $\mathcal{R}'$-extension $\mathcal{R}$, to obtain $d_{\mathcal{R}} \geq d_{\mathcal{R}'}$.

Thus, $d_{\mathcal{R}} \geq d_{\mathcal{R}'} \geq \sigma_{\mathcal{R}'}(\mathcal{P}) = \sigma_{\mathcal{R}}(\mathcal{Q})$, which concludes the proof. ◀

Recall that for each arrow (x, y), with $x \in T_1$ and $y \in T_2$, of $\mathcal{P}$, both x and y are critical points. However, only the point x has a "specified" target in the other tree. We say a point $x \in T_1$ or $y \in T_2$ is *specified* by $\mathcal{P}$, if $x \in \mathrm{dom}(\mathcal{P}.\varphi)$ or $y \in \mathrm{dom}(\mathcal{P}.\psi)$ respectively (see Figure 11). For our last invariant, we argue that $\mathcal{P}$ specifies all critical points that are not vertices, i.e., $C_1[\mathcal{P}] \setminus V(T_1) \subseteq \mathrm{dom}(\mathcal{P}.\varphi)$ and $C_2[\mathcal{P}] \setminus V(T_2) \subseteq \mathrm{dom}(\mathcal{P}.\psi)$.

▶ **Lemma 22.** *Let $\mathcal{P}$ be a finite interleaving and suppose $d_{\mathcal{P}} > 0$. Let $\mathcal{Q}$ be a finite minimal $\mathcal{P}$-augmentation. If $\mathcal{P}$ is dominant and specifies all critical points that are not vertices, then so does $\mathcal{Q}$. Moreover, $\mathcal{Q}$ specifies at least one vertex that $\mathcal{P}$ does not specify.*

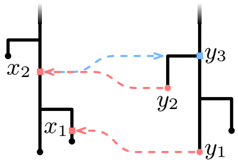

Figure 11 The points x_1, x_2, y_1, y_2 and y_3 are all critical, but only x_2, y_1 and y_2 are specified.

Proof. By Lemma 20, we know that the bottleneck $\mathcal{Q} \setminus \mathcal{P}$ uses only $\mathcal{P}$-critical pairs. Consider such a pair b. If b is an arrow critical pair, then without loss of generality assume $b = (x, y)$ for a point $x \in T_1$ and a point $y \in T_2$; the case $b = (y, x)$ is symmetric. By definition of an arrow critical pair, both x and y are $\mathcal{P}$-critical points. So the pair b does not create any additional $\mathcal{Q}$-critical points. As $\mathcal{P}$ already specifies all critical points that are not vertices and $\mathcal{Q}$ uses $\mathcal{P}$, we know that x must be a vertex. Since $\mathcal{P}$ is dominant, the shift of any arrow of $\mathcal{P}$ is strictly greater than $d_{\mathcal{P}}$. Moreover, the shift of any arrow of $\mathcal{Q} \setminus \mathcal{P}$ is exactly $d_{\mathcal{P}}$. Hence, the vertex x cannot be specified by $\mathcal{P}$: at least one new vertex is specified.

If b is a zigzag critical pair, then without loss of generality assume $b = (x_1, x_2)$ for two points x_1 and x_2 of T_1; if it is a pair of points of T_2 the argument is analogous. By definition of a zigzag critical pair, we know that both $x_1 \in V(T_1)$ and $x_2 \in V(T_1)$. Moreover, there is a point $y \in T_2$ such that $Q.\varphi(x_1) = y$ and $Q.\psi(y) = x_2$. So, x_1 is a vertex that is specified by $\mathcal{Q}$ but not by $\mathcal{P}$. Lastly, b creates a single $\mathcal{Q}$-critical point y, which is immediately specified. ◀

Finishing the proof of Theorem 15. Let $\mathcal{P}$ be the empty interleaving. Trivially, $\mathcal{P}$ is dominant and locally correct. Moreover, there are no (unspecified) critical points that are not vertices. Now, we iteratively replace $\mathcal{P}$ with a finite minimal $\mathcal{P}$-augmentation until the $\mathcal{P}$-residual interleaving distance becomes zero. By applying Lemmas 21 and 22 inductively, we maintain that $\mathcal{P}$ is dominant and locally correct, and specifies all critical points that are not vertices. Additionally, by Lemma 22, each iteration specifies at least one additional vertex. So after i iterations, at least i vertices are specified, and the total number of iterations is at most $|V(T_1)| + |V(T_2)|$. Afterwards, the residual interleaving distance is zero. Lemma 23 guarantees that a locally correct interleaving exists; its proof is similar to that of Lemma 21.

▶ **Lemma 23.** *Let $\mathcal{P}$ be a locally correct partial interleaving with $d_{\mathcal{P}} = 0$. Any optimal $\mathcal{P}$-extension is locally correct.*

Computing a locally correct interleaving. We sketch how the above construction can be translated into an incremental algorithm. To compute the residual interleaving distance, we modify the dynamic program by Touli and Wang [15] to also incorporate constraints induced by partial interleavings. We use the output to extract some augmentation. We then iteratively reduce the set of critical pairs by repeatedly calling the modified dynamic program until we obtain a (finite) minimal augmentation. This minimal augmentation, in turn, reduces the residual interleaving distance; we iterate until it has dropped to zero. To obtain a locally correct interleaving, we select suitable entries from the table filled by the modified dynamic program. A naive implementation of this algorithm takes n iterations, each of which needs to compute the residual interleaving distance $O(n^2)$ times to obtain a minimal augmentation (since an augmentation contains $O(n^2)$ critical pairs). So, the total running time of our algorithm is $O(n^3 T(n))$, where $T(n)$ is the running time of computing the residual interleaving distance. The algorithm by Touli and Wang [15] takes $O(n^4 2^\tau \tau^{\tau+2} \log n)$ time, where τ is a parameter related to the maximum degrees in the input trees.

5 Conclusion

We introduced locally correct interleavings between merge trees, which are optimal interleavings that are "tight". For this, we introduced the residual interleaving distance: a generalized version of the interleaving distance that can accommodate constraints on the interleavings. We presented a constructive proof that there always exists a locally correct interleaving.

We plan to further investigate the complexity of computing a locally correct interleaving, and we will study the effect of partial interleavings on the running time of the modified dynamic program. Furthermore, we will explore the effects of replacing the dynamic program with the approximation algorithm [1] or one of the variants or heuristics. Moreover, it would be interesting to study other (local) criteria of interleavings. Since locally correct interleavings are not necessarily unique, we could minimize their *total* shift, for a suitable definition of total. We could also consider a *lexicographic* optimization, inspired by lexicographic Fréchet matchings [12]. Lastly, we want to explore in which way locally correct interleavings can be translated to meaningful matchings between terrains.

References

1 P.K. Agarwal, K. Fox, A. Nath, A. Sidiropoulos, and Y. Wang. Computing the Gromov-Hausdorff distance for metric trees. *ACM Transactions on Algorithms*, 14(2):1–20, 2018. doi:10.1145/3185466.

2 T. Beurskens, T. Ophelders, B. Speckmann, and K. Verbeek. Relating interleaving and Fréchet distances via ordered merge trees. *Journal of Computational Geometry*, 17(1):1–36, 2026. doi:10.20382/jocg.v17i1a1.

3 K. Buchin, M. Buchin, W. Meulemans, and B. Speckmann. Locally correct Fréchet-matchings. *Computational Geometry*, 76(1):1–18, 2019. doi:10.1016/j.comgeo.2018.09.002.

4 J. Curry, H. Hang, W. Mio, T. Needham, and O. B. Okutan. Decorated merge trees for persistent topology. *Journal of Applied and Computational Topology*, 6(3):371–428, 2022. doi:10.1007/s41468-022-00089-3.

5 E. Gasparovic, E. Munch, S. Oudot, K. Turner, B. Wang, and Y. Wang. Intrinsic interleaving distance for merge trees. arXiv:1908.00063.

6 E. Gasparovic, E. Munch, S. Oudot, K. Turner, B. Wang, and Y. Wang. Intrinsic interleaving distance for merge trees. *La Matematica*, 4(1):40–65, 2025. doi:10.1007/s44007-024-00143-9.

7 W. Köpp and T. Weinkauf. Temporal merge tree maps: A topology-based static visualization for temporal scalar data. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1157–1167, 2022. doi:10.1109/TVCG.2022.3209387.

8 D. Morozov, K. Beketayev, and G. Weber. Interleaving distance between merge trees. Manuscript (accessed on 06-03-2025), 2013. URL: https://mrzv.org/publications/interleaving-distance-merge-trees/manuscript/.

9 E. Munch and A. Stefanou. The ℓ^∞-cophenetic metric for phylogenetic trees as an interleaving distance. In *Research in Data Science*, volume 17 of *Association for Women in Mathematics Series*, pages 109–127. Springer International Publishing, 2019. doi:10.1007/978-3-030-11566-1_5.

10 M. Pegoraro. A graph-matching formulation of the interleaving distance between merge trees. *AIMS Mathematics*, 10(6):13025–13081, 2025. doi:10.3934/math.2025586.

11 M. Pont, J. Vidal, J. Delon, and J. Tierny. Wasserstein distances, geodesics and barycenters of merge trees. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):291–301, 2022. doi:10.1109/TVCG.2021.3114839.

12 G. Rote. Lexicographic Fréchet matchings. In *Proc. 30th European Workshop on Computational Geometry (EuroCG'14)*, 2014. doi:10.17169/refubium-19845.

13 R. Sridharamurthy, T. B. Masood, A. Kamakshidasan, and V. Natarajan. Edit distance between merge trees. *IEEE Transactions on Visualization and Computer Graphics*, 26(3):1518–1531, 2020. doi:10.1109/TVCG.2018.2873612.

14 R. Sridharamurthy and V. Natarajan. Comparative analysis of merge trees using local tree edit distance. *IEEE Transactions on Visualization and Computer Graphics*, 29(2):1518–1530, 2021. doi:10.1109/TVCG.2021.3122176.

15 E.F. Touli and Y. Wang. FPT-algorithms for computing the Gromov-Hausdorff and interleaving distances between trees. *Journal of Computational Geometry*, 13:89–124, 2022. doi:10.20382/jocg.v13i1a4.

16 L. Yan, T. B. Masood, F. Rasheed, I. Hotz, and B. Wang. Geometry aware merge tree comparisons for time-varying data with interleaving distances. *IEEE Transactions on Visualization and Computer Graphics*, 29(8):3489–3506, 2022. doi:10.1109/TVCG.2022.3163349.

17 L. Yan, T. B. Masood, R. Sridharamurthy, F. Rasheed, V. Natarajan, I. Hotz, and B. Wang. Scalar field comparison with topological descriptors: Properties and applications for scientific visualization. *Computer Graphics Forum*, 40(3):599–633, 2021. doi:10.1111/cgf.14331.

18 L. Yan, Y. Wang, E. Munch, E. Gasparovich, and B. Wang. A structural average of labeled merge trees for uncertainty visualisation. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):832–842, 2020. doi:10.1109/TVCG.2019.2934242.