Algorithmica
DOI 10.1007/s00453-009-9319-y

Seriation in the Presence of Errors: A Factor 16 Approximation Algorithm for l_∞-Fitting Robinson Structures to Distances

Victor Chepoi · Morgan Seston

Received: 12 November 2007 / Accepted: 28 April 2009

Abstract The classical seriation problem consists in finding a permutation of the rows and the columns of the distance (or, more generally, dissimilarity) matrix d on a finite set X so that small values should be concentrated around the main diagonal as close as possible, whereas large values should fall as far from it as possible. This goal is best achieved by considering the *Robinson property*: a distance d_R on X is Robinsonian if its matrix can be symmetrically permuted so that its elements do not decrease when moving away from the main diagonal along any row or column. If the distance d fails to satisfy the Robinson property, then we are lead to the problem of finding a reordering of d which is as close as possible to a Robinsonian distance.

In this paper, we present a factor 16 approximation algorithm for the following NP-hard fitting problem: given a finite set X and a dissimilarity d on X, we wish to find a Robinsonian dissimilarity d_R on X minimizing the l_∞-error $\|d - d_R\|_\infty = \max_{x,y \in X}\{|d(x,y) - d_R(x,y)|\}$ between d and d_R.

Keywords Robinsonian dissimilarity · Approximation algorithm · Fitting problem

---

An extended abstract of this paper appeared in the proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009.

V. Chepoi (✉) · M. Seston
Laboratoire d'Informatique Fondamentale de Marseille, Faculté des Sciences de Luminy, Université d'Aix-Marseille, 13288 Marseille Cedex 9, France
e-mail: chepoi@lif.univ-mrs.fr

M. Seston
e-mail: seston@lif.univ-mrs.fr

1 Introduction

1.1 Seriation Problem

A major issue in classification and data analysis is to visualize simple geometrical and relational structures between objects. Necessary for such an analysis is a similarity or a dissimilarity measure on a set of objects, which is either measured directly or computed from a data matrix. Many applied algorithmic problems ranging from archeological dating through DNA sequencing and numerical ecology to sparse matrix reordering and overlapping clustering involve ordering a set of objects so that closely coupled elements are placed near each other. The rearranged data may then speak for themselves. For example, the classical *seriation problem* [39, 44, 47] is to find a simultaneous ordering (or permutation) of the rows and the columns of the similarity matrix with the objective of revealing an underlying one-dimensional structure. The basic idea is that large values should be concentrated around the main diagonal as closely as possible, whereas small values should fall as far from it as possible. This goal is best achieved by considering the so-called *Robinson property* [50]. A similarity matrix is said to have this property if its values decrease monotonically in the rows and the columns when moving away from the main diagonal in both directions. In case of (0, 1)-matrices, the Robinson property is best known as the Consecutive One Property. Seriation is of importance in archeological dating [39, 44, 50], clustering hypertext orderings [11], numerical ecology [48, 58], sparse matrix ordering [5], musicology [37], matrix visualization methods [18, 19], and DNA sequencing [10, 17, 49]. For example, if biological experiments were error-free, the genomic reconstruction problem would be precisely testing the Consecutive One Property. In ecology, seriation techniques are used under the name *ordination* [36]. Matrix visualization methods are employed to analyze multivariate data by means of permutation matrices [12], to visualize large scale genome expression data using heat maps [30], and to visualize clusterings using matrix shading [46, 53, 54]. A package seriation containing an implementation of various seriation methods is described in a recent paper by Hahsler, Hornik, and Buchta [36].

The most common methods for clustering provide a visual display of data in the form of hierarchical structures (or *dendrograms*). Dissimilarity matrices which are in perfect agreement with dendrograms satisfy the following Robinson property: the distances from a given object increase when moving from this object to the left or to the right in the underlying dendrogram. These dissimilarities are best known under the name of *ultrametrics*. Generalizing the correspondence between ultrametrics and dendrograms, it has been shown by Diday [27] and Durand and Fichet [29] that there exists a one-to-one correspondence between the Robinson dissimilarities and the hierarchical structures called *pyramids*. Analogously to dendrograms, the objects belonging to a single cluster in a pyramid appear in consecutive order in the Robinsonian matrix, however in pyramids two clusters may overlap. Clustering methods using pyramids have been considered in [13, 28, 54] (for recent developments and applications of pyramids and overlapping clustering see some articles of the volume [15]).

Real experimental or archeological data always contain errors, therefore the (dis)similarity between the objects can be measured only approximatively. As a consequence, any simultaneous permutation of the rows and the columns of the (dis)similarity matrix gives a matrix which fails to satisfy the Robinson property, and we are lead to the problem of finding a matrix reordering which is as close as possible to a Robinson matrix. As an error measure one can use the usual l_p-distance between two matrices of equal size. Various heuristics for approximate seriation using Robinson matrices have been considered in [16, 39, 40] and papers cited therein (the package `seriation` [36] contains an implementation of several such methods). In particular, Hubert [39] noticed: "Whether or not certain seriation techniques produce 'better' approximation to a Robinson matrix is essentially an open question and requires ... some more or less arbitrary measures of fit when the matrix is not initially Robinson." In this paper, we provide a constant factor polynomial time algorithm for the NP-hard problem of optimal fitting a (dis)similarity matrix by a Robinson matrix under the l_∞-error.

1.2 Basic Definitions, Notations, and Problems

Let X be a set of n elements to sequence, endowed with a dissimilarity function that reflects the desire for two elements to be near or far from each other in the sequence. Recall that a *dissimilarity* is a symmetric function d from X^2 to the nonnegative real numbers and vanishing on the diagonal, i.e. $d(x, y) = d(y, x) \geq 0$ and $d(x, y) = 0$ if and only if $x = y$. We call $d(x, y)$ the *distance* between the objects $x, y \in X$. If d satisfies, in addition, the triangle inequality $d(x, y) \leq d(x, z) + d(z, y)$ for all $x, y, z \in X$, then d is called a *metric*. Denote by d_0 the standard distance of the complete graph on X, i.e., $d_0(x, y) = 1$ if $x \neq y$ and $d_0(x, y) = 0$, otherwise. Denote by $\mathcal{D}$ the set of all dissimilarities on X. Notice that $\mathcal{D}$ is a convex cone, because the convex combination $d := cd' + (1 - c)d''$ of two dissimilarities $d', d'' \in \mathcal{D}$ also belongs to $\mathcal{D}$; for $x, y \in X$, $d(x, y) = cd'(x, y) + (1 - c)d''(x, y)$. Additionally, if $d \in \mathcal{D}$ and $c > 0$ is a constant such that $d - cd_0$ is non-negative, then $d - cd_0 \in \mathcal{D}$.

A dissimilarity d and a total order $\prec$ on a set X are said to be *compatible* if $x \prec z \prec y$ implies that $d(x, y) \geq \max\{d(x, z), d(z, y)\}$. A dissimilarity d on X is said to be *Robinsonian* if it admits a compatible order. Equivalently, d is Robinsonian if its matrix can be symmetrically permuted so that its elements do not decrease when moving away from the main diagonal along any row or column. Such a matrix is called *Robinson* [24, 27, 34, 39] or *linear* in the terminology of Mirkin and Rodin [49]. Notice that if d is Robinsonian, then every monotone transform φ of d is also Robinsonian (i.e., $d(x, y) \leq d(x', y')$ implies that $\varphi(d(x, y)) \leq \varphi(d(x', y'))$). In particular, for any Robinsonian dissimilarity d and any constant c, if $d + cd_0$ is a dissimilarity, then $d + cd_0 \in \mathcal{R}$.

Basic examples of Robinson dissimilarities are the ultrametrics and the standard *line-distance* between n points on the line. Recall, that d is an *ultrametric* if $d(x, y) \leq \max\{d(x, z), d(y, z)\}$ for all $x, y, z \in X$. Every line-distance or ultrametric is a *tree-distance* (i.e., a distance satisfying the four-point condition $d(x, y) + d(z, w) \leq \max\{d(x, z) + d(y, w), d(x, w) + d(y, z)\}$), nevertheless not every Robinson distance is a tree-distance and not every tree-distance is Robinsonian.

Let X be a set on n elements. Denote by $\mathcal{R}$, $\mathcal{R}^{\prec}$, and $\mathcal{U}$ the cone of all Robinson dissimilarities on X, the convex cone of all Robinson dissimilarities compatible with a total order $\prec$ on X, and the convex cone of all ultrametrics on X, respectively. For two dissimilarities $d, d' \in \mathcal{D}$, and $p \geq 1$, define the l_p-error or l_p-norm between d and d' by

$$\|d - d'\|_p = \left(\sum_{x,y \in X} |d(x, y) - d'(x, y)|^p \right)^{\frac{1}{p}},$$

$$\|d - d'\|_\infty = \max_{x,y \in X} \{|d(x, y) - d'(x, y)|\}.$$

Endow the set of all dissimilarities $\mathcal{D}$ with a partial order $\leq$, where $d \leq d'$ if and only if $d(x, y) \leq d'(x, y)$ for all $x, y \in X$. For a dissimilarity d and a subset $\mathcal{D}'$ of $\mathcal{D}$, a dissimilarity $\hat{d} \in \mathcal{D}'$ is called a *sub-dominant* of d in $\mathcal{D}'$ if $\hat{d}$ is the (necessarily unique) maximum of the set $\{d' \in \mathcal{D}' : d' \leq d\}$. Analogously, a dissimilarity $\check{d} \in \mathcal{D}'$ is called a *super-dominant* of d in $\mathcal{D}'$ if $\check{d}$ is the (necessarily unique) minimum of the set $\{d' \in \mathcal{D}' : d \leq d'\}$. The dissimilarities $\hat{d}$ and/or $\check{d}$ do not always exist, nevertheless they exist in some important cases, in particular if $\mathcal{D}'$ is the set of all ultrametrics [52]. In this case, the ultrametric sub-dominant $\hat{d}_u$ can be defined in the following way: construct the minimum spanning tree T in the complete graph on X in which the length of the edge xy is $d(x, y)$, then $\hat{d}_u(x, y)$ is the length of the longest edge on the unique path of T connecting the vertices x and y.

In this paper, we study the following optimization problem, which can be viewed as the seriation problem in the presence of errors:

Problem l_∞-FITTING-BY-ROBINSON: *Given a dissimilarity $d \in \mathcal{D}$ find a Robinson dissimilarity $d_R \in \mathcal{R}$ minimizing the l_∞-error $\|d - d_R\|_\infty$.*

In other words, we are searching for a minimal value of ϵ such that for each pair x, y of different elements of X one can pick a value $d_R(x, y) \in [d(x, y) - \epsilon, d(x, y) + \epsilon]$ so that the resulting dissimilarity d_R is Robinsonian. To formulate the underlying decision problem, we relax the notions of compatible order and Robinson dissimilarity. Given $\epsilon \geq 0$, a total order $\prec$ on X is called *ϵ-compatible* if $u \prec x \prec y \prec v$ implies $d(u, v) + 2\epsilon \geq d(x, y)$; here u and x may coincide as well as y and v. To show that a total order $\prec$ is ϵ-compatible it suffices to show that $d(x, z) \geq \max\{d(x, y), d(y, z)\} + \epsilon$ for any $x, y, z \in X$ such that $x \prec y \prec z$ (note that this condition is not necessary but is all we actually use). An *ϵ-Robinsonian dissimilarity* is a dissimilarity admitting an ϵ-compatible order. We are lead to the following recognition problem:

Problem ϵ-ROBINSON: *Given a dissimilarity d and a real number $\epsilon > 0$, is d ϵ-Robinsonian?*

In the companion paper [22], we show that the problem ϵ-ROBINSON is NP-complete and that, unless P = NP, it is NP-hard to approximate l_∞-FITTING-BY-ROBINSON within a factor smaller than 3/2. On the other hand, the main result of present paper is a factor 16 approximation algorithm for the problem l_∞-FITTING-BY-ROBINSON. (A polynomial time algorithm is called an *α-factor* *approximation* algorithm for a minimization problem Π if for each instance I of Π, it returns a solution whose value is at most α times the optimal value $\mathrm{OPT}_{\Pi}(I)$ of Π on I plus a constant not depending of I; see [57].)

1.3 Related Work

We briefly recall here what is known about Robinsonian structures and about l_∞-fitting of distances by simpler distances.

In the original Robinson's paper [50] and in some other papers [5, 43, 48], compatible orders and Robinson matrices are defined for similarities; the elements of a Robinson similarity matrix not increase when moving away from the main diagonal. In particular, Robinson (0, 1)-similarity matrices are exactly the symmetric matrices having the *Consecutive One Property* [14, 35, 43]. Atkins, Boman and Hendrickson [5] call a similarity matrix pre-Robinson if its lines and columns can be permuted to become a Robinson matrix. Kendall [43] established that if A is a (0, 1)-matrix that can be permuted to obey the Consecutive One Property, then AA^T can be permuted to become a Robinson matrix applying the same sequence of permutations as for A. Atkins et al. [5] showed that if S is a Robinson similarity matrix, then the coordinates of the eigenvector of its smallest nonzero eigenvalue of the Laplacian of S (the so-called Fiedler vector of S) constitute a monotone sequence of numbers. They use this result to design an algorithm of complexity $O(nT(n) + n^2 \log n)$ to recognize if a similarity matrix of size $n \times n$ is pre-Robinson, where $T(n)$ is the complexity of computing the Fiedler vector of a matrix. The algorithm of [5] also uses the PQ-trees of Booth and Lueker [14]. PQ-trees have been employed by Klinz, Rudolf, and Woeginger [45] to recognize in $O(nm(n+m))$ time the permuted bottleneck Monge $n \times m$ matrices, which can be viewed as another generalization of matrices with the Consecutive One Property. Mirkin and Rodin [49] describe an $O(n^4)$ algorithm for testing if a dissimilarity d on n points is Robinsonian. For this, they build up the hypergraph of all balls of d and test using the PQ-tree algorithm if this hypergraph is an interval hypergraph. A simple divide-and-conquer $O(n^3)$-time algorithm for the same recognition problem has been designed in [20]. Barthelemy and Brucker [9] established that the problem of an optimal l_p-approximation of a dissimilarity by a particular Robinsonian dissimilarity (namely, strongly Robinsonian dissimilarity) is NP-hard for $p < \infty$. Answering a question from [9], in [22] we establish the NP-hardness of this problem in the remaining case $p = \infty$.

The problem of fitting distances by simpler distances is a classical problem in data analysis, phylogeny, and, more recently, in computer science. Fitting distances and low-distortion embedding of distances have been also a subject of extensive mathematical studies. The combinatorial and algorithmic results on these problems have been surveyed in [41]; see also [8] for an account of most recent algorithmic developments. We review here only the results about l_∞-fitting of distances (this error measure is also known as the *maximum additive distortion* [7] or the *maximum additive two-sided error* [8]). Farach, Kannan, and Warnow [31] showed that l_∞-fitting of a distance d by an ultrametric can be done in polynomial time. This result has been used by Agarwala et al. [1] to design a factor 3 approximation algorithm for l_∞-fitting of distances by tree-distances, a problem which has been shown to be NP-hard in the same paper [1]. The idea behind the approximation algorithm of [1] is to show that the best l_∞-ultrametric for d can be used to compute an optimal (in the l_∞-sense) tree-distance which preserves the distances to a fixed point p (called *pivot*), and that this tree-distance provides the required factor 3 approximation to the initial fitting problem. A unified and simplified treatment of these results of [1, 31] using sub-dominants was given in [21]. Namely, it is shown in [21] that any dissimilarity d on X admits a sub-dominant $\hat{d}_p$ in the set of tree-distances on X preserving the distances to a pivot p, that $\hat{d}_p$ can be derived from the ultrametric sub-dominant $\hat{d}_u$ of d via the so-called Farris transform or Gromov product, and that $\hat{d}_u + \frac{1}{2}\|d - \hat{d}_u\|_\infty \cdot d_0$ and $\hat{d}_p + \frac{1}{2}\|d - \hat{d}_p\|_\infty \cdot d_0$ are optimal (in the l_∞-sense) ultrametric and tree-distance preserving the distances to p. l_∞-Fitting of a dissimilarity by a line-distance (the so-called Matrix-To-Line problem) has been proven NP-hard by Saxe [51] in 1979. More recently, Håstad, Ivansson, and Lagergren [38] showed that this problem can be approximated within factor of 2, but unless $\mathrm{P} = \mathrm{NP}$ cannot be approximated within a factor smaller than $7/5$ (notice that Agarwala et al. [1] establish a similar non-approximability result for tree-distances with $7/5$ replaced by $9/8$). The algorithm of Håstad et al. [38] assumes known the leftmost point p in the optimal arrangement of the points of X on the line (to find it one can simply run the algorithm with all n possible choices of the pivot p and return the best arrangement). The algorithm places p at the origin of the line and locate every other point x of X at distance $d(p,x)$ to the right of the origin. It can be easily shown that already this arrangement provides a factor 3 algorithm for the Matrix-To-Line problem (in fact this part of the algorithm is similar to the algorithm of [1, 31] because the distances to p remain all unchanged). To improve the factor from 3 to 2, each point of the obtained arrangement is moved by ϵ to the left or to the right, where ϵ is taken from a compact list of admissible values. To decide for a given ϵ how to move the points, a 2-SAT formula Φ_ϵ is defined. Taking each admissible value of ϵ in increasing order, the algorithm solves the corresponding 2-SAT problem Φ_ϵ, until the first value ϵ^* is found for which Φ_{ϵ^*} has a satisfying assignment. The algorithm returns the arrangement resulting from this assignment. Bădoiu [6] extended the results of [38] to l_∞-fitting of distances by rectilinear (l_1-) distances in the 2-dimensional space and proposed a constant-factor algorithm for this problem. The similar l_p-fitting problems with $p < \infty$ are much harder and finding a constant factor approximation algorithm for either of these problems is a challenging task. Notice only that Dhamdhere [26] obtained a factor $O(\log n)$ algorithm for l_1-fitting of a distance by a line distance and Ailon and Charikar [2] developed a factor $O((\log n \log\log n)^{1/p})$ algorithm for l_p-fitting of a distance by a ultrametric or a tree distance.

Motivated by monotone mappings of distances used in the literature on *multidimensional scaling* [25, 32, 33, 55], Alon, Bădoiu, Demaine, Farach-Colton, Hajiaghayi, and Sidiropoulos [3] introduced the notion of ordinal embedding with multiplicative and additive relaxation. Although the authors investigate only the multiplicative relaxation, ordinal embeddings with additive relaxation—only briefly discussed in [3]—turn out to be closely related with the subject of our paper. Given a dissimilarity d on X and a class of target dissimilarities $\mathcal{D}'$ on X, an *ordinal embedding with additive relaxation* $\alpha \geq 0$ is a choice of $d' \in \mathcal{D}'$ such that if $d(x,y) + \alpha < d(u,v)$, then $d'(x,y) < d'(u,v)$. In other words, significantly different distances have their relative order preserved. For $d \in \mathcal{D}$ and $\mathcal{D}' \subset \mathcal{D}$, the ORDINAL ADDITIVE RELAXATION PROBLEM asks for computing the least α^* such that d admits an ordinal embedding with additive relaxation α^* into $\mathcal{D}'$. It can be easily seen that an embedding with l_∞-error ϵ is also an ordinal embedding with additive relaxation at most 2ϵ. On the other hand, Alon et al. [3] provide an example of a 4-point metric space which can be ordinally embedded into the line with additive relaxation α and additive distortion 10, showing thus that already in the simple case of line-distances the additive relaxation α^* cannot be bounded by the optimal l_∞-error ϵ^*. As we will show below, if the target set is the set $\mathcal{R}$ of all Robinsonian dissimilarities on X, then $\alpha^* = 2\epsilon^*$. Since computing ϵ^* is NP-hard [22], computing the minimum additive relaxation α^* is hard as well. On the other hand, the Robinsonian dissimilarity returned by our factor 16 algorithm provides an ordinal embedding of d into $\mathcal{R}$ with relaxation $\leq 16\alpha^*$.

1.4 Our Result and Techniques

The main result of the present paper is a factor 16 approximation algorithm for the problem l_∞-FITTING-BY-ROBINSON. The algorithm runs in $O(n^6 \log n)$ time. The basic setting of our algorithm goes as follows. First we show that the optimal error ϵ^* of l_∞-FITTING-BY-ROBINSON belongs to a well-defined list Δ of size $O(n^4)$ and that for a given total order $\prec$, a Robinsonian dissimilarity compatible with $\prec$ and best fitting d can be found in $O(n^2)$ time. As in some other minmax problems [57], our approximation algorithm tests the entries of Δ, using a parameter ϵ, which is the "guess" of the optimal error. For current $\epsilon \in \Delta$, the algorithm either finds that no ϵ-compatible order exist, in which case the input dissimilarity d is not ϵ-Robinsonian, or it returns a 16ϵ-compatible order. Now, if ϵ is the least value for which the algorithm does not return the negative answer, then $\epsilon^* \geq \epsilon$, and the returned 16ϵ-Robinsonian dissimilarity has l_∞-error at most $16\epsilon^*$, establishing that we have a factor 16 approximation algorithm.

For $\epsilon \in \Delta$, the algorithm constructs a canonical binary relation $\preccurlyeq$ such that every ϵ-compatible total order refines either $\preccurlyeq$ or its dual. If $\preccurlyeq$ is not a partial order, then the algorithm returns the negative answer. If $\preccurlyeq$ is a total order, then we are done. Otherwise, we select a maximal chain $P = (a_1, a_2, \ldots, a_p)$ of the partial order $\preccurlyeq$ and we search to fit each element of $X^\circ := X \setminus P$ between two consecutive elements of P. We say that two consecutive elements $a_i, a_{i+1} \in P$ form a *hole* H_i and that all elements $x \in X^\circ$ assigned between a_i and a_{i+1} are *located* in H_i. The distribution of the elements to holes is performed so that (a) any pair of elements $x, y \in X^\circ$ located in the same hole H_i must "fit" in this hole, i.e., at least one of the orders $a_i \prec x \prec y \prec a_{i+1}$ or $a_i \prec y \prec x \prec a_{i+1}$ must be $c\epsilon$-compatible for some $c \leq 12$. Partitioning X° into sets X_i of elements which will be located in the same hole H_i, $i = 1, \ldots, p-1$, is not obvious. Even if such a partition is available, we cannot directly apply a recursive call to each X_i, because (b) the elements located outside the hole H_i will impose a certain order on the elements of X_i and, since we tolerate some errors, (c) we cannot ensure that X_i is exactly the set of elements which must be located in H_i in some ϵ-compatible total order.

To deal with (a), we give a classification of admissible and pairwise admissible holes for elements of X°. This allows to conclude that each element $x \in X^\circ$ can be located in the leftmost or rightmost admissible hole for x (we call them *bounding holes* of x); see Fig. 1. Both locations are allowable except the case when there exist several elements having the same pair of bounding holes. For $i < j$, let X_{ij} be the set of all elements of X° having H_i and H_{j-1} as bounding holes. To deal with (b) and (c), we define a directed graph $\overrightarrow{\mathcal{L}_{ij}}$ on each set X_{ij}. We draw an arc between two elements x and y of X_{ij}, if either the distance $d(x, y)$ is "small enough" (in which case both arcs between x and y are present) or if x and y must be located in the same hole and for any ϵ-compatible order $\prec$ we will have $a_{i+1} \prec x \prec y$ or $y \prec x \prec a_{j-1}$. The strongly connected components (which we call *cells*) of $\overrightarrow{\mathcal{L}_{ij}}$ have the property that in each ϵ-compatible order all elements of the same component must be located in the same hole. In fact the cells (and not the sets X_i) are the units to which we apply the recursive calls in the algorithm. To decide in which hole H_i or H_{j-1} to locate each cell of $\overrightarrow{\mathcal{L}_{ij}}$ and to define the relative order between the cells assigned to the same hole, we define another directed graph $\overrightarrow{\mathcal{G}_{ij}}$ displaying the "macroscopic" structure of $\overrightarrow{\mathcal{L}_{ij}}$. The vertices of $\overrightarrow{\mathcal{G}_{ij}}$ are the cells of $\overrightarrow{\mathcal{L}_{ij}}$ and an arc between two cells C and C' is defined (again using distance constraints only) so that to ensure that for any ϵ-compatible order $\prec$ either $a_{i+1} \prec u \prec v$ or $v \prec u \prec a_{j-1}$ holds for all $u \in C$ and $v \in C'$. Moreover, if the cells C and C' contain elements which we know that they must be located in different holes in all ϵ-compatible orders, then in $\overrightarrow{\mathcal{G}_{ij}}$ we draw an arc between C and C' and an arc between C' and C, thus creating a 2-cycle. Furthermore, $\overrightarrow{\mathcal{G}_{ij}}$ satisfies the following crucial properties: (i) if some $\overrightarrow{\mathcal{G}_{ij}}$ does not have a partition into two acyclic subgraphs then no ϵ-compatible order exist and (ii) if $\overrightarrow{\mathcal{G}_{ij}}$ has a partition into two acyclic subgraphs $\overrightarrow{\mathcal{G}_{ij}^-}$ and $\overrightarrow{\mathcal{G}_{ij}^+}$, then all cells of $\overrightarrow{\mathcal{G}_{ij}^-}$ will be located in H_i, all cells of $\overrightarrow{\mathcal{G}_{ij}^+}$ will be located in H_{j-1}, and the topological ordering of each of these graphs defines the relative order between the cells. To partition the graphs $\overrightarrow{\mathcal{G}_{ij}}$ into two acyclic subgraphs (this problem in general is NP-complete [42]), we investigate the specific properties of graphs in question, allowing us to define a 2-SAT formula Φ_{ij} which is satisfiable if and only if the required bipartition of $\overrightarrow{\mathcal{G}_{ij}}$ exists. Finally, to locate in each hole H_i the cells coming from different subgraphs $\overrightarrow{\mathcal{G}_{j'i}^+}$, $\overrightarrow{\mathcal{G}_{ij}^-}$, and $\overrightarrow{\mathcal{G}_{ij''}^-}$ with $j' < i < j < j''$, we use the following separation rule: the cells of $\overrightarrow{\mathcal{G}_{j'i}^+}$ are located to the left of the cells of $\overrightarrow{\mathcal{G}_{ij}^-}$ and the cells of $\overrightarrow{\mathcal{G}_{ij}^-}$ are located to the right of the cells of $\overrightarrow{\mathcal{G}_{ij''}^-}$.

1.5 The Structure of the Paper

In Sect. 2, we establish some auxiliary results used to design the approximation algorithm. In Sects. 3 and 4 we develop the principal technical tools used in the algorithm. The approximation algorithm itself is described in Sect. 5 and its performance analysis is given in Sect. 6. In Sect. 7 we establish that the Robinsonian dissimilarity provided by our algorithm also provides an ordinal embedding with additive relaxation at most $16\alpha^*$. For reader convenience, in the final Sect. 8 we give a glossary of main terms used in our paper. Throughout the rest of this paper we consider only Robinson dissimilarities; all our results can be also re-phrased in terms of similarities.

2 Preliminary Results

In this section, we establish some auxiliary results used in the approximation algorithm.

2.1 Optimal $\prec$-Restricted Robinsonian Approximation

$\prec$-RESTRICTED l_∞-FITTING-BY-ROBINSON problem is obtained from the problem l_∞-FITTING-BY-ROBINSON by replacing $\mathcal{R}$ by $\mathcal{R}_\prec$. Following the approach of [21], here we show that this restricted problem can be solved in polynomial time.

Lemma 2.1 *A dissimilarity d on X is Robinsonian if and only if there exists a total order $\prec$ on X, such that $d(x, y) \geq d(u, v)$ holds for any four (not necessarily distinct) elements $u, v, x, y \in X$ such that $x \prec u \prec v \prec y$.*

Proof First suppose that the inequality $d(x, y) \geq d(u, v)$ holds for all $x \prec u \prec v \prec y$. To show that $\prec$ is compatible with d, pick the elements $x, y, z \in X$ such that $x \prec z \prec y$. Applying twice the 4-point inequality, first with $u = z, v = y$ and then with $u = x, v = z$, we deduce that $d(x, y) \geq d(z, y)$ and $d(x, y) \geq d(x, z)$, respectively. Thus $\prec$ is compatible with d, yielding that d is Robinsonian. Conversely, if d is Robinsonian, then there exists a total order $\prec$ compatible with d. If $x \prec u \prec v \prec y$, then we have $d(x, y) \geq d(x, v)$ and $d(x, v) \geq d(u, v)$, therefore $d(x, y) \geq d(u, v)$, and the required 4-point inequality is proven. □

Lemma 2.2 *If $\prec$ is a total order on X satisfying $d(u, w) \geq \max\{d(u, v), d(v, w)\} - \epsilon$ for any three elements $u, v, w \in X$ such that $u \prec v \prec w$, then $\prec$ is ϵ-compatible.*

Proof It suffices to show that $d(u, v) \geq d(x, y) - 2\epsilon$ for any four (not necessarily distinct) elements $u, x, y, v \in X$ such that $u \prec x \prec y \prec v$. This is indeed the case, because $d(u, v) \geq d(u, y) - \epsilon$ and $d(u, y) \geq d(x, y) - \epsilon$. □

Lemma 2.3 *For a total order $\prec$ on X and $d \in \mathcal{D}$, let $\check{d}_\prec$ be a dissimilarity defined in the following way: for $x, y \in X$ with $x \prec y, x \neq y$, set*

$$\check{d}_\prec(x, y) = \max\{d(u, v) : u, v \in X \text{ and } x \prec u \prec v \prec y\}$$

(we suppose here that $a \prec a$ for any $a \in X$). Then $\check{d}_\prec$ is the super-dominant of d in $\mathcal{R}_\prec$.

Proof From the definition of $\check{d}_\prec$ we infer that $d \leq \check{d}_\prec$. We assert that $\check{d}_\prec$ is Robinsonian, namely that $\prec$ is a compatible order for $\check{d}_\prec$. Let $E(u, v) = \{d(u', v') : u \prec u' \prec v' \prec v\}$. For a triplet $x \prec z \prec y$, the distance-sets $E(x, z)$ and $E(z, y)$ are obviously contained in $E(x, y)$. Taking the maximum in each of the sets $E(x, z)$, $E(z, y)$ and $E(x, y)$, we deduce that $x \prec z \prec y$ implies $\check{d}_\prec(x, y) \geq \max\{\check{d}_\prec(x, z), \check{d}_\prec(z, y)\}$. It remains to show that if $d' \in \mathcal{R}_\prec$ and $d \leq d'$, then $\check{d}_\prec \leq d'$. Pick $x, y \in X$ with $x \prec y$. Let u, v be defined so that $x \prec u \prec v \prec y$ and $\check{d}_\prec(x, y) = d(u, v)$. Lemma 2.1

yields $d'(x,y) \geq d'(u,v)$. Since $d \leq d'$, we also have $d(u,v) \leq d'(u,v)$. Putting all together, we obtain $\check{d}_\prec(x,y) = d(u,v) \leq d'(u,v) \leq d'(x,y)$, thus $\check{d}_\prec \leq d'$. This shows that indeed $\check{d}_\prec$ is the super-dominant of d in $\mathcal{R}_\prec$. □

Let $2\tilde{\epsilon}_\prec = \|d - \check{d}_\prec\|_\infty$ and set $\delta = \min\{d(x,y) : x \neq y\} > 0$. Let $\tilde{d}_\prec$ be the (Robinsonian) dissimilarity obtained from $\check{d}_\prec$ by setting $\tilde{d}_\prec(x,y) = \max\{\check{d}_\prec(x,y) - \tilde{\epsilon}_\prec, \delta\} > 0$ for all $x, y \in X$, $x \neq y$, and $\tilde{d}_\prec(x,x) = 0$.

Lemma 2.4 *For a total order $\prec$ on X and $d \in \mathcal{D}$, $\tilde{d}_\prec$ is a Robinsonian dissimilarity that minimizes $\|d - d'\|_\infty$ for $d' \in \mathcal{R}_\prec$.*

Proof By definition, $\tilde{d}_\prec$ is a dissimilarity. From Lemma 2.3 we know that $\prec$ is a compatible order for $\check{d}_\prec$. We assert that $\prec$ is a compatible order also for $\tilde{d}_\prec$. Let $x \prec z \prec y$ and $x \neq z \neq y$. By symmetry, it suffices to show that $\tilde{d}_\prec(x,y) \geq \tilde{d}_\prec(x,z)$. If $\tilde{d}_\prec(x,z) = \delta$, then we are done because $\tilde{d}_\prec(x,y) \geq \delta$. Otherwise, if $\tilde{d}_\prec(x,z) = \check{d}_\prec(x,z) - \tilde{\epsilon}_\prec > \delta$, then the inequality $\check{d}_\prec(x,y) \geq \check{d}_\prec(x,z)$ implies that $\tilde{d}_\prec(x,y) = \check{d}_\prec(x,y) - \tilde{\epsilon}_\prec > \delta$, whence $\tilde{d}_\prec(x,y) = \check{d}_\prec(x,y) - \tilde{\epsilon}_\prec \geq \check{d}_\prec(x,z) - \tilde{\epsilon}_\prec = \tilde{d}_\prec(x,z)$. This shows that $\tilde{d}_\prec \in \mathcal{R}_\prec$.

To prove that $\tilde{d}_\prec$ is an optimal l_∞-approximation for d in $\mathcal{R}_\prec$, first note that the l_∞-distance between d and $\tilde{d}_\prec$ is one-half of the l_∞-distance between d and $\check{d}_\prec$. Indeed, if $\tilde{d}_\prec(x,y) = \check{d}_\prec(x,y) - \tilde{\epsilon}_\prec$, then, since $0 \leq \check{d}_\prec(x,y) - d(x,y) \leq 2\tilde{\epsilon}_\prec$, we obtain that $0 \leq \tilde{d}_\prec(x,y) + \tilde{\epsilon}_\prec - d(x,y) \leq 2\tilde{\epsilon}_\prec$, yielding $|\tilde{d}_\prec(x,y) - d(x,y)| \leq \tilde{\epsilon}_\prec$ in this case. If $\tilde{d}_\prec(x,y) = \delta$, then $\check{d}_\prec(x,y) - \tilde{\epsilon}_\prec \leq \delta$, therefore $\delta + \tilde{\epsilon}_\prec \geq \check{d}_\prec(x,y) \geq d(x,y) \geq \delta$, whence $|d(x,y) - \tilde{d}_\prec(x,y)| \leq \tilde{\epsilon}_\prec$. On the other hand, there exist two distinct elements u, v of X such that $\check{d}_\prec(u,v) - d(u,v) = 2\tilde{\epsilon}_\prec$. For this pair, necessarily we have $\tilde{d}_\prec(u,v) = \check{d}_\prec(u,v) - \tilde{\epsilon}_\prec$, thus indeed $\|d - \tilde{d}_\prec\|_\infty = \tilde{\epsilon}_\prec$. Now, suppose by way of contradiction that there exists $d' \in \mathcal{R}_\prec$ such that $\epsilon' = \|d - d'\|_\infty < \tilde{\epsilon}_\prec$. Let u, v be a pair of X for which $\check{d}_\prec(u,v) - d(u,v) = \tilde{\epsilon}_\prec$. Then $\check{d}_\prec(u,v) = \tilde{d}_\prec(u,v) + \tilde{\epsilon}_\prec$. Since $d \leq d' + \epsilon' d_0$ and $d' + \epsilon' d_0 \in \mathcal{R}_\prec$, Lemma 2.3 implies that $d'(u,v) + \epsilon' \geq \check{d}_\prec(u,v) = \tilde{d}_\prec(u,v) + \tilde{\epsilon}_\prec$, whence $d'(u,v) > \tilde{d}_\prec(u,v)$. Since $\tilde{d}_\prec(u,v) = d(u,v) + \tilde{\epsilon}_\prec$, we conclude that $d'(u,v) > \tilde{d}_\prec(u,v) = d(u,v) + \tilde{\epsilon}_\prec > d(u,v) + \epsilon'$, contrary to the definition of $\epsilon' = \|d - d'\|_\infty$.

□

2.2 Compact List of Breakpoints

We prove here that the optimal error ϵ^* in the problem l_∞-FITTING-BY-ROBINSON belongs to a compact list Δ of size $O(n^4)$, whose entries can be derived from the matrix of d.

Lemma 2.5 *For a dissimilarity $d \in \mathcal{D}$, the optimal error ϵ^* of l_∞-fitting of d by a Robinsonian dissimilarity belongs to the set $\Delta = \{|d(x,y) - d(x',y')|/2 : x, y, x', y' \in X\}$.*

Proof Let d^* be an optimal Robinsonian dissimilarity for d and let $\prec$ be a total order compatible with d^*. From Lemma 2.4 we obtain

$$\epsilon^* = \|d - d^*\|_\infty = \|d - \tilde{d}_\prec\|_\infty = \frac{1}{2}\|d - \check{d}_\prec\|_\infty = \frac{1}{2}\max_{x,y\in X} |d(x,y) - \check{d}_\prec(x,y)|.$$

By Lemma 2.3, $\check{d}_\prec(x,y) = d(u,v)$ for some $u, v \in X$ such that $x \prec u \prec v \prec y$, whence ϵ^* has the form $\frac{1}{2}|d(x,y) - d(u,v)|$. This proves that $\epsilon^* \in \Delta$. □

2.3 A False Start

Lemma 2.4 shows that an optimal solution of the problem l_∞-FITTING-BY-ROBINSON can be selected among $n!$ Robinsonian dissimilarities of the form $\tilde{d}_\prec$. We show here that the natural heuristic similar to the approximation algorithms of Håstad et al. [38] and Agarwala et al. [1] (which instead of $n!$ total orders considers only n such orders) does not provide a constant-factor approximation algorithm for our problem. Let $X = \{x_1, \ldots, x_n\}$ and, for simplicity, let d be a dissimilarity on X taking pairwise distinct nonzero values. For each $i = 1, \ldots, n$, the algorithm computes a total order $\prec_i$ by considering x_i as the leftmost point (i.e., $x_i \prec_i x_j$ for all $x_j \neq x_i$) and ordering the remaining points according to their distances to x_i, i.e., by setting $x_j \prec_i x_k$ if and only if $d(x_i, x_j) < d(x_i, x_k)$. For each total order $\prec_i$, the algorithm computes the Robinsonian dissimilarity $\tilde{d}_{\prec_i}$. Finally, from the list of n Robinsonian dissimilarities $\tilde{d}_{\prec_1}, \ldots, \tilde{d}_{\prec_n}$, the algorithm returns the dissimilarity $\tilde{d}_{\prec_i}$ minimizing the l_∞-error $\tilde{\epsilon}_{\prec_i}$ to d. The dissimilarity (see Table 1) on 6 points shows that this heuristic is not a factor c approximation algorithm for any constant $c > 1$.

For total order $\prec$ such that $x_1 \prec x_2 \prec x_3 \prec x_4 \prec x_5 \prec x_6$, we have $\hat{\epsilon}_\prec = 1$, therefore $\epsilon^* \leq 1$. On the other hand, the heuristic selects the following 6 total orders:

$$\prec_1 = (x_1, x_2, x_3, x_5, x_4, x_6) \quad \text{with } \tilde{\epsilon}_{\prec_1} = \left|\frac{d(x_3,x_4) - d(x_3,x_5)}{2}\right| = 1/2 + c;$$

$$\prec_2 = (x_2, x_3, x_4, x_5, x_1, x_6) \quad \text{with } \tilde{\epsilon}_{\prec_2} = \left|\frac{d(x_4,x_6) - d(x_1,x_6)}{2}\right| = 1 + c;$$

$$\prec_3 = (x_3, x_2, x_4, x_5, x_1, x_6) \quad \text{with } \tilde{\epsilon}_{\prec_3} = \left|\frac{d(x_4,x_6) - d(x_1,x_6)}{2}\right| = 1 + c;$$

$$\prec_4 = (x_4, x_5, x_3, x_2, x_1, x_6) \quad \text{with } \tilde{\epsilon}_{\prec_4} = \left|\frac{d(x_4,x_6) - d(x_1,x_6)}{2}\right| = 1 + c;$$

$$\prec_5 = (x_5, x_4, x_3, x_2, x_1, x_6) \quad \text{with } \tilde{\epsilon}_{\prec_5} = \left|\frac{d(x_4,x_6) - d(x_1,x_6)}{2}\right| = 1 + c;$$

$$\prec_6 = (x_6, x_5, x_3, x_4, x_2, x_1) \quad \text{with } \tilde{\epsilon}_{\prec_6} = \left|\frac{d(x_4,x_5) - d(x_3,x_5)}{2}\right| = 3/2 + c.$$

As a result, the heuristic returns the total order $\prec_1$ and the Robinsonian dissimilarity $\tilde{d}_{\prec_1}$ for which the l_∞-error is $\tilde{\epsilon}_{\prec_1} = 1/2 + c > c\epsilon^*$. Similar examples can be constructed to show that several other natural heuristics fail as well.

Table 1 The dissimilarity d on 6 points from Sect. 2.3

	x_1	x_2	x_3	x_4	x_5	x_6
x_1	0	$6+2c$	$7+2c$	$10+2c$	$8+2c$	$15+4c$
x_2		0	2	$3+2c$	$5+2c$	$14+2c$
x_3			0	3	$4+2c$	$12+2c$
x_4				0	1	$13+2c$
x_5					0	$11+2c$
x_6						0

2.4 Canonical Betweenness and Canonical Partial Order

Given a value $\epsilon \in \Delta$, we define two invariants of all (potential) ϵ-compatible orders: the canonical betweenness relation $\mathcal{B}$ and the canonical partial order $\preccurlyeq$ such that every ϵ-compatible total order satisfies $\mathcal{B}$ and refines either $\preccurlyeq$ or its dual.

For a total order $\prec$ on X and two elements $x, y \in X$, the *segment* $[x, y]$ consists of x, y and all $z \in X$ which are *between* x and y, i.e., $[x, y] = \{z \in X : x \prec z \prec y\}$ if $x \prec y$ and $[x, y] = \{z \in X : y \prec z \prec x\}$ if $y \prec x$. Let $\mathcal{B}$ be the set of all triplets (x, z, y) such that $z \in [x, y]$ holds for all ϵ-compatible with d total orders on X. We will show how to derive $\mathcal{B}$ from the distance information. (Note that the problem of deciding if a set of ordered triplets come from a common total order is NP-complete and a polynomial algorithm returning either the negative answer or a total order satisfying at least one half of triplets has been designed in [23].)

Lemma 2.6 *If $d(x, y) > \max\{d(x, z), d(z, y)\} + 2\epsilon$, then $z \in [x, y]$ in any ϵ-compatible with d total order $\prec$.*

Proof Suppose by way of contradiction that $x \prec y \prec z$ for the ϵ-compatible total order $\prec$. Then $d(x, z) + 2\epsilon \geq d(x, y) > d(x, z) + 2\epsilon$, a contradiction. □

Therefore we can insert in $\mathcal{B}$ all triplets (x, x, y), (x, y, y) for $x, y \in X$ and all triplets of the form (x, z, y) for all $x, y, z \in X$ such that $d(x, y) > \max\{d(x, z), d(z, y)\} + 2\epsilon$. Then we close $\mathcal{B}$ by adding all triplets of the form (x, u, y) which follow from the identities satisfied by a betweenness relation in a total order [56]:

(B1) If $z \in [x, y]$ and $u \in [x, z] \cup [y, z]$, then $u \in [x, y]$;
(B2) If $x', y' \in [x, y]$ and $u \in [x', y']$, then $u \in [x, y]$;
(B3) If $z \in [x, y]$, $t \in [z, y]$, and $u \in [t, z]$, then $u \in [x, y]$;
(B4) If $z \in [t, y]$, $t \in [x, z]$, and $u \in [t, z]$, then $u \in [x, y]$.

If we apply the algorithm of Chor and Sudan [23] to $\mathcal{B}$ and it returns a negative answer, then d does not admit any ϵ-compatible total order. Otherwise, however, even if a total order is compatible with all triplets of $\mathcal{B}$, this order is not necessarily ϵ-compatible with d. Therefore we need to refine the constraints imposed by $\mathcal{B}$ by introducing the partial order $\preccurlyeq$ such that every ϵ-compatible total order $\prec$ refines either $\preccurlyeq$ or the dual of $\preccurlyeq$ (a total order $\prec$ *refines* a partial order $\preccurlyeq$ if $x \preccurlyeq y$ implies $x \prec y$). In order to give an orientation to $\mathcal{B}$, we set $p \preccurlyeq q$ for two arbitrary elements $p, q \in X$, and close $\preccurlyeq$ using the following rules:

(PO1) If $x \preccurlyeq y$ and $z \in [x, y]$, then $x \preccurlyeq z \preccurlyeq y$;
(PO2) If $x \preccurlyeq z$ and $z \in [x, y]$, then $x \preccurlyeq y$ and $z \preccurlyeq y$;
(PO3) If $x \preccurlyeq y$ and $d(x, z) > d(y, z) + 2\epsilon$, then $x \preccurlyeq z$;
(PO4) If $x \preccurlyeq y$ and $d(y, z) > d(x, z) + 2\epsilon$, then $z \preccurlyeq y$;
(PO5) If $x \preccurlyeq y$, $x \preccurlyeq z$, and $d(x, z) > d(x, y) + 2\epsilon$, then $x \preccurlyeq y \preccurlyeq z$;
(PO6) If $x \preccurlyeq z$, $y \preccurlyeq z$, and $d(x, z) > d(y, z) + 2\epsilon$, then $x \preccurlyeq y \preccurlyeq z$.

Every ϵ-compatible total order $\prec$ with $p \prec q$ satisfies the conditions (PO1)–(PO6), therefore any such order refines $\preccurlyeq$. In particular, if the resulting binary relation $\preccurlyeq$ is not a partial order, then d does not admit an ϵ-compatible total order. So, further we will assume that $\preccurlyeq$ is a partial order. Notice that for the correctness of our algorithm we only use the rules (PO3)–(PO6), the first two properties (PO1), (PO2) of total orders can be used to derive a richer partial order $\preccurlyeq$. For two disjoint subsets A, B of X, we use the notation $A \preccurlyeq B$ if $a \preccurlyeq b$ for any $a \in A$ and $b \in B$. We write $x?y$ if neither $x \preccurlyeq y$ nor $y \preccurlyeq x$ hold. Let $G = (X, A)$ be the directed graph in which $(x, y) \in A$ if and only if $x \preccurlyeq y$ and $x \preccurlyeq z \preccurlyeq y$ implies that either $z = x$ or $z = y$. Since $\preccurlyeq$ is antisymmetric, the graph G is acyclic. We continue with basic properties of the partial order $\preccurlyeq$ and graph G.

For two numbers α and β we will use the following notations (i) $\alpha \approx_c \beta$ if $|\alpha - \beta| \le c\epsilon$, (ii) $\beta \gtrsim_c \alpha$ if $\beta \ge \alpha - c\epsilon$, and (iii) $\beta \gg_c \alpha$ if $\beta > \alpha + c\epsilon$. The following simple properties will be used throughout the paper without specific references:

(1) If $\alpha \circledast_c \beta$ and $\beta \circledast_{c'} \gamma$, then $\alpha \circledast_{(c+c')} \gamma$ for any $\circledast \in \{\approx, \lesssim, \ll\}$;
(2) If $\alpha \lesssim_c \beta$ and $\beta \ll_{c'} \gamma$, then $\alpha \ll_{(c'-c)} \gamma$;
(3) If $\alpha \ll_c \beta$ and $c' < c$, then $\alpha \ll_{c'} \beta$;
(4) If $\alpha \lesssim_{c'} \beta$ and $\beta \lesssim_{c''} \alpha$, then $\alpha \approx_c \beta$, where $c = \max\{c', c''\}$.

Lemma 2.7 *If $x, y, z \in X$ satisfy $\{x, y\} \preccurlyeq z$ or $z \preccurlyeq \{x, y\}$, and $x?y$, then we have*:

(1) $d(x, z) \approx_2 d(y, z)$;
(2) $d(x, y) \lesssim_2 \min\{d(x, z), d(y, z)\}$.

Proof Since $x \preccurlyeq z$ and $y \preccurlyeq z$, if $|d(x, z) - d(y, z)| > 2\epsilon$ then condition (PO6) implies that either $x \preccurlyeq y$ or $y \preccurlyeq x$ holds, contrary to the assumption that $x?y$. This establishes (1). The proof of (2) is similar, using only (PO3) and (PO4) instead of (PO6). □

Lemma 2.8 *If $u, v, w, z \in X$ satisfy $\{w, u\} \preccurlyeq \{v, z\}$ and $u?w$, $v?z$, then we have*:

(1) $d(w, z) \approx_4 d(u, v)$;
(2) $\max\{d(u, w), d(v, z)\} \lesssim_2 \min\{d(w, z), d(u, z), d(v, w), d(u, z)\}$.

Proof Since $|d(w, z) - d(u, v)| \le |d(w, z) - d(w, v)| + |d(w, v) - d(u, v)|$, while $d(w, z) \approx_2 d(w, v)$ and $d(w, v) \approx_2 d(u, v)$ hold by Lemma 2.7(1), we conclude that $d(w, z) \approx_4 d(u, v)$. The property (2) is a direct consequence of Lemma 2.7(2). □

Lemma 2.9 *If $u, v, w, z \in X$ satisfy $w \preccurlyeq \{v, z\}$, $u \preccurlyeq v$, $u?z$, $u?w$, and $v?z$, then we have*:

(1) $d(w, z) \approx_4 d(u, v)$, $d(u, z) \approx_4 d(w, z)$;

(2) $d(v,z) \lesssim_2 \min\{d(w,z), d(w,v)\}$;
(3) $d(w,u) \lesssim_2 \min\{d(w,v), d(u,v)\}$.

Proof The properties (2) and (3) are immediate consequences of Lemma 2.7(2). Applying Lemma 2.7(1), we obtain $d(w,z) \approx_2 d(w,v)$ and $d(u,v) \approx_2 d(w,v)$, therefore $d(w,z) \approx_4 d(u,v)$. □

3 Admissible and Pairwise Admissible Holes

For a given $\epsilon \in \Delta$, let $\preccurlyeq$ be the canonical partial order on X and let G be the acyclic graph of this order. Denote by $P = (a_1, a_2, \ldots, a_{p-1}, a_p)$ any maximal chain of G. We say that two consecutive elements $a_i, a_{i+1} \in P$ form a *hole* H_i and that all elements $x \in X^\circ$ assigned between a_i and a_{i+1} are *located* in the hole H_i. In this section, we introduce and investigate admissible and pairwise admissible holes for elements of X°. For $x \in X^\circ$, these are the holes H_i into which x can be located such that any other element $y \in X^\circ$ can be located into an admissible hole H_j such that the resulting total order on $P \cup \{x, y\}$ is $c\epsilon$-compatible for some constant $c \leq 12$. We establish that if we tolerate the error 12ϵ (instead of ϵ), then any $x \in X^\circ$ can be located in its leftmost or rightmost admissible holes and that both locations are allowable except the case when there exist several elements having the same pair of bounding holes.

3.1 Admissible Holes

For notational convenience, we consider two additional elements $\curlyvee$ and $\curlywedge$ so that $\curlyvee \preccurlyeq x$ and $x \preccurlyeq \curlywedge$ for all elements of X. For this, we simply set $d(\curlyvee, x) = d(x, \curlywedge) := 4\max\{d(u,v) : u, v \in X\}$. Thus, the path P has the form $(a_1 := \curlyvee, a_2, \ldots, a_{p-1}, \curlywedge := a_p)$, and this way, every element of X° must be located in a hole.

For an element $x \in X^\circ$, define $\underline{a}_x = \max\{a_i \in P : a_i \preccurlyeq x\}$ and $\overline{a}_x = \min\{a_i \in P : x \preccurlyeq a_i\}$. Let $\underline{b}_x$ be the next to $\underline{a}_x$ vertex of P and let $\overline{b}_x$ be the previous to $\overline{a}_x$ vertex of the chain P. Let $I(x) = \{a_i \in P : \underline{a}_x \preccurlyeq a_i \preccurlyeq \overline{a}_x\}$ and $I^\circ(x) = I(x) \setminus \{\underline{a}_x, \overline{a}_x\}$. Denote by $H(x)$ the union of all holes $H_i = [a_i, a_{i+1}]$ comprised between $\underline{a}_x$ and $\overline{a}_x$ (i.e., such that $\underline{a}_x \preccurlyeq a_i \preccurlyeq a_{i+1} \preccurlyeq \overline{a}_x$) and call $H(x)$ the *segment* of x. The holes $\underline{H}_x = [\underline{a}_x, \underline{b}_x]$ and $\overline{H}_x = [\overline{b}_x, \overline{a}_x]$ are called the *bounding holes* of $H(x)$, all other holes of $H(x)$ are called *inner holes*. Since $x \notin P$, the segment $H(x)$ contains at least two holes, in particular the bounding holes of $H(x)$ are different. The hole H_i of $H(x)$ is called *x-admissible,* if the total order on $P \cup \{x\}$ obtained from $\preccurlyeq$ by adding the relation $a_i \preccurlyeq x \preccurlyeq a_{i+1}$ is ϵ-compatible with d. Notice that the bounding holes of $H(x)$ must be x-admissible. Otherwise, if say $[\underline{a}_x, \underline{b}_x]$ is not x-admissible, then we must have $\underline{b}_x \prec x$ in all ϵ-compatible orders $\prec$ extending $\preccurlyeq$. This yields $\underline{b}_x \preccurlyeq x$, and we are lead to a contradiction with the definition of $\underline{a}_x$.

Lemma 3.1 *If $a_i, a_j \in I^\circ(x)$, then $d(x, a_i) \approx_2 d(x, a_j)$.*

Proof Let $a_i \preccurlyeq a_j$. Since we can locate x in the bounding holes of $H(x)$, the orders $x \prec a_i \prec a_j$ and $a_i \prec a_j \prec x$ are ϵ-compatible. Therefore $d(x,a_i) \lesssim_2 d(x,a_j)$ and $d(x,a_j) \lesssim_2 d(x,a_i)$, yielding $d(x,a_i) \approx_2 d(x,a_j)$. ☐

Denote by d_x the mean value of $\min\{d(x,a_i) : a_i \in I^\circ(x)\}$ and $\max\{d(x,a_i) : a_i \in I^\circ(x)\}$. From Lemma 3.1 we infer that $d_x \approx_1 d(x,a_i)$ for all $a_i \in I^\circ(x)$. For a hole $H_i = [a_i, a_{i+1}]$, we call $\delta_i = d(a_i, a_{i+1})$ the *size* of H_i.

Lemma 3.2 *If an inner hole H_i of $H(x)$ is x-admissible, then $\delta_i \approx_2 d(x,a_i)$, $d(x,a_{i+1})$. In particular, $\delta_i \approx_3 d_x$.*

Proof Since the total order $a_i \prec x \prec a_{i+1}$ is ϵ-compatible, we conclude that $d(a_i,a_{i+1}) \gtrsim_2 d(a_i,x)$ and $d(a_i,a_{i+1}) \gtrsim_2 d(x,a_{i+1})$. On the other hand, if say $d(a_i,a_{i+1}) \gg_2 d(a_i,x)$, then the order $x \prec a_i \prec a_{i+1}$ is not ϵ-compatible, contrary to the fact that the leftmost bounding hole of $H(x)$ is x-admissible. Thus indeed $\delta_i \approx_2 d(x,a_i)$ and $\delta_i \approx_2 d(x,a_{i+1})$. ☐

From Lemma 3.2 we conclude that any inner x-admissible hole has size 3-approximatively equal to d_x. This fact can be generalized in the following way:

Lemma 3.3 *If $a_i, a_j \in I^\circ(x)$, then $d_x \gtrsim_3 d(a_i,a_j)$.*

Proof Let $a_i \preccurlyeq a_j$. Since the orders $x \prec a_i \prec a_j$ and $a_i \prec a_j \prec x$ are ϵ-compatible, we deduce that $d(a_i,a_j) \lesssim_2 d(x,a_j) \approx_1 d_x$ and $d(a_i,a_j) \lesssim_2 d(x,a_i) \approx_1 d_y$. ☐

It can be easily shown that the size of a bounding hole is approximatively equal to the distance from x to the respective endpoint and maybe only larger than d_x and $d(\underline{b}_x, \overline{b}_x)$:

Lemma 3.4 *For bounding holes of $H(x)$ we have*

$$d(\underline{a}_x, x) \approx_2 d(\underline{a}_x, \underline{b}_x) \gtrsim_2 d(x, \underline{b}_x) \gtrsim_1 d_x \gtrsim_3 d(\underline{b}_x \overline{b}_x),$$

$$d(\overline{a}_x, x) \approx_2 d(\overline{a}_x, \overline{b}_x) \gtrsim_2 d(x, \overline{b}_x) \gtrsim_1 d_x \gtrsim_3 d(\underline{b}_x \overline{b}_x).$$

Lemma 3.5 *If $a_i, a_j \in I^\circ(x)$, $a_k \in P \setminus I^\circ(x)$, then $d(a_k,x) \approx_2 d(a_k,a_i), d(a_k,a_j)$.*

Proof Suppose without loss of generality that $a_k \preccurlyeq a_i \preccurlyeq a_j$. Then $a_k \preccurlyeq x$. Since $x?a_i$ and $x?a_j$, the orders $a_k \prec x \prec a_i \prec a_j$ and $a_k \prec a_i \prec a_j \prec x$ are ϵ-compatible, yielding the required relationships. ☐

3.2 Pairwise Admissible Holes

A pair $\{H_i, H_j\}$ of holes is called (x,y,c)*-admissible* if H_i is x-admissible, H_j is y-admissible, and the total order on $P \cup \{x,y\}$ obtained from $\preccurlyeq$ by adding the relations $a_i \preccurlyeq x \preccurlyeq a_{i+1}$ and $a_j \preccurlyeq y \preccurlyeq a_{j+1}$ is $c\epsilon$-compatible. Denote by $AH(x)$ the set of all x-admissible holes H_i so that for each element $y \in X^\circ$ different from x there exists an y-admissible hole H_j such that $\{H_i, H_j\}$ is a $(x, y, 1)$-admissible pair. Further we can assume that for any element $x \in X^\circ$ the bounding holes $[\underline{a}_x, \underline{b}_x]$ and $[\overline{b}_x, \overline{a}_x]$ of $H(x)$ belong to $AH(x)$. Otherwise, if say $[\underline{a}_x, \underline{b}_x] \notin AH(x)$, then $\underline{b}_x \prec x$ in any ϵ-compatible total order $\prec$ extending $\preccurlyeq$, therefore we can augment the canonical partial order $\preccurlyeq$ by setting $\underline{b}_x \preccurlyeq x$ and by reducing the sets $H(x)$ and $I(x)$ accordingly.

For two elements $x, y \in X^\circ$, we distinguish the following cases of the mutual geometric location of the segments $H(x)$ and $H(y)$:

(i) $H(x)$ and $H(y)$ are disjoint ($H(x) \cap H(y) = \emptyset$);
(ii) $H(x)$ and $H(y)$ overlap and $|H(x) \cap H(y)| \geq 2$ (notation $H(x) \circ H(y)$);
(iii) $H(x)$ and $H(y)$ overlap and $|H(x) \cap H(y)| = 1$ (notation $H(x) * H(y)$);
(iv) $H(y)$ is a proper subinterval of $H(x)$ (notation $H(y) \Subset H(x)$);
(v) $H(x) = H(y)$.

The rest of this subsection is devoted to the proof of the following important result:

Proposition 3.6 *For two elements $x, y \in X^\circ$, any location of x in a bounding hole of $H(x)$ and any location of y in a bounding hole of $H(y)$ is $(x, y, 12)$-admissible, unless $H(x) = H(y)$ and $d(x, y) \ll_3 \max\{d_x, d_y\}$ or $d(x, y) \gg_3 \max\{d_x, d_y\}$, subject to the following three constraints*:

(i) *if $H(x) \Subset H(y)$, x and y are located in a common bounding hole, then x is between y and $I^\circ(y)$*;
(ii) *if $H(x) * H(y)$, then $\underline{a}_x \preccurlyeq \underline{a}_y$ implies $x \prec y$*;
(iii) *if $H(x) = H(y)$, x and y are located in the same bounding hole, and $d_y \ll_4 d_x$, then y is between x and $I^\circ(y) = I^\circ(x)$.*

If $H(x) = H(y)$ and $d(x, y) \gg_3 \max\{d_x, d_y\}$, then the only $(x, y, 1)$-admissible locations are the two locations of x and y in different bounding holes. If $H(x) = H(y)$ and $d(x, y) \ll_3 \max\{d_x, d_y\}$, then any $(x, y, 1)$-admissible location is in common x- and y-admissible holes.

The proof of Proposition 3.6 is split into several lemmata.

Lemma 3.7 *If H_i is x-admissible, H_j is y-admissible, $i < j$, and $\max\{d(x, a_j), d(y, a_{i+1})\} \lesssim_c d(x, y) \lesssim_c \min\{d(y, a_i), d(x, a_{j+1})\}$ for $c \geq 2$, then the location of x in H_i and of y in H_j is $(x, y, c + 2)$-admissible.*

Proof Suppose x is located in H_i to the left of y. According to Lemma 2.2, it suffices to show that for any $a_k \in P$ we have $d(a_k, y) \gtrsim_{(c+2)} \max\{d(a_k, x), d(x, y)\}$ if $k \leq i$, $d(a_k, x) \gtrsim_{(c+2)} \max\{d(a_k, y), d(x, y)\}$ if $j \leq k$, and $d(x, y) \gtrsim_{(c+2)} \max\{d(x, a_k), d(a_k, y)\}$ if $i \neq j$ and $i < k < j$. Let $k \leq i$. Since H_j is y-admissible and H_i is x-admissible, we deduce that $d(a_k, y) \gtrsim_2 d(a_k, a_j) \gtrsim_2 d(a_k, x)$. On the other hand, $d(a_k, y) \gtrsim_2 d(y, a_i) \gtrsim_c d(x, y)$. Thus $d(a_k, y) \gtrsim_{(c+2)} \max\{d(a_k, x), d(x, y)\}$. The case $j \leq k$ is analogous. Now suppose that $i < k < j$. Then $d(x, y) \gtrsim_c d(x, a_j) \gtrsim_2 d(x, a_k)$, and we obtain that $d(x, y) \gtrsim_{(c+2)} d(x, a_k)$. The inequality $d(x, y) \gtrsim_{(c+2)} d(a_k, y)$ is completely analogous. This completes the proof. $\square$

The proof of the following result is similar to the proof of Lemma 3.7.

Lemma 3.8 *If H_i is x- and y-admissible,* $d(x, y) \lesssim_c \min\{d(y, a_i), d(x, a_{i+1})\}$, $d(a_k, y) \gtrsim_{(c+2)} d(a_k, x)$ *for any* $k \le i$*, and* $d(a_k, x) \gtrsim_{(c+2)} d(a_k, y)$ *for any* $k > i$*, then the location of x and y in H_i so that x is to the left of y is* $(x, y, c+2)$*-admissible.*

Next we investigate the pairwise admissible locations of x and y in function of the mutual geometric location of the segments $H(x)$ and $H(y)$ and of the values $d(x, y), d_x$, and d_y.

Lemma 3.9 *If* $H(x) \cap H(y) = \emptyset$*, then* $d(x, y) \gtrsim_5 \max\{d_x, d_y\}$ *and any x-admissible hole H_i and any y-admissible hole H_j define a* $(x, y, 4)$*-admissible pair.*

Proof Suppose without loss of generality that $x \preccurlyeq y$, in particular, $i < j$. Denote by $\prec$ the resulting total order on $P \cup \{x, y\}$. Since there exists a $(x, y, 1)$-admissible location in which x belongs $\underline{H}_x$, the order $x \prec \overline{b}_x \prec \overline{a}_x \prec y$ is ϵ-compatible, therefore $d(x, y) \gtrsim_2 d(\overline{b}_x, \overline{a}_x)$ holds. On the other hand, Lemma 3.4 implies that $d(\overline{b}_x, \overline{a}_x) \gtrsim_3 d_x$, thus yielding $d(x, y) \gtrsim_5 d_x$. Analogously, we can prove that $d(x, y) \gtrsim_5 d_y$. To establish the second assertion, by Lemma 2.2 it suffices to show that $d(u, w) \gtrsim_2 \max\{d(u, v), d(v, w)\}$ for any three elements $u, v, w \in P \cup \{x, y\}$ such that $u \prec v \prec w$. Since H_i is x-admissible and H_j is y-admissible, it suffices to establish this inequality in the case when $\{u, v, w\} = \{x, y, a_k\}$ for any element $a_k \in P$. If $a_k \notin I^\circ(x) \cup I^\circ(y)$, then either $a_k \preccurlyeq x \preccurlyeq y$, or $x \preccurlyeq a_k \preccurlyeq y$, or $x \preccurlyeq y \preccurlyeq a_k$, and the required inequality follows from the ϵ-compatibility of the canonical partial order $\preccurlyeq$. Now suppose that $a_k \in I^\circ(x) \cup I^\circ(y)$, say $a_k \in I^\circ(x)$. Then $x?a_k$ and $a_k \preccurlyeq y$. Since $x \preccurlyeq y$, from Lemma 2.7 we infer that $d(x, y) \approx_2 d(a_k, y)$ and $d(x, a_k) \lesssim_2 \min\{d(x, y), d(a_k, y)\}$. If $a_k \prec x$, then $d(a_k, y) \gtrsim_2 d(a_k, x)$ and $d(a_k, x) \gtrsim_2 d(x, y)$. Analogously, if $x \prec a_k$, we obtain $d(x, y) \gtrsim_2 d(x, a_k)$ and $d(x, y) \gtrsim_2 d(a_k, y)$, establishing the result. □

Lemma 3.10 *If* $H_x \cap H_y = \emptyset$*, then for all* $a_j \in I^\circ(x)$*,* $a_i \in I^\circ(y)$ *we have* $d(a_i, a_j) \approx_8 d(x, y)$*.*

Proof Let $y \preccurlyeq x$. By previous lemma, any location of x and y in bounding holes is pairwise 4ϵ-admissible, therefore the orders $\underline{a}_y \prec y \prec \underline{b}_y \prec a_j \prec a_i \prec \overline{b}_x \prec x \prec \overline{a}_x$ and $a_j \prec \overline{b}_y \prec y \prec \overline{a}_y \prec \underline{a}_x \prec x \prec \underline{b}_x \prec a_i$ are both 4ϵ-compatible, yielding $d(a_i, a_j) \approx_8 d(x, y)$. □

Lemma 3.11 *If $H(x)$ and $H(y)$ overlap, then* $d(x, y) \gtrsim_3 \max\{d_x, d_y\}$ *and any x-admissible hole of* $H(x) \setminus H(y)$ *and any y-admissible hole of* $H(y) \setminus H(x)$ *define a* $(x, y, 4)$*-admissible pair.*

Proof Suppose without loss of generality that $I^\circ(x) \setminus I^\circ(y)$ is to the left of $I^\circ(y) \setminus I^\circ(x)$. The leftmost hole of $H(x)$ belongs to a $(x, y, 1)$-admissible pair, thus the order $x \prec \underline{a}_y \prec y$ is ϵ-admissible. Since $\underline{a}_y \in I^\circ(x)$, we obtain $d(x, y) \gtrsim_2 d(x, \underline{a}_y) \approx_1 d_x$, i.e. $d(x, y) \gtrsim_3 d_x$. Analogously, one can show that $d(x, y) \gtrsim_3 d_y$.

To establish the second assertion, we will show that any triplet $\{x, y, a_k\}$ satisfies the conditions of Lemma 2.2. First, let $a_k \preccurlyeq \underline{a}_x$. If $x?y$, from Lemma 2.7 we infer that $d(a_k, x) \approx_2 d(a_k, y)$ and $d(x, y) \lesssim_2 \min\{d(a_k, x), d(a_k, y)\}$, yielding the required inequality. Otherwise, we must have $x \preccurlyeq y$ and therefore $a_k \preccurlyeq x \preccurlyeq y$ is ϵ-compatible by the definition of $\preccurlyeq$. This settles the case $a_k \preccurlyeq \underline{a}_x$. The case $\overline{a}_y \preccurlyeq a_k$ is analogous. Now, suppose without loss of generality that $a_k \in I^\circ(x)$. If $a_k \in I^\circ(x) \cap I^\circ(y)$, then $d(x, a_k) \approx_1 d_x$ and $d(y, a_k) \approx_1 d_y$. Also $x \prec a_k \prec y$ by the conditions of the lemma. Since $d(x, y) \gtrsim_3 \max\{d_x, d_y\}$ by first part, we conclude that $d(x, y) \gtrsim_4 \max\{d(x, a_k), d(a_k, y)\}$. Finally, suppose that $a_k \in I^\circ(x) \setminus I^\circ(y)$. Then $x?a_k$, $x?y$, $y?\overline{a}_x$, $a_k \preccurlyeq \{y, \overline{a}_x\}$, and $x \preccurlyeq \overline{a}_x$. By Lemma 2.9, we have $d(x, y) \approx_4 d(a_k, y)$. To complete the proof, it suffices to show that $d(a_k, x) \lesssim_4 \min\{d(a_k, y), d(x, y)\}$. Since the location of x is x-admissible, the order $a_k \prec x \prec \overline{a}_x$ is ϵ-compatible. Analogously, the order $a_k \prec \overline{a}_x \prec y$ is ϵ-compatible. Thus $d(a_k, x) \lesssim_2 d(a_k, \overline{a}_x) \lesssim_2 d(a_k, y)$, yielding $d(a_k, x) \lesssim_4 d(a_k, y)$. Since $a_k \in I^\circ(x)$, we have $d(x, a_k) \approx_1 d_x \lesssim_3 d(x, y)$, thus $d(a_k, x) \lesssim_4 d(x, y)$. ☐

In the next two results, we assume that $H(x) \circ H(y)$. We will suppose without loss of generality that $H(y) \setminus H(x)$ is to the left of $H(x) \setminus H(y)$. Then either $y \preccurlyeq x$ or $x?y$ hold, otherwise the location of y in the leftmost hole of $H(y)$ cannot be extended to a $(x, y, 1)$-admissible location. Then the bounding hole $H' := [\overline{b}_y, \overline{a}_y]$ of $H(y)$ is an inner hole of $H(x)$ and, vice-versa, the bounding hole $H'' := [\underline{a}_x, \underline{b}_x]$ of $H(x)$ is an inner hole of $H(y)$. Notice also that these bounding holes are distinct and that H'' is to the left of H'.

Lemma 3.12 *If $H(x) \circ H(y)$, then $d(x, y) \approx_3 d_x, d_y$. Moreover, if $a_k \in I^\circ(x) \cup I^\circ(y)$, then $d(x, y) \approx_6 d(a_k, x), d(a_k, y)$.*

Proof Since $H' \in AH(y)$, there exists an x-admissible hole H_j such that H' and H_j define a pairwise ϵ-admissible location of y and x. If H_j is different from H', we obtain that either $d(x, \overline{b}_y) \lesssim_2 d(x, y) \lesssim_2 d(x, \overline{a}_y)$ or $d(x, \overline{a}_y) \lesssim_2 d(x, y) \lesssim_2 d(x, \overline{b}_y)$. Since H' is an inner hole of x, from Lemma 3.1 we infer that $d(x, \overline{b}_y) \approx_2 d(x, \overline{a}_x)$. In both cases we conclude that $d(x, y) \approx_3 d_x$. Now suppose that $H_j = H'$. Then $d(x, y)$ as well as the distances from x and y to $\overline{b}_y$ and $\overline{a}_y$ are all smaller or equal than $d(\overline{b}_y, \overline{a}_y) + 2\epsilon$. At least one of the inequalities $d(x, \overline{b}_y) \gtrsim_2 d(x, y)$ or $d(x, \overline{a}_y) \gtrsim_2 d(x, y)$ holds, because y is located between x and one of the elements $\overline{b}_y$ and $\overline{a}_y$. Since H' is an inner hole of $H(x)$, from Lemma 3.1 we conclude that $d_x \gtrsim_1 d(x, \overline{a}_y)$ and $d_x \gtrsim_1 d(x, \overline{b}_y)$, yielding $d_x \gtrsim_3 d(x, y)$. Analogously, using H'' instead of H', one can show that $d_y \gtrsim_3 d(x, y)$. Since $d(x, y) \gtrsim_3 d_x$ and $d(x, y) \gtrsim_3 d_y$ by Lemma 3.11, we conclude that $d(x, y) \approx_3 d_x$ and $d(x, y) \approx_3 d_y$, thus establishing the first assertion.

To prove the second assertion, pick any $a_k \in I^\circ(x) \cup I^\circ(y)$. If $a_k \in I^\circ(x)$, then Lemma 3.1 and previous assertion yield $d(a_k, x) \approx_1 d_x$ and $d_x \approx_3 d(x, y)$, whence $d(a_k, x) \approx_4 d(x, y)$. Analogously, if $a_k \in I^\circ(y)$ we conclude that $d(a_k, y) \approx_4 d(x, y)$. So, suppose that $a_k \in I^\circ(y) \setminus I^\circ(x)$. We assert that in this case $d(a_k, x) \approx_6 d(x, y)$. Since the hole $[\underline{a}_y, \underline{b}_y]$ is y-admissible, there exists a pairwise $(x, y, 1)$-admissible location such that y belongs to $[\underline{a}_y, \underline{b}_y]$ and $y \prec a_k \prec x$. Therefore

$d(x,y) \gtrsim_2 d(x,a_k)$. On the other hand, since $a_k \prec \underline{b}_x \prec x$ and $\underline{b}_x \in I^\circ(x)$, we infer that $d(x,a_k) \gtrsim_2 d(\underline{b}_x, x) \gtrsim_1 d_x \gtrsim_3 d(x,y)$, yielding $d(x,a) \approx_6 d(x,y)$. □

Lemma 3.13 *If $H(x) \circ H(y)$, then any x-admissible hole H_i and any y-admissible hole H_j define a $(x,y,12)$-admissible pair.*

Proof Denote the resulting total order on $P \cup \{x,y\}$ by $\prec$. To establish the result, by Lemma 2.2 we must show that $d(u,w) \gtrsim_{12} \max\{d(u,v), d(v,w)\}$ for any u, v, w such that $u \prec v \prec w$ and $\{u,v,w\} = \{x,y,a_k\}$ for $a_k \in P$.

If $a_k \notin I^\circ(x) \cup I^\circ(y)$, then either $a_k \preccurlyeq \{x,y\}$ or $\{x,y\} \preccurlyeq a_k$. Then $a_k \prec \{x,y\}$ in the first case and $\{x,y\} \prec a_k$ in the second case. If $x?y$, then from Lemma 2.7 we infer that $d(a_k,x) \approx_2 d(a_k,y)$ and $d(x,y) \lesssim_2 \min\{d(a_k,x), d(a_k,y)\}$. Since either $a_k \prec \{x,y\}$ or $\{x,y\} \prec a_k$, we conclude that independently of the relative location of x and y we will obtain the required inequality. Now suppose that $y \preccurlyeq x$. Since the partial order $\preccurlyeq$ is ϵ-compatible, if $a_k \preccurlyeq y$, then $d(a_k,x) \gtrsim_2 \max\{d(a_k,y), d(y,x)\}$, and if $x \preccurlyeq a_k$, then $d(y,a_k) \gtrsim_2 \max\{d(y,x), d(x,a_k)\}$. Therefore, if $y \prec x$, then we will obtain the stronger inequality $d(u,w) \gtrsim_2 \max\{d(u,v), d(v,w)\}$. Notice also that independently of the position of a_k we will also have $d(a_k,x) \gtrsim_2 d(x,y)$ and $d(a_k,y) \gtrsim_2 d(x,y)$. It remains to establish a relationship between $d(a_k,x)$ and $d(a_k,y)$ if $x \prec y$. Pick an element $a_i \in I^\circ(x) \cap I^\circ(y)$ (such an element exists because $H(x) \circ H(y)$). Since $a_k \notin I^\circ(x) \cup I^\circ(y)$, by Lemma 3.5 $d(a_k,x) \approx_2 d(a_k,a_i)$ and $d(a_k,y) \approx_2 d(a_k,a_i)$, yielding $d(a_k,x) \approx_4 d(a_k,y)$. This concludes the proof in the case $a_k \notin I^\circ(x) \cup I^\circ(y)$.

Now assume that $a_k \in I^\circ(x) \cup I^\circ(y)$. By Lemma 3.12 we have $d(a_k,x) \approx_6 d(x,y)$, $d(a_k,y) \approx_6 d(x,y)$, and therefore $d(a_k,x) \approx_{12} d(a_k,y)$. Thus we will obtain the required inequality independently of the relative location of x, y and a_k. □

Now, let $H(x) * H(y)$. Suppose that the hole $H_i = H(x) \cap H(y)$ is the leftmost hole of $H(y)$ (and the rightmost hole of $H(x)$). Then either $x \preccurlyeq y$ or $x?y$ holds. Indeed, if $y \preccurlyeq x$, then there is no $(x,y,1)$-admissible pair where x is located in the leftmost hole of $H(x)$.

Lemma 3.14 *If $H(x) * H(y)$, then $d(x,y) \gtrsim_3 \max\{d_x, d_y\}$ and $d(x,y) \approx_4 \delta_i$. Additionally, if $x?y$, then $d(x,y) \approx_3 d_x, d_y$.*

Proof First, let $x \preccurlyeq y$. Since $H_i \in AH(y)$, there exists $H_j \in H(x)$ such that the pair H_j, H_i is $(x,y,1)$-admissible. Hence $d(x,y) \lesssim_2 d(x,a_{i+1}) \lesssim_2 d(a_i,a_{i+1}) = \delta_i$ in view of Lemma 3.4. On the other hand, there exists a $(x,y,1)$-admissible pair such that x is placed in the leftmost hole of $H(x)$ and y in some hole of $H(y)$. Then a_i is between x and y, and we conclude that $d(x,y) \gtrsim_2 \max\{d(x,a_i), d(y,a_i)\}$. Since $d(x,a_i) \approx_1 d_x$, we infer that $d(x,y) \gtrsim_3 d_x$ (analogously one can show that $d(x,y) \gtrsim_3 d_y$). On the other hand, since $d(y,a_i) \approx_2 \delta_i$ by Lemma 3.4, we conclude that $d(x,y) \gtrsim_4 \delta_i$, hence $d(x,y) \approx_4 \delta_i$.

If $x?y$, then there exists a $(x,y,1)$-admissible pair such that x is placed to the right of y. This is possible only if both x and y are located in H_i, i.e., $a_i \prec x \prec y \prec a_{i+1}$ is ϵ-compatible. This implies that $d(x,y) \lesssim_2 d(a_i,a_{i+1}) = \delta_i$. Since $d(x,y) \gtrsim_4$

$d(a_i, a_{i+1})$ can be shown as above, we conclude that $d(x, y) \approx_4 \delta_i$. Now, since $a_i \in I^\circ(x)$ and $a_{i+1} \in I^\circ(y)$, by Lemma 3.1 we obtain that $d(x, y) \lesssim_2 d(x, a_i) \lesssim_1 d_x$ and $d(x, y) \lesssim_2 d(y, a_{i+1}) \lesssim_1 d_y$. Since $d(x, y) \gtrsim_3 \max\{d_x, d_y\}$ by Lemma 3.14, we deduce that $d(x, y) \approx_3 d_x, d_y$. □

Lemma 3.15 *If $H(x) * H(y)$, then any location of x and y in their bounding holes is $(x, y, 12)$-admissible, subject to the constraint that if they are located in H_i, then x is to the left of y.*

Proof If x is located in its leftmost hole and y is located in its rightmost hole, then this location is $(x, y, 4)$-admissible by Lemma 3.11. Now suppose that both x and y are located in H_i, x to the left of y. From Lemmata 3.14 and 3.4 we conclude that $d(x, y) \lesssim_4 d(a_i, a_{i+1}) \lesssim_2 d(x, a_{i+1})$ and $d(x, y) \lesssim_4 d(a_i, a_{i+1}) \lesssim_2 d(y, a_i)$, thus $d(x, y) \lesssim_6 \min\{d(x, a_{i+1}), d(y, a_i)\}$. Pick a_k with $k \le i$. Since $a_k \notin I^\circ(y)$ and $a_{i+1} \in I^\circ(y)$, from Lemma 3.5 we infer that $d(a_k, y) \approx_2 d(a_k, a_{i+1})$. Now, $d(a_k, a_{i+1}) \gtrsim_2 d(a_k, x)$ because the hole H_i is x-admissible. This shows that $d(a_k, y) \gtrsim_4 d(a_k, x)$. Analogously, we can show that $d(a_k, x) \gtrsim_4 d(a_k, y)$ for any $k > i$. By Lemma 3.8, the location of x and y in H_i, x to the left of y, is $(x, y, 8)$-admissible.

Now suppose that x is located in H_i and y is located in its rightmost hole H_j (the case when y is located in H_i and x in its leftmost hole is analogous). To apply Lemma 3.7, we must show that $\max\{d(x, a_j), d(y, a_{i+1})\} \lesssim_{10} d(x, y) \lesssim_{10} \min\{d(y, a_i), d(x, a_{j+1})\}$. Since $d(x, y) \lesssim_6 \min\{d(x, a_{i+1}), d(y, a_i)\}$ and $d(x, a_{i+1}) \lesssim_2 d(x, a_{j+1})$, we obtain $d(x, y) \lesssim_8 \min\{d(y, a_i), d(x, a_{j+1})\}$. On the other hand, since $a_{i+1} \in I^\circ(y)$, from Lemma 3.14 we obtain $d(x, y) \gtrsim_3 d_y \gtrsim_1 d(a_{i+1}, y)$, yielding $d(x, y) \gtrsim_4 d(y, a_{i+1})$. Finally, since $a_i \in I^\circ(x)$ and $a_j \notin I^\circ(x)$, from Lemma 3.5 we infer that $d(x, a_j) \approx_2 d(a_i, a_j)$. Now, $d(y, a_i) \gtrsim_2 d(a_j, a_i)$, because H_j is y-admissible, yielding $d(y, a_i) \gtrsim_4 d(x, a_j)$. Note that $d(y, a_i) \approx_2 d(a_i, a_{i+1})$ by Lemma 3.4. Since $d(x, y) \approx_4 d(a_i, a_{i+1})$ by Lemma 3.14, thus $d(x, y) \gtrsim_{10} d(x, a_j)$. □

Lemma 3.16 *If $H(y) \Subset H(x)$, then $d(x, y) \approx_3 d_x \gtrsim_6 d_y$ and $d(x, y) \gtrsim_7 d_y$. Any location of x and y in their bounding holes is $(x, y, 12)$-admissible, subject to the constraint that if x and y are located in the same bounding hole, then x is between y and the elements of $I^\circ(y)$.*

Proof Denote by H_i the leftmost hole of $H(x)$ and by H_j the leftmost hole of $H(y)$. To establish the first assertion, suppose without loss of generality that H_i does not belong to $H(y)$. This implies that H_j is an inner hole of $H(x)$. Since H_i is x-admissible, the order $x \prec a_{i+1} \prec y$ is ϵ-compatible, whence $d(x, y) \gtrsim_2 d(x, a_{i+1}) \approx_1 d_x$, yielding $d(x, y) \gtrsim_3 d_x$. On the other hand, the hole H_j is y-admissible, therefore one of the total orders $x \prec y \prec a_{j+1}$ or $a_j \prec y \prec x$ is ϵ-compatible. Since $a_j, a_{j+1} \in I^\circ(x)$, in the first case we obtain $d_x \approx_1 d(x, a_{j+1}) \gtrsim_2 d(x, y)$ while in the second case we obtain $d_x \approx_1 d(a_j, x) \gtrsim_2 d(x, y)$, therefore $d_x \gtrsim_3 d(x, y)$ in both cases. This shows that $d(x, y) \approx_3 d_x$. On the other hand, since the hole H_j is y-admissible, we have $d(a_j, a_{j+1}) \gtrsim_2 \max\{d(y, a_j), d(y, a_{j+1})\}$. Since $a_{j+1} \in I^\circ(y)$, from previous inequalities we deduce that $d_x \approx_1 d(x, a_{j+1}) \gtrsim_2 d(a_j, a_{j+1}) \gtrsim_2 d(y, a_{j+1}) \approx_1 d_y$, yielding $d_x \gtrsim_6 d_y$. To show that $d(x, y) \gtrsim_7 d_y$, notice that if we consider the $(x, y, 1)$-admissible pair $\{H_i, H_k\}$, then $d(x, y) \gtrsim_2 d(a_k, y)$. Now, if $H_k \neq H_j$, then $a_k \in I^\circ(y)$, and we conclude that $d(x, y) \gtrsim_3 d_y$. On the other hand, if $H_k = H_j$, then $d(a_k, y) \gtrsim_5 d_y$ by Lemma 3.4, yielding $d(x, y) \gtrsim_7 d_y$. This establishes the first part of the lemma.

To establish the second assertion, assume that x is located in its leftmost hole H_i (the case when x is located in its rightmost hole is analogous). First suppose that y is also located in H_i, to the right of x (i.e., $H_j = H_i$). We will show that this location is $(x, y, 10)$-admissible by using Lemma 3.8. Pick any element a_k. If $a_k \notin I^\circ(x)$, then $a_k \notin I^\circ(y)$, and Lemma 3.5 yields $d(a_k, a_{i+1}) \approx_2 d(a_k, x)$ and $d(a_k, a_{i+1}) \approx_2 d(a_k, y)$, therefore $d(a_k, x) \approx_4 d(a_k, y)$. If $a_k \in I^\circ(x) \setminus I^\circ(y)$, then $d(a_k, y) \approx_2 d(a_k, a_{i+1})$ by Lemma 3.5, and $d(a_k, x) \gtrsim_2 d(a_k, a_{i+1})$ because H_i is x-admissible, yielding $d(a_k, x) \gtrsim_4 d(a_k, y)$. Finally, if $a_k \in I^\circ(x) \cap I^\circ(y)$, then $d(a_k, x) \approx_1 d_x \gtrsim_6 d_y \approx_1 d(a_k, y)$, whence $d(a_k, x) \gtrsim_8 d(a_k, y)$. This establishes the first condition of Lemma 3.8. It remains to show that $d(x, y) \lesssim_8 \min\{d(x, a_{i+1}), d(y, a_i)\}$. Since $d(x, a_{i+1}) \approx_1 d_x \approx_3 d(x, y)$, we conclude that $d(x, a_{i+1}) \gtrsim_4 d(x, y)$. On the other hand, Lemma 3.4 yields $d(a_i, y) \approx_2 d(a_i, a_{i+1}) \gtrsim_2 d(x, a_{i+1}) \approx_1 d_x \approx_3 d(x, y)$, whence $d(a_i, y) \gtrsim_8 d(x, y)$. From Lemma 3.8 we infer that the location of x and y in H_i is $(x, y, 10)$-admissible.

Now suppose that x is located in H_i and y is located in the bounding hole H_k, which is different from H_i. To apply Lemma 3.7, we must show that

$$\max\{d(x, a_k), d(y, a_{i+1})\} \lesssim_{10} d(x, y) \lesssim_{10} \min\{d(x, a_{k+1}), d(y, a_i)\}.$$

Since $a_k \in I^\circ(x)$ and $d_x \approx_3 d(x, y)$, we have $d(x, a_k) \approx_1 d_x \approx_3 d(x, y)$, thus $d(x, a_k) \approx_4 d(x, y)$. If $a_{i+1} \notin I^\circ(y)$, then $a_{i+1} \preccurlyeq y$. Since the hole H_i is x-admissible, there exists an ϵ-compatible order such that $x \prec a_{i+1} \prec y$, whence $d(x, y) \gtrsim_2 d(a_{i+1}, y)$. On the other hand, if $a_{i+1} \in I^\circ(y)$, then $d(y, a_{i+1}) \approx_1 d_y \lesssim_9 d(x, y)$, yielding $d(x, y) \gtrsim_{10} d(y, a_{i+1})$. To show that $d(x, y) \lesssim_{10} d(y, a_i)$, notice that if $a_k \in I^\circ(y)$, then $d(y, a_i) \approx_2 d(a_i, a_k)$ by Lemma 3.5, and if $a_k \notin I^\circ(y)$, then $d(a_i, y) \gtrsim_2 d(a_i, a_k)$ because $a_i \preccurlyeq a_k \preccurlyeq y$ holds. Since H_i is x-admissible, the order $a_i \prec x \prec a_k$ is ϵ-compatible, thus $d(a_i, a_k) \gtrsim_2 d(x, a_k) \approx_1 d_x \approx_3 d(x, y)$. Hence $d(y, a_i) \gtrsim_8 d(x, y)$ in both cases. It remains to show that $d(x, a_{k+1}) \gtrsim_{10} d(x, y)$. If $a_{k+1} \in I^\circ(x)$, then $d(x, a_{k+1}) \approx_1 d_x \approx_3 d(x, y)$, and we are done. Otherwise, since $a_{k+1} \in H(y) \Subset H(x)$, a_{k+1} is the rightmost element of $H(x)$. From Lemma 3.4 we infer that $d(x, a_{k+1}) \gtrsim_5 d_x \approx_3 d(x, y)$, thus $d(x, a_{k+1}) \gtrsim_8 d(x, y)$. The conditions of Lemma 3.7 are satisfied, whence H_i, H_k is $(x, y, 12)$-admissible pair. □

Finally, consider the case $H(x) = H(y)$. Then $I^\circ(x) = I^\circ(y)$. If $a_k \notin I^\circ(x)$ and $a_l \in I^\circ(x)$, then Lemma 3.5 implies that $d(a_k, x) \approx_2 d(a_k, a_l)$ and $d(a_k, y) \approx_2 d(a_k, a_l)$, thus $d(a_k, x) \approx_4 d(a_k, y)$. In the next three results, we will assume without loss of generality that $d_x \geq d_y$.

Lemma 3.17 *If* $H(x) = H(y)$ *and* $d(x, y) \ll_3 \max\{d_x, d_y\}$, *then any* $(x, y, 1)$-*admissible location of* x *and* y *is in common* x- *and* y-*admissible holes. On the other hand, any location of* x *and* y *in a common bounding hole is* $(x, y, 6)$-*admissible*, *subject to the constraint that if* $d_y \ll_4 d_x$, *then* y *is located between* x *and the elements of* $I^\circ(y) = I^\circ(x)$.

Proof Suppose by way of contradiction that an x-admissible hole H_i and an y-admissible hole H_j with $i < j$ define a $(x, y, 1)$-admissible location. Then $d_x > d(x, y) + 3\epsilon \geq d(x, a_{i+1}) + 3\epsilon - 2\epsilon \geq d_x$ because $a_{i+1} \in I^\circ(x)$, a contradiction.

Now, let x and y be both located in the leftmost hole H_i of $H(x) = H(y)$. For any $a_k \notin I^\circ(x)$ we know that $d(a_k, x) \approx_4 d(a_k, y)$. On the other hand, if $a_k \in I^\circ(x)$, then $d(a_k, x) \approx_1 d_x$ and $d(a_k, y) \approx_1 d_y$. First, let y be located to the left of x. From the second condition of the lemma we infer that $d_y \gtrsim_4 d_x$. Thus $d(a_k, y) \gtrsim_6 d(a_k, x)$ in this case. To use Lemma 3.8 we must show that $d(x, y) \lesssim_2 \min\{d(x, a_i), d(y, a_{i+1})\}$. Notice that $d(y, a_{i+1}) \approx_1 d_y \gtrsim_4 d_x \gg_3 d(x, y)$, whence $d(y, a_{i+1}) \gtrsim_2 d(x, y)$. On the other hand, Lemma 3.4 yields $d(x, a_i) \gtrsim_5 d_x \gg_3 d(x, y)$, whence $d(x, a_i) \gtrsim_2 d(x, y)$. Now assume that x is located to the left of y. Then $d_x \gtrsim_4 d_y$, and the proof is obtained by switching the roles of x and y. □

Lemma 3.18 *If* $H(x) = H(y)$ *and* $d(x, y) \gg_3 \max\{d_x, d_y\}$, *then the only* $(x, y, 1)$-*admissible locations are the two locations of* x *and* y *in different bounding holes of* $H(x) = H(y)$.

Proof For any other location of x and y, one of these two elements, say y, necessarily will be located between x and an element a_k of $I^\circ(x) = I^\circ(y)$. Then $d(x, a_k) \lesssim_1 d_x \ll_3 d(x, y)$, contrary to our assumption that this location is $(x, y, 1)$-admissible. Now, if the location of x in the leftmost hole H_i and of y in the rightmost hole H_j of $H(x)$ is not $(x, y, 1)$-admissible, then the unique $(x, y, 1)$-admissible location is that of x in $[a_j, a_{j+1}]$ and of y in $[a_i, a_{i+1}]$, therefore $\overline{b}_x = a_j \preccurlyeq x$ and $y \preccurlyeq a_{i+1} = \underline{b}_y$ hold, contrary to the definition of $\underline{a}_y$ and $\overline{a}_x$. □

Lemma 3.19 *If* $H(x) = H(y)$ *and* $d(x, y) \approx_3 d_x \gg_6 d_y$, *then any* $(x, y, 1)$-*admissible location is that of* x *in one of the bounding holes and of* y *in any* y-*admissible hole*, *subject to the constraint that if* x *and* y *are located in the same bounding hole*, *then* y *must be located between* x *and the elements of* $I^\circ(x) = I^\circ(y)$. *Vice versa*, *any such location of* x *and* y *in their bounding holes is* $(x, y, 10)$-*admissible.*

Proof Consider a $(x, y, 1)$-admissible location, such that x is located in an inner hole H_k and suppose that y is located to the right of x. Then $a_k \in I^\circ(x) = I^\circ(y)$, therefore $d_y \geq d(a_k, y) - \epsilon \geq d(x, y) - 3\epsilon \geq d_x - 6\epsilon > d_y$, a contradiction. Conversely, let x be located in the leftmost hole H_i of $H(x) = H(y)$. First assume that y is also located in H_i to the right of x. Then $d(x, a_{i+1}) \approx_1 d_x \approx_3 d(x, y)$, whence $d(x, a_{i+1}) \gtrsim_4 d(x, y)$. On the other hand, from Lemma 3.4 we infer that $d(y, a_i) \approx_2 d(a_i, a_{i+1}) \gtrsim_2 d(x, a_{i+1}) \approx_1 d_x \approx_3 d(x, y)$, whence $d(y, a_i) \gtrsim_8 d(x, y)$. Since $d(a_k, x) \approx_4 d(a_k, y)$ for any $a_k \notin I^\circ(x)$ and $d(a_k, x) \approx_1 d_x \gtrsim_6 d_y \approx_1 d(a_k, y)$ for any $a_k \in I^\circ(x)$, Lemma 3.8 yields that this location is $(x, y, 10)$-admissible. Now assume that y is located in the rightmost hole H_j of $H(x) = H(y)$. We know already that $d(y, a_i) \gtrsim_8 d(x, y)$. Analogously, from Lemma 3.4 we deduce that

$d(x, a_{j+1}) \gtrsim_5 d_x \approx_3 d(x, y)$, yielding $d(x, a_{j+1}) \gtrsim_8 d(x, y)$. It remains to show that $d(x, y) \gtrsim_8 \max\{d(x, a_j), d(y, a_{i+1})\}$. Since $a_{i+1}, a_j \in I^\circ(x) = I^\circ(y)$, we have $d(x, a_j) \approx_1 d_x$ and $d(y, a_{i+1}) \approx_1 d_y$. Now, since $d(x, y) \approx_3 d_x \gg_6 d_y$, we obtain the required inequality. ☐

Lemma 3.20 *If* $H(x) = H(y)$, $d(x, y) \approx_3 d_x, d_y$, *and then any location of* x *and* y *in their bounding holes is* $(x, y, 10)$*-admissible.*

Proof The proof is analogous to the second part of the proof of Lemma 3.19. ☐

4 Distributing the Elements to Holes

In this section, we describe how, for each hole $H_i = [a_i, a_{i+1}]$, to compute the set X_i of the elements of X° which will be located in H_i. The set X_i consists of certain elements x such that H_i is the leftmost or the rightmost hole of $H(x)$. Additionally, each set X_i will be partitioned into an ordered list of cells, and to each cell we perform a recursive call of the algorithm.

For two vertices a_i, a_j $(i < j)$ of the chain P, denote by H_{ij} the union of all holes comprised between a_i and a_j. Let X_{ij} be the set of all $x \in X^\circ$ such that $H(x) = H_{ij}$. The sets X_{ij} constitute a partition of X°. In the next two subsections, we will describe how to partition (if possible) the set X_{ij} into two subsets X^-_{ij} and X^+_{ij}, so that all elements of X^-_{ij} will be located in the hole H_i and all elements of X^+_{ij} will be located in the hole H_{j-1}; see Fig. 1.

4.1 Linked and Separated Pairs

Two elements $x, y \in X_{ij}$ are called *linked* (notation $x \sim y$) if in all $(x, y, 1)$-admissible locations x and y must be placed in the same hole and *separated* (notation $x|y$) if in all $(x, y, 1)$-admissible locations x and y must be placed in distinct bounding holes. For two subsets A and B of X_{ij}, we write $A \sim B$ if $x \sim y$ for all $x \in A, y \in B$, and say that A and B are linked (instead of $A \sim A$, we write $\sim A$). Analogously, we write $A|B$ if $x|y$ for all $x \in A$ and $y \in B$, and say that A and B are separated. Obviously, $\sim$ is an equivalence relation, while $A|B$ implies that $\sim A$ and $\sim B$. Notice that $\sim A, \sim B$ and $x|y$ for some $x \in A, y \in B$ imply $A|B$. Finally, $A|B$ and $A'|B$ imply $A \sim A'$.

In general, we can compute in polynomial time all linked and all separated pairs of X_{ij}. Therefore we can compute the equivalence classes of $\sim$. Hypothetically, on these equivalence classes we can define a graph $\mathcal{S}^*_{ij}$, in which we draw an edge

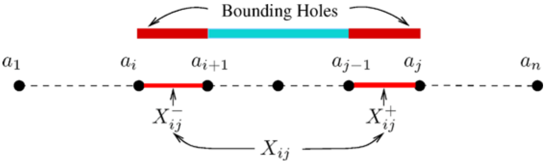

Fig. 1 Bounding holes and the partition of X_{ij} into X^-_{ij} and X^-_{ij}

between two classes A, B if and only if A and B are located in different bounding holes in any ϵ-compatible order. The graph $\mathcal{S}^*_{ij}$ must be bipartite, otherwise no ϵ-compatible order on X_{ij} exist. If $\mathcal{S}^*_{ij}$ is also connected, then (up to symmetry) we have only one bipartition of X_{ij}. However, if $\mathcal{S}^*_{ij}$ is bipartite but contains several connected components, then we have to decide which bipartition of X_{ij} must be selected. On the other hand, we cannot construct the graph $\mathcal{S}^*_{ij}$ in polynomial time, therefore we have to discover its edges (or at least a part of them) in an efficient way. To address these questions, we will employ the quantitative conditions for two elements to be linked or separated provided by Lemmata 3.17 and 3.18. Namely, let

$$S_{ij} = \{xy : x, y \in X_{ij} \text{ and } d(x,y) \gg_3 \max\{d_x, d_y\}\},$$

$$L_{ij} = \{xy : x, y \in X_{ij} \text{ and } d(x,y) \ll_3 \max\{d_x, d_y\}\}.$$

By Lemma 3.18 we have $x|y$ for all $xy \in S_{ij}$, and by Lemma 3.17, we have $x \sim y$ for all $xy \in L_{ij}$. Since $\sim$ is an equivalence relation, necessarily all vertices of the same connected component of the graph $\mathcal{L}_{ij} = (X_{ij}, L_{ij})$ are linked, and therefore all they must be located in the same in all ϵ-compatible orders. We call the connected components of $\mathcal{L}_{ij}$ *blocks*. We continue by investigating in which cases two blocks of $\mathcal{L}_{ij}$ are separated or linked.

Two elements $x, y \in X_{ij}$ are called *strongly separated* if $d(x,y) \gg_9 \max\{d_x, d_y\}$. For two elements $x, y \in X_{ij}$, set $x \twoheadrightarrow y$ if and only if one of the following two conditions is satisfied:

(A1) $d_x \ll_4 d_y$;
(A2) $d_x \gtrsim_4 d_y$ and there exists an element $z \in X_{ij}$ such that $xz, yz \notin L_{ij}$ and $d(x,z) \ll_{16} d(y,z)$.

As we will show below in Lemma 4.2, the arc $x \twoheadrightarrow y$ prescribes the order between the elements x and y in all ϵ-compatible orders in which x and y are not located in different bounding holes H_i and H_j.

Lemma 4.1 *If $x, y, z \in X_{ij}$ satisfy the condition* (A2), *then y and z are strongly separated.*

Proof Since $xz, yz \notin L_{ij}$, we have $d(x,z) \gtrsim_3 \max\{d_x, d_z\}$ and $d(y,z) \gtrsim_3 \max\{d_y, d_z\}$. On the other hand, $d_x \gtrsim_4 d_y$. Putting these things together, we conclude that $d(y,z) \gg_{16} d(x,z) \gtrsim_3 \max\{d_x, d_z\} \gtrsim_4 \max\{d_y, d_z\}$. Thus $d(y,z) \gg_9 \max\{d_y, d_z\}$. □

Lemma 4.2 *If $x \twoheadrightarrow y$, then $x \prec y$ in all ϵ-compatible orders $\prec$ such that $a_{i+1} \prec \{x, y\}$ and $y \prec x$ in all ϵ-compatible orders $\prec$ such that $\{x, y\} \prec a_{j-1}$.*

Proof Let $\prec$ be an ϵ-compatible order such that $a_{i+1} \prec \{x, y\}$. First, assume that $d_x \ll_4 d_y$. Since $a_{i+1} \in I^\circ(x) \cap I^\circ(y)$, we have $d(a_{i+1}, x) \approx_1 d_x$ and $d(a_{i+1}, y) \approx_1 d_y$, thus $d(a_{i+1}, x) \ll_2 d(a_{i+1}, y)$. Since $a_{i+1} \prec \{x, y\}$, from Lemma 2.7 we infer that $x \prec y$. Otherwise, if $d_x \gtrsim_4 d_y$, then the arc $x \twoheadrightarrow y$ has type (A2), thus $y|z$ by Lemma 4.1. Since $a_{i+1} \prec y$, this shows that $z \prec a_{i+1}$, whence $z \prec \{x, y\}$. But $d(x,z) \ll_{16} d(y,z)$, therefore we must have $z \prec x \prec y$. □

Now, on X_{ij} we define a directed graph $\overrightarrow{\mathcal{L}_{ij}}$: we draw an arc $x \to y$ with tail x and head y if and only if one of the following two conditions is satisfied:

(L1) $x \twoheadrightarrow y$ and x, y belong to a common block of $\mathcal{L}_{ij}$;
(L2) $d(x, y) \ll_5 \max\{d_x, d_y\}$.

Obviously, if (L2) is satisfied, then $xy \in L_{ij}$ and $y \to x$ holds as well. We call the strongly connected components of $\overrightarrow{\mathcal{L}_{ij}}$ *cells* and we say that two (not necessarily distinct) vertices $x, y \in X_{ij}$ are *strongly linked* if they belong to a common cell of $\overrightarrow{\mathcal{L}_{ij}}$. The following lemma shows that the partition of X_{ij} into cells refines the partition into blocks.

Lemma 4.3 *If x and y belong to a common cell of $\overrightarrow{\mathcal{L}_{ij}}$, then they belong to a common block of $\mathcal{L}_{ij}$. In particular, every block of $\mathcal{L}_{ij}$ is a disjoint union of cells of $\overrightarrow{\mathcal{L}_{ij}}$.*

Proof Since x and y are strongly linked, there exists a directed path of $\overrightarrow{\mathcal{L}_{ij}}$ running from x to y. Pick any arc $u \to v$ on this path. If it has type (L2), then uv is an edge of $\mathcal{L}_{ij}$. Otherwise, if $u \to v$ has type (L1), then, by definition, u and v belong to a common block. Thus the extremities of all arcs $u \to v$ of any directed path between x, y belong to a common block B of $\mathcal{L}_{ij}$. By transitivity, we conclude that x and y also belong to B. ☐

Lemma 4.4 *Let $x, x', y \in X_{ij}$. If x, x' belong to a common cell of $\overrightarrow{\mathcal{L}_{ij}}$, but $\{x, x'\}$ and y belong to different blocks of $\mathcal{L}_{ij}$, then there does not exist an ϵ-compatible order such that $x \prec y \prec x'$.*

Proof Suppose by way of contradiction that $x \prec y \prec x'$ in some ϵ-compatible order $\prec$ on X. Since x and x' belong to a common cell of $\overrightarrow{\mathcal{L}_{ij}}$, there exists a directed path R connecting x to x' in this graph. Since $x \prec y$ and $y \prec x'$, the path R necessarily contains an arc $u \to u'$ such that $u \prec y$ and $y \prec u'$. Since x, x', u, u' are strongly linked, by Lemma 4.3 they belong to a common block of $\mathcal{L}_{ij}$. Therefore $\{u, u'\}$ and y belong to different blocks of $\mathcal{L}_{ij}$, thus we may assume without loss of generality that $x = u$ and $x' = u'$, i.e., that $x \to x'$.

First, assume that the arc $x \to x'$ has type (L2), i.e., $d(x, x') \ll_5 \max\{d_x, d_{x'}\}$. Since $x \prec y \prec x'$, we infer that $d(x, x') \gtrsim_2 \max\{d(x, y), d(x', y)\}$, thus $\max\{d_x, d_{x'}\} \gg_3 \max\{d(x, y), d(x', y)\}$. Now, if $d_x \le d_{x'}$, then $\max\{d_{x'}, d_y\} \ge \max\{d_{x'}, d_x\} \gg_3 d(x', y)$, therefore $x'y \in L_{ij}$. Otherwise, if $d_{x'} \le d_x$, we conclude in the same way that $xy \in L_{ij}$. In each case, we obtain that $\{x, x'\}$ and y belong to the same block of $\mathcal{L}_{ij}$, contrary to our assumption.

Now, assume that $x \to x'$ has type (L1), i.e., $x \twoheadrightarrow x'$ and x, x' belong to a common block of $\mathcal{L}_{ij}$. Since x and x' belong to a common cell of $\overrightarrow{\mathcal{L}_{ij}}$, there exists a directed path Q from x' to x. Since $y \prec x'$ and $x \prec y$, while running along Q from x' to x we will necessarily meet an arc $u \to v$ such that $y \prec u$ and $v \prec y$. Hence $v \prec u$. The elements x, x', u, v are linked, thus they must be located in the same hole. Since $x \twoheadrightarrow x'$ and $x \prec x'$, from Lemma 4.2 we infer that $a_{i+1} \prec \{x, x'\}$. As u, v are located in the same hole with x and x', we deduce that $a_{i+1} \prec \{u, v\}$. Since $v \prec u$, Lemma 4.2

shows that we cannot have $u \twoheadrightarrow v$. Thus $u \to v$ has type (L2), and we can employ the same proof as in the case when $x \to x'$ satisfies (L2). $\square$

Lemma 4.5 *Let B', B'' be two distinct blocks of $\mathcal{L}_{ij}$. If there exist $x, x' \in B'$ and $y, y' \in B''$, such that the pairs xx' and yy' are strongly linked, $x \twoheadrightarrow y$ and $y' \twoheadrightarrow x'$, then $B'|B''$.*

Proof Suppose by way of contradiction that there exists an ϵ-compatible order $\prec$ such that either $a_{i+1} \prec \{x, x', y, y'\}$ or $\{x, x', y, y'\} \prec a_{j-1}$ holds, say the first. Since $x \twoheadrightarrow y$ and $y' \twoheadrightarrow x'$, Lemma 4.2 implies that $x \prec y$ and $y' \prec x'$. Therefore, if $y \prec x'$, then $x \prec y \prec x'$, otherwise, if $x' \prec y$, then $y' \prec x' \prec y$. Analogously, if $y = y'$, we have $x \prec y \prec x'$, and if $x = x'$, we have $y' \prec x' \prec y$. In other words, either y is located between x and x' or x is located between y and y', say the first. Since $\{x, x'\}$ and y belong to different blocks of $\mathcal{L}_{ij}$, by Lemma 4.4 there is no ϵ-compatible order such that y is located between x and x', contrary to our assumption. $\square$

Lemma 4.6 *Let C', C'' be two distinct cells of $\overrightarrow{\mathcal{L}_{ij}}$. If there exist $x, x' \in C'$ and $y, y' \in C''$ such that $x \twoheadrightarrow y$ and $y' \twoheadrightarrow x'$, then C', C'' belong to different blocks and $C'|C''$.*

Proof By Lemma 4.3, x, x' belongs to a common block B' and y, y' belongs to a common block B''. Therefore, if x and y belong to a common block, by transitivity we conclude that $B' = B''$. Since $x \twoheadrightarrow y$ and $y' \twoheadrightarrow x'$, we deduce that they are arcs of type (L1), thus $x \to y$ and $y' \to x'$ hold. This is impossible, because $\{x, x'\}$ and $\{y, y'\}$ belong to distinct cells. Thus, $B' \neq B''$. From Lemma 4.5 we infer that $B'|B''$. This shows that $C'|C''$. $\square$

Now, define a graph Ψ_{ij} whose vertices are the cells of $\overrightarrow{\mathcal{L}_{ij}}$ and we draw an edge between two cells C' and C'' if and only if at least one of the following two conditions is satisfied:

(S1) there exist an element x belonging to the same block as C' and an element y belonging to the same block as C'' such that $xy \in S_{ij}$;
(S2) there exist two elements x, x' belonging to the same block as C' and two elements y, y' belonging to the same block as C'' such that the pairs xx' and yy' are strongly linked, and $x \twoheadrightarrow y$, $y' \twoheadrightarrow x'$.

According to Lemmata 3.18 and 4.6, in both cases (S1) and (S2) we have $C'|C''$. Analogously to the graph $\mathcal{S}^*_{ij}$, the graph Ψ_{ij} must be bipartite, otherwise no ϵ-compatible order exist. Notice also, that all cells from the same block have one and the same neighborhood in Ψ_{ij}. Now, we complete the bipartite graph Ψ_{ij} in the following way: for each connected component of Ψ_{ij} consider its canonical bipartition (two vertices are in the same part if they can be connected by a path of even length) $\{A', A''\}$, and draw an edge between any two cells, one from A' and another from A''. Denote the obtained graph also by Ψ_{ij} (notice that every connected component of this graph is a complete bipartite graph). Call the union of all cells from A' (respectively, from A'') a *cluster*. The clusters $\mathcal{K}'$ and $\mathcal{K}''$ of A' and A'' are called *twins*. From the construction, we immediately obtain the following result:

Lemma 4.7 *All elements of a cluster $\mathcal{K}$ are linked and each pair of twin clusters $\mathcal{K}'$ and $\mathcal{K}''$ is separated, i.e., we have $\sim \mathcal{K}$ and $\mathcal{K}'|\mathcal{K}''$.*

A connected bipartite component $\{\mathcal{K}', \mathcal{K}''\}$ of Ψ_{ij} is called a *principal component* if there exist $x \in \mathcal{K}'$ and $y \in \mathcal{K}''$ such that x and y are strongly separated.

Lemma 4.8 *If two strongly separated pairs xy and $x'y'$ belong to different principal components of the bipartite graph Ψ_{ij} and $d(x,y) \geq d(x',y')$, then $d(x,y) \gg_4 d(x',y')$, $d_x \gg_4 d_{x'}, d_{y'}$, and $d_y \gg_4 d_{x'}, d_{y'}$. In particular, for any ϵ-compatible order $\prec$, the elements x', y' are located between x and y.*

Proof By Lemma 3.18, in any ϵ-compatible order $\prec$, the elements x and y as well as x' and y' must be located in different bounding holes of H_{ij}, say x, x' in the leftmost hole and y, y' in the rightmost hole (i.e. $\{x, x'\} \prec \{y, y'\}$). Up to symmetry, we distinguish two cases: $x \prec x' \prec y \prec y'$ and $x \prec x' \prec y' \prec y$. If $x \prec x' \prec y \prec y'$, then we assert that $xy' \in S_{ij}$. Indeed, $d(x,y') \gtrsim_2 \max\{d(x,y), d(x',y')\} \gg_9 \max\{d_x, d_{y'}\}$, yielding $x'y \in S_{ij}$. Since x and y' are separated, we conclude that the elements x, y', x', and y all belong to the same connected component of Ψ_{ij}, contrary to our hypothesis.

Now, let $x \prec x' \prec y' \prec y$. Notice that $d(x,y') \gg_7 \max\{d_{y'}, d_{x'}\}$: indeed, since x' is located between x and y', we obtain $d(x,y') \gtrsim_2 d(x',y') \gg_9 \max\{d_{x'}, d_{y'}\}$. Now, if $d_x \lesssim_4 d_{x'}$ or $d_x \lesssim_4 d_{y'}$, we deduce that $d(x,y') \gg_3 d_x$, yielding that $xy' \in S_{ij}$. Since this is impossible, we conclude that $d_x \gg_4 \max\{d_{x'}, d_{y'}\}$. Analogously, we can show that $d_y \gg_4 \max\{d_{x'}, d_{y'}\}$. It remains to prove that $d(x,y) \gg_4 d(x',y')$. Indeed, if $d(x',y') \gtrsim_4 d(x,y)$, then we obtain $d(x,y') \gtrsim_2 d(x',y') \gtrsim_4 d(x,y) \gg_9 d_x$, thus $d(x,y') \gg_3 d_x$. Since $d(x,y') \gg_7 d_{y'}$, we deduce that $xy' \in S_{ij}$. This contradiction concludes the proof of the lemma. □

4.2 Partitioning X_{ij} into X^-_{ij} and X^+_{ij}

We describe how to partition the set X_{ij} into the subsets X^-_{ij} and X^+_{ij} which will be located, respectively, in the bounding holes H_i and H_{j-1} of H_{ij}. For this, we define (yet another!) directed graph $\overrightarrow{\mathcal{G}_{ij}}$ whose vertices are the cells of the directed graph $\overrightarrow{\mathcal{L}_{ij}}$, and there exists an arc $C' \rightarrowtail C$ with tail C' and head C if and only if at least one of the following three conditions is satisfied:

- (G1) C' and C belong to twin clusters of the bipartite graph Ψ_{ij};
- (G2) C' and C are not connected by (G1)-arcs and there exist $x \in C$ and $x' \in C'$ such that $d_{x'} \ll_4 d_x$;
- (G3) C' and C are not connected by (G1) and (G2)-arcs and there exist $x \in C$, $x' \in C'$, and $z \in X_{ij}$ such that $xz, x'z \notin L_{ij}$ and $d(x',z) \ll_{16} d(x,z)$.

In other words, there is an arc $C' \rightarrowtail C$ if and only if either these cells are separated in Ψ_{ij} (in this case, we also have $C \rightarrowtail C'$ and both $C' \rightarrowtail C$ and $C \rightarrowtail C'$ are (G1)-arcs) or there exists $x \in C$ and $x' \in C'$ such that $x' \twoheadrightarrow x$. Notice that between two cells C' and C'' we can have at most one arc of type (G2) or (G3). Indeed, otherwise, if we have two such arcs $C' \rightarrowtail C$ and $C \rightarrowtail C'$, then (S2) implies that C and C' belong to twin clusters, thus these arcs have type (G1).

A cell which is a head of a (G3)-arc is called a (G3)-*cell*. A directed cycle $\mathcal{C}$ of $\overrightarrow{\mathcal{G}_{ij}}$ is called a (Gi)-*cycle* if all its arcs have type (Gi), $i = 1, 2, 3$. From the definition, we immediately conclude that any (G1)-cycle has length 2. Conversely, from Lemma 4.6 we infer that all 2-cycles of $\overrightarrow{\mathcal{G}_{ij}}$ consist of two edges of type (G1), thus the (G1)-cycles are exactly the cycles of length 2. A *mixed cycle* is a directed cycle of $\overrightarrow{\mathcal{G}_{ij}}$ containing arcs of both types (G2) and (G3). Finally, an *induced cycle* is a directed cycle $\mathcal{C}$ such that for two cells $C, C' \in \mathcal{C}$ we have $C' \rightarrowtail C$ if and only if C is the successor of C' in $\mathcal{C}$.

Our next goal is to establish that either no ϵ-compatible order exist or to show that the set of cells can be partitioned into two subsets such that the subgraphs of $\overrightarrow{\mathcal{G}_{ij}}$ induced by these subsets do not contain directed cycles. Notice that the general problem of deciding if the vertex-set of a directed graph can be partitioned into two subsets which induce acyclic subgraphs is NP-complete (problem CUT INTO ACYCLIC SUBGRAPHS from [42]). In our case, this can be done in polynomial time, using the special structure of the graph $\overrightarrow{\mathcal{G}_{ij}}$.

Lemma 4.9 *If $\mathcal{C}$ is a directed cycle of $\overrightarrow{\mathcal{G}_{ij}}$, then for any ϵ-compatible total order, $\mathcal{C}$ has a cell located in the bounding hole $H_i = [a_i, a_{i+1}]$ and a cell located in the bonding hole $H_{j-1} = [a_{j-1}, a_j]$.*

Proof The assertion is obvious if $\mathcal{C}$ is a (G1)-cycle. So, suppose that all arcs of $\mathcal{C}$ have type (G2) or (G3). Let $\mathcal{C} = (C_1, C_2, \ldots, C_k, C_1)$. The definition of cells implies that $\mathcal{C}$ contains two consecutive cells, say C_1 and C_k, which belong to different blocks. Suppose that there exists an ϵ-compatible order $\prec$ such that no element of $\bigcup_{l=1}^{k} C_l$ is located in the hole $H_j = [a_{j-1}, a_j]$, i.e., $a_{i+1} \prec \bigcup_{l=1}^{k} C_l$. In each C_l pick two elements x_l, y_l such that $x_l \twoheadrightarrow y_{l+1}$. From Lemma 4.2 we infer that $x_l \prec y_{l+1}$ for all $l = 1, \ldots, k$ (all indices here are modulo k).

We divide the cells of $\mathcal{C}$ into groups: a group consists of all consecutive cells of $\mathcal{C}$ belonging to one and the same block. The first group starts with C_1, while the last group ends with C_k. We assert that if $\{C_{l-q}, \ldots, C_l\}$ and $\{C_{l+1}, \ldots, C_{l+r}\}$ are two consecutive groups of $\mathcal{C}$, then $C_l \prec C_{l+1} \cup \cdots \cup C_{l+r}$ (all additions here are modulo k). Indeed, pick $u \in C_l$ and $v \in C_{l+1}$. Since $\{x_l, u\}$ and $\{y_{l+1}, v\}$ belong to different blocks while each of these pairs belong to a common cell, applying Lemma 4.4 to each of the triplets of the quadruplet x_l, u, y_{l+1}, v, we infer that in the total order $\prec$ none of y_{l+1}, v is located between x_l and u and none of x_l, u is located between y_{l+1} and v. Since $x_l \prec y_{l+1}$, we conclude that $\{x_l, u\} \prec \{y_{l+1}, v\}$, yielding $C_l \prec C_{l+1}$. Now, consider the cell C_{l+2}. The element y_{l+2} must be located to the right of x_{l+1}, therefore to the right of C_l. Since C_{l+2} and C_l belong to different blocks, we can show that $C_l \prec C_{l+2}$ by using exactly the same reasoning as for the cells C_l and C_{l+1}. Continuing this way, we obtain the required relationship $C_l \prec C_{l+1} \cup \cdots \cup C_{l+r}$. This establishes the assertion. Suppose that $[1, i_1], [i_1 + 1, i_2], \ldots, [i_r + 1, k]$ are the indices of cells defining the beginning and the end of each group. From our assertion we infer that $C_k \prec C_{i_1} \prec C_{i_2} \prec \cdots \prec C_{i_r} \prec C_k$, contrary that $\prec$ is a total order. □

Lemma 4.10 *If C and C' are two* (G3)*-cells, then either C and C' belong to the same principal component of Ψ_{ij} or there is a* (G2)*-arc from C to C' or from C' to C.*

Proof By definition of (G3)-arcs, there exist $x \in C$ and $x' \in C'$ which are heads of two (A2)-arcs whose tails are outside C and C', respectively. By Lemma 4.1, we can find two elements z and z' such that xz and $x'z'$ are strongly separated pairs. If xz and $x'z'$ belong to distinct principal components of $\mathcal{L}_{ij}$ and $d(x,z) \geq d(x',z')$, from Lemma 4.8 we infer that $d_x \gg_4 d_{x'}, d_{z'}$, and $d_z \gg_4 d_{x'}, d_{z'}$. Therefore x' and z' are connected by arcs of type (A1) with both x and z. In particular, this shows that there is a (G2)-arc from C' to C. □

Lemma 4.11 *If $C \rightarrowtail C'$ is a* (G3)*-arc and the cell C belongs to a principal component of Ψ_{ij}, then C and C' belong to the same cluster of this component. In particular, either $\overrightarrow{\mathcal{G}_{ij}}$ does not contain* (G3)*-cycles or no ϵ-compatible order exist.*

Proof Suppose that xy is a strongly separated pair with $x \in C$. Since $C \rightarrowtail C'$ is a (G3)-arc, there exist $y' \in C$ and $x' \in C'$ such that $y' \twoheadrightarrow x'$ is an (A2)-arc. By Lemma 4.1, there exists z' such that $x'z'$ is a strongly separated pair. If xz and $x'y'$ belong to different principal components of Ψ_{ij}, from Lemma 4.8 we infer that either there exists a (G2)-arc from C' to C or from C to C'. In the first case, C and C' obey the condition (S2), thus we cannot have a (G3)-arc from C to C'. Analogously, in the second case, we deduce that we have at the same time a (G3)-arc and a (G2)-arc from C to C'. Since this is impossible, C and C' belong to a common principal component and to the same cluster of this component.

Now, if $\overrightarrow{\mathcal{G}_{ij}}$ contains a (G3)-cycle, then the first assertion implies that its cells all belong to one and the same cluster, therefore they are linked, contrary to Lemma 4.9. □

Lemma 4.12 *The graph $\overrightarrow{\mathcal{G}_{ij}}$ does not contain* (G2)*-cycles.*

Proof Suppose by way of contradiction that $\mathcal{C} = (C_1, C_2, \ldots, C_k, C_1)$ is a (G2)-cycle of $\overrightarrow{\mathcal{G}_{ij}}$. In each cell C_i we can select two elements x_i, y_i such that $d_{x_i} \ll_4 d_{y_{i+1}}$ for all $i = 1, \ldots, k$ (here and throughout this proof all indices are modulo k). On the other hand, since there is no arc of type (G2) or (G3) running from C_{i+1} to C_i, we obtain that $d_{y_i} \lesssim_4 d_{x_{i+1}}$ for all $i = 1, \ldots, k$. Putting all these $2k$ inequalities together, we conclude that $d_{x_i} \ll_4 d_{y_{i+1}} \lesssim_4 d_{x_{i+2}}$, yielding that $d_{x_i} < d_{x_{i+2}}$ for $i = 1, \ldots, k$. If k is even, then $d_{x_1} < d_{x_3} < \cdots < d_{x_{k-1}} < d_{x_1}$ and if k is odd, then $d_{x_1} < d_{x_3} < \cdots < d_{x_k} < d_{x_2} < d_{x_4} < \cdots < d_{x_{k-1}} < d_{x_1}$, a contradiction. □

Lemma 4.13 *Let $C'_1 \rightarrowtail C_1$ and $C'_2 \rightarrowtail C_2$ be two* (G2)*-arcs. If $C'_1 \neq C_2$, $C'_2 \neq C_1$ and neither C'_1, C_2 nor C'_2, C_1 do not define* (G1)*-cycles, then there exists either a* (G2)*-arc from C'_1 to C_2 or a* (G2)*-arc from C'_2 to C_1. In particular, any induced cycle $\mathcal{C}$ of $\overrightarrow{\mathcal{G}_{ij}}$ contains exactly one or two* (G2)*-arcs, and if $\mathcal{C}$ contains two such arcs, then they are consecutive.*

Proof Since $C'_1 \rightarrowtail C_1$ and $C'_2 \rightarrowtail C_2$ are (G2)-arcs, we can find the elements $y_1 \in C'_1, x_1 \in C_1, y_2 \in C'_2$, and $x_2 \in C_2$ such that $d_{y_1} \ll_4 d_{x_1}$ and $d_{y_2} \ll_4 d_{x_2}$. Suppose by way of contradiction that there do not exist (G2)-arcs neither between C'_1 and C_2, nor between C'_2 and C_1. This implies that $d_{x_2} \lesssim_4 d_{y_1}$ and $d_{x_1} \lesssim_4 d_{y_2}$, whence $d_{y_1} \ll_4 d_{x_1} \lesssim_4 d_{y_2} \ll_4 d_{x_2} \lesssim_4 d_{y_1}$, thus $d_{y_1} < d_{y_1}$, and we obtain a contradiction. This establishes the first assertion. To prove the second assertion, let $\mathcal{C}$ be an induced cycle containing two non-consecutive (G2)-arcs $C'_1 \rightarrowtail C_1$ and $C'_2 \rightarrowtail C_2$. Since $C'_1 \neq C_2$, $C'_2 \neq C_1$, either C'_1, C_2 or C'_2, C_1 define a (G1)-cycle, or there exists a (G2)-arc from C'_1 to C_2 or a (G2)-arc from C'_2 to C_1. In all cases we obtain a contradiction with the assumption that the cycle $\mathcal{C}$ is induced. □

Lemma 4.14 *Let* $C' \rightarrowtail C$ *be a* (G3)-*arc*, $C \rightarrowtail C''$ *be a* (G2)-*arc*, *and suppose that there is no* (G2)-*arc from* C' *to* C''. *Let* $x' \in C', x, y \in C, y'' \in C''$, *and* $z \in \overline{C}$ *be the elements defining these arcs*, *i.e.*, *satisfying* $x'z, xz \notin L_{ij}, d_{x'} \gtrsim_4 d_x, d(x',z) \ll_{16} d(x,z)$, *and* $d_y \ll_4 d_{y''}$. *Then the following holds for any* ϵ-*compatible order*:

- (i) *if* $d(x,x') \gtrsim_2 d(x,z)$, *then* $C \mid C'$ *and the cells* C, C' *belong to distinct twin clusters of the same principal component of* Ψ_{ij};
- (ii) *if* $d(x,x') \ll_2 d(x,z)$ *and* $d(y'',z) \gtrsim_2 d(x,z)$, *then* $C \sim C''$ *and the cells* C, C'' *belong to the same cluster of* Ψ_{ij};
- (iii) *if* $d(x,x') \ll_2 d(x,z)$ *and* $d(y'',z) \ll_2 d(x,z)$, *then* $C \mid C''$.

In particular, *if* C, C' *do not belong to distinct twin clusters and* C, C'' *do not belong to the same cluster*, *then* $C \mid C''$.

Proof By Lemma 4.1 the elements x and z are strongly separated, thus $C \mid \overline{C}$. First, let $d(x,x') \gtrsim_2 d(x,z)$, i.e., $d(x,x') \gg_{14} d(x',z)$. Since $x'z \notin L_{ij}$, we have $d(x',z) \gtrsim_3 \max\{d_{x'}, d_z\}$. Since $d_{x'} \gtrsim_4 d_x$, we deduce that $d(x,x') \gg_{14} d(x',z) \gtrsim_7 \max\{d_{x'}, d_x\}$, whence $d(x,x') \gg_7 \max\{d_{x'}, d_x\}$. This shows that $x \mid x'$, yielding $C \mid C'$. This settles case (i). Further suppose that $d(x,x') \ll_2 d(x,z)$. Since $d(x',z) \ll_{16} d(x,z)$, in any ϵ-compatible order the cell C' must be located between the cells C and $\overline{C}$. Now, additionally assume that $d(y'',z) \gtrsim_2 d(x,z)$. Then $d(y'',z) \gg_{14} d(x',z)$. Since there is no (G2)-arc from C' to C'', we have $d_{x'} \gtrsim_4 d_{y''}$. On the other hand, since $x'z \notin L_{ij}$, we have $d(x',z) \gtrsim_3 \max\{d_{x'}, d_z\}$. Combining all these inequalities, we deduce that $d(y'',z) \gg_7 \max\{d_{y''}, d_z\}$, establishing that $y'' \mid z$. This shows that $C'' \mid \overline{C}$, thus $C'' \sim C$, establishing the assertion (ii). Finally, suppose that $d(y'',z) \ll_2 d(x,z)$. Since $d_y \ll_4 d_{y''}$, if C'' is located in the same bounding hole as C or in an inner hole, then C'' cannot be located between C and $\overline{C}$. Thus C must be located between C'' and $\overline{C}$, contrary to the assumption that $d(y'',z) \ll_2 d(x,z)$. This shows that C'' must be located in the same bounding hole as $\overline{C}$, i.e., $C \mid C''$, thus establishing (iii). □

Notice that in case (iii), the fact that $C \mid C''$ does not necessarily mean that C and C'' belong to distinct twin clusters. It may happen that C and C'' belong to distinct components of the graph Ψ_{ij}. Lemma 4.14 can be viewed as a new separation condition for cells.

To complete the partition of cells into two subsets which induce acyclic subgraphs of $\overrightarrow{\mathcal{G}_{ij}}$, it remains to deal with mixed cycles, in particular with induced mixed cycles. Let $\mathcal{C}$ be an induced cycle. From Lemma 4.13 we infer that $\mathcal{C}$ contains either one (G2)-arc (we call such a $\mathcal{C}$ a 1-*cycle*) or two consecutive (G2)-arcs (we call such a $\mathcal{C}$ a 2-*cycle*), all remaining arcs of $\mathcal{C}$ being (G3)-arcs. By Lemma 4.11, the heads of all (G3)-arcs of $\mathcal{C}$ are (G3)-cells of one and the same cluster $\mathcal{K}$. In this case, we say that the cycle $\mathcal{C}$ *intersects* the cluster $\mathcal{K}$.

For each (G2)-arc $C_0 \rightarrowtail C$ and for each cluster $\mathcal{K}$, we will show how to detect if there exists a 1-cycle or a 2-cycle $\mathcal{C}$ passing via $C_0 \rightarrowtail C$ and intersecting $\mathcal{K}$. First, we consider the case of 1-cycles. Then necessarily C_0 must be a (G3)-cell of $\mathcal{K}$. Notice that a required induced cycle, if it exists, cannot contain cells C' such that $C_0 \rightarrowtail C'$ is an arc of type (G2) or (G3). Therefore, we can remove all such cells of $\mathcal{K}$ from our consideration. Analogously, we should remove all cells C' such that $C' \rightarrowtail C$ is an arc. Now, in the subgraph induced by the remaining cells of $\mathcal{K}$ we search for a shortest directed path $Q = C \rightarrowtail C_1 \rightarrowtail C_2 \rightarrowtail \cdots \rightarrowtail C_k \rightarrowtail C_0$ subject to the constraint that the first arc $C \rightarrowtail C_1$ and the last arc $C_k \rightarrowtail C_0$ of this path are (G3)-arcs. This can be done in polynomial time by testing all possible choices for C_1 and C_k and applying for each such pair a shortest path finding algorithm in an acyclic graph. If such a path Q does not exist, then no required induced cycle $\mathcal{C}$ exist. Otherwise, we claim that the arcs of the path Q together with the arc $C_0 \rightarrowtail C$ define an induced cycle $\mathcal{C}$ having exactly one (G2)-arc. Indeed, if $C_i \rightarrowtail C_j$ is an arc of type (G2) or (G3) and $|i - j| \geq 2$, since the subgraph induced by $\mathcal{K}$ is acyclic, we must have $i < j$. But this contradicts the minimality of the path Q. So, the resulting cycle is indeed induced. It remains to notice that $\mathcal{C}$ does not contain other (G2)-arcs, because by Lemma 4.13 in an induced cycle the (G2)-arcs are consecutive. Analogously, we can decide efficiently if there exists a 2-cycle passing via $C_0 \rightarrowtail C$, intersecting $\mathcal{K}$, and having a second (G2)-arc $C \rightarrowtail C_0'$. For this we must remove the cells C' which are tails or heads of arcs running between C and C', the cells which are heads of arcs with tails C_0, and the cells which are tails of arcs with head C_0'. In the subgraph induced by the remaining cells of $\mathcal{K}$, we search a shortest directed path Q from C_0' to C_0 starting and finishing with arc of type (G3). The remaining part of analysis is similar. Finally, in the same way we can decide efficiently if there exists a (G2)-cycle passing via $C_0 \rightarrowtail C$, intersecting $\mathcal{K}$, and having a second (G2)-arc $C_0' \rightarrowtail C_0$. Therefore, we have proven the following result:

Lemma 4.15 *For a* (G2)*-arc* $C_0 \rightarrowtail C$ *and a cluster* $\mathcal{K}$, *one can decide in polynomial time if there exists an induced* 1*- or* 2*-cycle* $\mathcal{C}$ *passing via* $C' \rightarrowtail C$ *and intersecting* $\mathcal{K}$.

For a cell C, let $\Omega_1(C)$ denote the set of (G2)-arcs $C_0 \rightarrowtail C$ which belong to a 1-cycle intersecting a cluster $\mathcal{K}$ not containing C. Let $\Omega_2(C)$ denote the set of (G2)-arcs $C_0 \rightarrowtail C$ which belong to a 2-cycle $\mathcal{C}$ intersecting a cluster $\mathcal{K}$ not containing C and passing via $C_0 \rightarrowtail C$ in such a way that the arc of $\mathcal{C}$ entering C_0 is a (G3)-arc. Notice that in both cases, C_0 belongs to the cluster $\mathcal{K}$ because C_0 is a head of a (G3)-arc of $\mathcal{C}$, and all such heads belong to $\mathcal{K}$. Finally, denote by $\Omega_3(C)$ the set of (G2)-arcs $C \rightarrowtail C_0$ which belong to a 2-cycle $\mathcal{C}$ intersecting a cluster $\mathcal{K}$, such that C belongs to $\mathcal{K}$ and the arc of $\mathcal{C}$ entering C has type (G2). Figure 2 illustrates this classification.

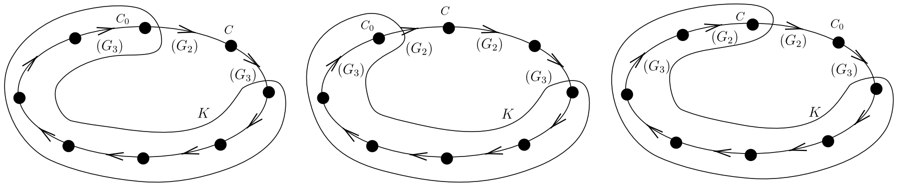

1-cycle: $C_0 \rightarrowtail C$ belongs to $\Omega_1(C)$ 2-cycle: $C_0 \rightarrowtail C$ belongs to $\Omega_2(C)$ 2-cycle: $C \rightarrowtail C_0$ belongs to $\Omega_3(C)$

Fig. 2 The classification of the arcs incident to a cell C

For each cell C of $\overrightarrow{\mathcal{G}_{ij}}$ we introduce a binary variable x_C. This set of variables x_C satisfies the following conditions:

(F1) $x_{C'} = x_{C''}$, if the cells C', C'' belongs to the same cluster;
(F2) $x_{C'} \neq x_{C''}$, if the cells C', C'' belong to distinct clusters of the same component;
(F3) $x_C \neq x_{C_0}$, if the arc $C_0 \rightarrowtail C$ belongs to $\Omega_1(C) \cup \Omega_2(C)$;
(F4) $x_C \neq x_{C_0}$, if the arc $C \rightarrowtail C_0$ belongs to $\Omega_3(C)$.

Now, we can define a 2-SAT formula Φ_{ij} by replacing every constraint of the form $a = b$ by two clauses $(a \vee \bar{b})$ and $(\bar{a} \vee b)$ and every constraint of the form $a \neq b$ by two clauses $(a \vee b)$ and $(\bar{a} \vee \bar{b})$ (Φ_{ij} can be solved in linear time of its length using the well known algorithm of [4]).

Proposition 4.16 *If the* 2-*SAT formula* Φ_{ij} *admits a satisfying assignment* A, *then the sets* $X_{ij}^- = \{C : A(x_C) = 0\}$ *and* $X_{ij}^+ = \{C : A(x_C) = 1\}$ *define a partition of cells of* $\overrightarrow{\mathcal{G}_{ij}}$ *into two acyclic subgraphs. Conversely, given an* ϵ-*compatible order on* X, *the assignment* A *defined by setting* $A(x_C) = 0$ *if* C *is located in the hole* H_i, $A(x_C) = 1$ *if* C *is located in the hole* H_{j-1}, *and* $A(x_{C'}) = A(x_{C''}) \in \{0, 1\}$ *if* C' *and* C'' *are located in a common inner hole, is a true assignment for* Φ_{ij}. *In particular, if* Φ_{ij} *is not satisfiable, then no* ϵ-*compatible order exist.*

Proof First suppose that A is a true assignment of Φ_{ij} and the partition X_{ij}^-, X_{ij}^+ of X_{ij} is defined as above. Denote by $\overrightarrow{\mathcal{G}_{ij}^-}$ and $\overrightarrow{\mathcal{G}_{ij}^+}$ the subgraphs induced by the sets X_{ij}^- and X_{ij}^+. Condition (F1) forces that every cluster will be included in the same set. Condition (F2) implies that any pair of twin clusters will be separated. Hence $\overrightarrow{\mathcal{G}_{ij}^-}$ and $\overrightarrow{\mathcal{G}_{ij}^+}$ do not contain (G1)-cycles: if C and C' are the two cells of a (G1)-cycle, then the constraint $(x_C \vee x_{C'}) \wedge (\bar{x}_C \vee \bar{x}_{C'})$ yields $A(x_C) \neq A(x_{C'})$. From Lemma 4.12 we know that $\overrightarrow{\mathcal{G}_{ij}}$ does not contain (G2)-cycles. Since the cells of every (G3)-cycle are all contained in the same cluster and we know that each cluster induces an acyclic subgraph, we conclude that $\overrightarrow{\mathcal{G}_{ij}^-}$ and $\overrightarrow{\mathcal{G}_{ij}^+}$ do not contain (G3)-cycles as well. It remain to deal with mixed cycles. Suppose by way of contradiction that $\overrightarrow{\mathcal{G}_{ij}^+}$ contains a mixed cycle. Then it also contains an induced mixed cycle $\mathcal{C}$. From Lemma 4.13 we infer that $\mathcal{C}$ contains either exactly one (G2)-arc $C_0 \rightarrowtail C$ or exactly two consecutive (G2)-arcs $C_0 \rightarrowtail C \rightarrowtail C''$. In the first case, we conclude that $C_0 \rightarrowtail C$ belongs to $\Omega_1(C)$, thus according to (F3) we must have $x_C \neq x_{C_0}$, contrary to the fact that $A(x_C) = A(x_{C'}) = 1$. Analogously, in the second case, we deduce that either $x_C \neq x_{C_0}$ and the arc $C_0 \rightarrowtail C$ belongs to $\Omega_2(C)$ or $x_C = x_{C_0}$ and the arc $C \rightarrowtail C''$ belongs to $\Omega_3(C)$, whence $x_C \neq x_{C''}$. But then we obtain a contradiction with the assumption that $A(x_{C_0}) = A(x_C) = A(x_{C''}) = 1$. This contradiction establishes that the subgraphs $\overrightarrow{\mathcal{G}_{ij}^-}$ and $\overrightarrow{\mathcal{G}_{ij}^+}$ obtained from the true assignment A of Φ_{ij} are both acyclic.

Conversely, suppose that A is an assignment obtained from an ϵ-compatible order as defined in the proposition. We assert that A is a true assignment for Φ_{ij}, i.e., it satisfies the constraints (F1)–(F4). This is obviously true for constraints of type (F1) and (F2), because if two cells C', C'' belong to the same cluster, then they will be located in the same hole and we must have $A(x_{C'}) = A(x_{C''})$. If two cells C' and C'' belong to distinct twin clusters, then they must be separated, therefore the unique ϵ-admissible location of C' and C'' will be in different bounding holes, thus we will have $A(x_{C'}) \neq A(x_{C''})$.

Now, pick an arc $C_0 \rightarrowtail C$ which belongs to $\Omega_1(C) \cup \Omega_2(C)$. If $C_0 \rightarrowtail C$ belongs to $\Omega_1(C)$, then there exists a 1-cycle $\mathcal{C}$ passing via $C_0 \rightarrowtail C$ and intersecting some cluster $\mathcal{K}$. Since all cells of $\mathcal{C}$, except C, are heads of (G3)-arcs, they all belong to $\mathcal{K}$, i.e. they all have the same value in the assignment. Thus, by Lemma 4.9, C must be separated from C_0 (namely C and C' must be located in different bounding holes), showing that $A(x_C) \neq A(x_{C_0})$. If the arc $C_0 \rightarrowtail C$ belongs to $\Omega_2(C)$, then let $\mathcal{C}$ be a 2-cycle passing via $C_0 \rightarrowtail C$ and intersecting the cluster $\mathcal{K}$ not containing C. Additionally, we know that the arc $C' \rightarrowtail C_0$ of $\mathcal{C}$ entering C_0 is a (G3)-arc, thus C_0 belongs to $\mathcal{K}$. Since C' cannot belong to the twin cluster of $\mathcal{K}$ (this contradicts that $C' \rightarrowtail C_0$ is a (G3)-arc) and since C does not belong to $\mathcal{K}$, from Lemma 4.14 we infer that the cells C_0 and C must be separated, thus $A(x_C) \neq A(x_{C_0})$. Finally, suppose that $C \rightarrowtail C_0$ belongs to $\Omega_3(C)$. Then there exists a 2-cycle $\mathcal{C}$ passing via $C \rightarrowtail C_0$ and intersecting the cluster $\mathcal{K}$, such that C belongs to $\mathcal{K}$ and the arc of $\mathcal{C}$ entering C has type (G2). Since all cells of $\mathcal{C}$ except C and C_0 are heads of (G3)-arcs, they all belong to the cluster $\mathcal{K}$. Since C also belongs to this cluster, by Lemma 4.9, C_0 must be separated from the remaining cells of $\mathcal{C}$, yielding $x_C \neq x_{C_0}$. Hence the assignment A satisfies all constraints (F1)–(F4). This shows, in particular, that if Φ_{ij} is not satisfiable, then no ϵ-compatible order exists. ☐

4.3 Sorting the Cells of X_{ij}^- and X_{ij}^+

Let X_{ij}^- and X_{ij}^+ be the partition of X_{ij} obtained from the true assignment of the 2-SAT formula Φ_{ij}. Let $\overrightarrow{\mathcal{G}_{ij}^-}$ and $\overrightarrow{\mathcal{G}_{ij}^+}$ be the subgraphs of the graph $\overrightarrow{\mathcal{G}_{ij}}$ induced by X_{ij}^- and X_{ij}^+. We will locate all cells of X_{ij}^- in the bounding hole $H_i = [a_i, a_{i+1}]$ and all cells of X_{ij}^+ in the bounding hole $H_{j-1} = [a_{j-1}, a_j]$ of H_{ij}. The elements from two different cells C', C'' located in the same hole will not be mixed, i.e. either all elements of C' will be placed to the right of all elements of C'', or vice versa, all elements of C'' will be placed to the left of all elements of C'. To specify this relative order of the cells of X_{ij}^- and X_{ij}^+, we use the fact that both directed graphs $\overrightarrow{\mathcal{G}_{ij}^-}$ and $\overrightarrow{\mathcal{G}_{ij}^+}$ are acyclic, therefore each of them admits a topological order of its cells (recall that a total order on vertices of a directed graph is called a *topological order* if the tail of any arc is to the left of its head). We find a topological order $C_{j_1} \prec C_{j_2} \prec \cdots \prec C_{j_p}$ on the cells of X^+_{ij} and we find a dual topological order $C_{i_q} \prec C_{i_{q-1}} \prec \cdots \prec C_{i_1}$ on the cells of X^-_{ij}. Then we locate the cells of X^+_{ij} in H_{j-1} and the cells of X^-_{ij} in H_i following these two orders. As a result we obtain the following total order on the elements $a_i, a_{i+1}, a_{j-1}, a_j$ of the path P and the cells of X_{ij}:

$$a_i \prec C_{i_q} \prec C_{i_{q-1}} \prec \cdots \prec C_{i_1} \prec a_{i+1} \prec \cdots \prec a_{j-1} \prec C_{j_1} \prec C_{j_2} \prec \cdots \prec C_{j_p} \prec a_j.$$

The following two results relay the topological orders on the cells of X^-_{ij} and X^+_{ij} with the order on distances between elements from such cells.

Lemma 4.17 *Let C', C'' be two distinct cells of X^+_{ij}. If $C' \prec C''$ in the topological order, then for any $y \in C'$, $z \in C''$ and $x \in X^-_{ij}$, we have $d_y \lesssim_4 d_z$ and $d(x, y) \lesssim_{16} d(x, z)$.*

Proof Since C', C'' belong to the same set X^+_{ij}, C' and C'' are not connected by (G1)-arcs. Since $C' \prec C''$ in the topological order on X^+_{ij}, we cannot have an arc from C'' to C'. In particular, since $C'' \rightarrowtail C'$ is not a (G2)-arc, we must have $d_z \gtrsim_4 d_y$. Analogously, since $C'' \rightarrowtail C'$ is not a (G3)-arc, we deduce that $d(x, y) \lesssim_{16} d(x, z)$. □

Notice that an analogous property holds for any two elements y, z of two different cells C', C'' of X^-_{ij} such that $C' \prec C''$ and any element x of X^+_{ij}.

Lemma 4.18 *Let C, C', C'' be three distinct cells of the graph $\overrightarrow{\mathcal{G}_{ij}}$. If the algorithm returns the total order $\prec$ and $C \prec C' \prec C''$, then for any elements $x \in C, y \in C'$, $z \in C''$ we have $d(x, z) \gtrsim_{16} \max\{d(x, y), d(y, z)\}$.*

Proof Up to symmetry, we have two cases: (1) C belongs to X^-_{ij} and C', C'' belong to X^+_{ij} and (2) all three cells C, C', and C'' belong to X^+_{ij}. Consider the case (1). Lemma 4.17 yields $d(x, y) \lesssim_{16} d(x, z)$. Suppose by way of contradiction that $d(y, z) \gg_{16} d(x, z)$. If $yz \in L_{ij}$, then $d(y, z) \lesssim_3 \max\{d_y, d_z\}$, i.e. $d(x, z) \ll_{13} d_z$. Thus $d(x, z) \ll_3 d_z$, i.e., $xz \in L_{ij}$. This shows that x and z belong to the same cluster, and the algorithm should locate them in the same hole, contrary to $x \in C \subseteq X^-_{ij}$ and $z \in C'' \subseteq X^+_{ij}$. Thus $yz \notin L_{ij}$, i.e., $d(y, z) \gtrsim_3 \max\{d_y, d_z\}$. Since C' and C'' belong to X^+_{ij} and $C' \prec C''$, we cannot have an arc from C'' to C', in particular a (G2)-arc. Thus $d_y \lesssim_4 d_z$. From these inequalities we conclude that $d(y, z) \gg_9 \max\{d_y, d_z\}$, thus C' and C'' must be separated, which is not the case because C' and C'' belong to X^+_{ij}.

Now consider the case when the cells C, C' and C'' belong to X^+_{ij}. Then $d(x, y) \lesssim_3 \max\{d_x, d_y\}$, $d(y, z) \lesssim_3 \max\{d_y, d_z\}$, and $d(x, z) \lesssim_3 \max\{d_x, d_z\}$, otherwise one of these pairs will belong to S_{ij}. Since $C \neq C''$, the elements x and z are not strongly linked, therefore we also have $d(x, z) \gtrsim_5 \max\{d_x, d_z\}$. By construction, the order $x \prec y \prec z$ follows the topological order on the cells C, C', and C'' of X^+_{ij}. Since we do not have arcs from C'' to C, C' and from C' to C, we deduce that $d_y \lesssim_4 d_z$, $d_x \lesssim_4 d_z$, and $d_x \lesssim_4 d_y$. From previous inequalities we obtain

$d(x, y) \lesssim_3 \max\{d_x, d_y\} \lesssim_4 \max\{d_x, d_z\} \lesssim_5 d(x, z)$ and $d(y, z) \lesssim_3 \max\{d_y, d_z\} \lesssim_4 \max\{d_x, d_z\} \lesssim_5 d(x, z)$. Hence $\max\{d(x, y), d(y, z)\} \lesssim_{12} d(x, z)$. □

After fixing the relative position of each cell C of X_{ij}, we make a recursive call to C. For this, we update the canonical order $\preccurlyeq$ in the following way: if C is located in X^+_{ij}, then we set $x \preccurlyeq^+ y$ if $x \twoheadrightarrow y$, otherwise, if C is located in X^-_{ij}, then set $x \preccurlyeq^- y$ if $y \twoheadrightarrow x$. Obviously $\preccurlyeq^+$ and $\preccurlyeq^-$ are dual, therefore if we will apply to each of them the rules (PO1)–(PO6), then we will obtain two dual partial orders, which also will be denoted by $\preccurlyeq^+$ and $\preccurlyeq^-$.

Lemma 4.19 *Let C be a cell and let $\prec$ be a ϵ-compatible total order on X. Then the restriction of $\prec$ to C is an extension of one of the partial orders $\preccurlyeq^+$ or $\preccurlyeq^-$.*

Proof We know that all elements of C will be placed in the same hole, i.e., either we have $a_{i+1} \prec C$ or we have $C \prec a_j$. If $a_{i+1} \prec C$, then according to Lemma 4.2 we will have $x \prec y$ for all $x, y \in C$ such that $x \twoheadrightarrow y$. Since the properties (PO1)–(PO6) hold for all ϵ-compatible orders, we conclude that the total order $\prec$ is a linear extension of $\preccurlyeq^+$. The case $C \prec a_j$ is completely analogous: then $\prec$ is a linear extension of $\preccurlyeq^-$. □

The recursive call to a cell C, either return the answer “not”, in which case, no ϵ-compatible order exist, or it return a total order on C, which is 16ϵ-compatible by virtue of induction hypothesis. The total order between the cells of $\overrightarrow{\mathcal{G}_{ij}}$ described above and the total orders on cells are concatenated to give a single total order $\prec$ on the set X_{ij}.

Lemma 4.20 *Let C, C' be two distinct cells of the graph $\overrightarrow{\mathcal{G}_{ij}}$. If the algorithm returns the total order $\prec$, then for any three elements $x, y, z \in C \cup C'$ such that $x \prec y \prec z$, we have $d(x, z) \gtrsim_{16} \max\{d(x, y), d(y, z)\}$.*

Proof Up to symmetry, we have two cases: (1) C belongs to X^-_{ij} and C' belongs to X^+_{ij} and (2) the cells C and C' belong to X^+_{ij}. The analysis of case (2) is similar to that of case (2) of Lemma 4.18. Now, consider the case (1). Suppose without loss of generality that $x \in C$ and $y, z \in C'$. Since C and C' are separated, they belong to different blocks, thus $xy, xz \notin L_{ij}$. Since $y \prec z$ and both y, z belong to a cell of X^+_{ij}, we cannot have an arc $z \twoheadrightarrow y$. Thus $d_y \lesssim_4 d_z$. Also $d(x, y) \lesssim_{16} d(x, z)$, because $xy, xz \notin L_{ij}$. Now, suppose by way of contradiction that $d(y, z) \gg_{16} d(x, z)$. Since y and z belong to the same cell, we have $yz \in L_{ij}$. Then $d(y, z) \lesssim_3 \max\{d_y, d_z\}$, thus $d(x, z) \ll_{13} d_z$, establishing that $xz \in L_{ij}$. This shows that x and z belong to the same cluster, whence the algorithm should locate them in the same hole, contrary to $x \in C \subseteq X^-_{ij}$ and $z \in C' \subseteq X^+_{ij}$. □

4.4 Defining the Total Order on X_i

Recall that X_i denote the set of all elements of X° located in the hole H_i. According to our algorithm, X_i is the disjoint union of all sets X^-_{ij} $(j > i + 1)$ and of all sets

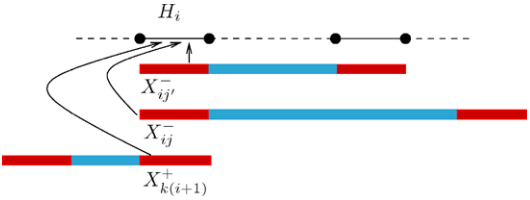

Fig. 3 Relative location of the cells of $X^+_{k(i+1)}$, $X^-_{ij'}$, and X^-_{ij} ($k < i$, $j' < j$) in the hole H_i

$X^+_{j'(i+1)}$ ($j' < i$), where some of these sets may be empty. In previous subsection, we defined a total order between the cells of each of the sets X^-_{ij}, $X^+_{j'(i+1)}$, and applying recursion we defined a total order on the elements of each cell. To obtain a total order on the whole set X_i it remains to define a total order between the sets X^-_{ij} ($j > i+1$) and $X^+_{j'(i+1)}$ ($j' < i$) and to extend it accordingly. First, we locate each set $X^+_{j'(i+1)}$ ($j' < i$) to the left of each set X^-_{ij} ($j > i$). Given two sets $X^+_{j'(i+1)}, X^+_{j''(i+1)}$ ($j', j'' < i$), we locate $X^+_{j'(i+1)}$ to the left of $X^+_{j''(i+1)}$ if and only if $j'' < j'$, i.e., iff $H_{j'(i+1)} \Subset H_{j''(i+1)}$. Analogously, given two sets $X^-_{ij}, X^-_{ij'}$ ($j, j' > i+1$), we locate $X^-_{ij'}$ to the right of X^-_{ij} if and only if $j' < j$, i.e., iff $H_{ij'} \Subset H_{ij}$. This location is justified by the Proposition 3.6 and is illustrated in Fig. 3.

5 The Algorithm

We have collected all necessary tools to describe the algorithm. It consists of three procedures l_∞-Fitting_by_Robinson, Refine, and Partition_and_Sort. The main procedure l_∞-Fitting_by_Robinson constructs the list Δ of feasible values for the optimal error ϵ^*. This list is sorted and its values are considered in a binary search fashion. The algorithm returns the smallest value $\epsilon \in \Delta$ occurring in this search for which a 16ϵ-compatible total order on X exists (i.e., the smallest ϵ for which the answer "not" is not returned). To decide, if for a given ϵ such an order exists, the procedure Refine$(X, \preccurlyeq, \epsilon)$ constructs (and/or updates) the canonical partial order $\preccurlyeq$ and computes a maximal chain $P = a_1 \preccurlyeq a_2 \preccurlyeq \cdots \preccurlyeq a_p$ of $(X, \preccurlyeq)$. For each element $x \in X^\circ := X \setminus P$, this procedure computes the set $AH(x)$ of all x-holes which participate in $(x, y, 1)$-admissible locations for all $y \in X^\circ$ and defines the segment $H(x)$. For each pair $i < j-1$, Refine constructs the set X_{ij} and makes a call of the procedure Partition_and_Sort(X_{ij}), which returns the bipartition $\{X^-_{ij}, X^+_{ij}\}$ of X_{ij} and two total orders $\preccurlyeq^-_{ij}$ and $\preccurlyeq^+_{ij}$ on the cells of X^-_{ij} and X^+_{ij}, respectively. Then Refine concatenates in a single total order on cells the total orders on cells coming from different sets assigned to the same hole. After this, Refine is recursively applied to each cell occurring in some graph $\overrightarrow{\mathcal{G}_{ij}}$. The returned total orders on cells are concatenated into a single total order $\prec$ on X according to the total orders between cells and between holes. The resulting total order is returned by the algorithm l_∞-Fitting_by_Robinson.

The procedure Partition_and_Sort constructs the graphs $\mathcal{L}_{ij}$ and $\overrightarrow{\mathcal{L}_{ij}}$. Using these graphs, the partition of X_{ij} into blocks and cells is derived. Further, the graph Ψ_{ij} and its clusters are constructed. Using the cells, the directed graph $\overrightarrow{\mathcal{G}_{ij}}$ is constructed. If

```
Algorithm $l_\infty$-Fitting_by_Robinson
Input: Finite set $X$ and dissimilarity $d$ on $X$
Output: The smallest $\epsilon$ such that $d$ is $16\epsilon$-Robinson and an $16\epsilon$-compatible order $\prec$ on $X$
Construct the list $\Delta = \{|\frac{d(x,y)-d(x',y')}{2}| : x, y, x', y' \in X\}$ and sort it
for each $\epsilon$ occurring in a binary search of the sorted list $\Delta$
    do Refine($X, \preccurlyeq, \epsilon$)
return $\tilde{d}_\prec$ and the total order $\prec$ for least tested $\epsilon$ for which Refine($X, \preccurlyeq, \epsilon$) $\neq$ "not"
```

```
Refine($X, \preccurlyeq, \epsilon$)
Input: A set $X$, $\epsilon \in \Delta$ and the canonical partial order $\preccurlyeq$
Output: The answer "not" or a $16\epsilon$-compatible total order $\prec$ on $X$
1.  if $\preccurlyeq$ is empty
2.      then set $p \preccurlyeq q$ for two arbitrary elements $p, q \in X$
3.          Update the betweenness set $\mathcal{B}$ and the canonical partial order $\preccurlyeq$ using the
            rules (B1)–(B4) and (PO1)–(PO6)
4.  if $\preccurlyeq$ is not a partial order
5.      then return "not"
6.  Find a maximal chain $P = a_1 \preccurlyeq a_2 \preccurlyeq \cdots \preccurlyeq a_p$ of $(X, \preccurlyeq)$
7.  for all $x \in X^\circ := X \setminus P$
8.      do compute $AH(x)$
9.      if $AH(x)$ consists of a single hole $H_i = [a_i, a_{i+1}]$
10.         then set $a_i \preccurlyeq x \preccurlyeq a_{i+1}$, $P := P \cup \{x\}$ and repeat steps 3–10
11. for all $i, j$ such that $1 \le i < j - 1 \le p - 1$
12.     do set $X_{ij} = \{x \in X \setminus P : a_i \preccurlyeq x \preccurlyeq a_j, x?a_{i+1} \text{ and } x?a_{j-1}\}$
13.     if Partition_and_Sort($X_{ij}$) = "not"
14.         then return "not"
15.         else set $\{X^-_{ij}, X^+_{ij}, \preccurlyeq^-_{ij}, \preccurlyeq^+_{ij}\} \Leftarrow$ Partition_and_Sort($X_{ij}$)
16.     for $i = 1 \ldots p - 1$
17.         do define on $X_i$ the total order
                $X^+_{1(i+1)} \prec X^+_{2(i+1)} \prec \cdots \prec X^+_{(i-1)(i+1)} \prec X^-_{ip} \prec \cdots \prec X^-_{i(i+2)}$        (1)
18.     for all $i = 1 \ldots p$, $j = 1 \ldots p$ and for each cell $C$ of $X^+_{ij}$ and for each cell $C'$ of $X^-_{ij}$
19.         do if Refine($C, \preccurlyeq^+_{ij}, \epsilon$) = "not" or Refine($C', \preccurlyeq^-_{ij}, \epsilon$) = "not"
20.             then return "not"
21.             else return the total order $\prec$ on $X$ which on each cell $C \in X^+_{ij}$ coincides with
                Refine($C, \preccurlyeq^+_{ij}$), on each cell $C' \in X^-_{ij}$ coincides with Refine($C', \preccurlyeq^-_{ij}$), on each set
                $X^-_{ij}$ coincides with (2), on each set $X^+_{ij}$ coincides with (3), and on each set
                $X_i$ coincides with (1)
```

Ψ_{ij} is not bipartite or $\overrightarrow{\mathcal{G}_{ij}}$ contains (G3)-cycles, the procedure Partition_and_Sort returns the answer "not". Otherwise, for each cell C and each cluster $\mathcal{K}$, it tests if there exists a 1-cycle and/or a 2-cycle passing via C and intersecting $\mathcal{K}$. Consequently, for each cell C, the lists $\Omega_1(C)$, $\Omega_2(C)$, and $\Omega_3(C)$ of (G2)-arcs are computed. These lists are employed to construct the 2-SAT formula Φ_{ij}, which is solved by the algorithm of [4]. If Φ_{ij} admits a true assignment A, then $X^-_{ij} = \{C : A(x_C) = 0\}$ and $X^+_{ij} = \{C : A(x_C) = 1\}$ define a bipartition of X_{ij} into two acyclic subgraphs $\overrightarrow{\mathcal{G}^-_{ij}}$, $\overrightarrow{\mathcal{G}^+_{ij}}$ of $\overrightarrow{\mathcal{G}_{ij}}$. Then Partition_and_Sort locates the cells from X^+_{ij} in the bounding hole H_{j-1}

```
Partition_and_Sort($X_{ij}$)
Input: A set $X_{ij} = \{x : H(x) = [a_i, a_j]\}$
Output: A bipartition $\{X^-_{ij}, X^+_{ij}\}$ of $X_{ij}$ and two total orders $\preccurlyeq^-_{ij}$ and $\preccurlyeq^+_{ij}$ on the cells of $X^-_{ij}$
        and $X^+_{ij}$, respectively
Construct the graphs $\mathcal{L}_{ij}$, $\overrightarrow{\mathcal{L}_{ij}}$, and $\Psi_{ij}$ using the rules (L1)–(L2) and (S1)–(S2)
if $\Psi_{ij}$ is not bipartite
    then return "not"
Construct the cells of $\overrightarrow{\mathcal{L}_{ij}}$ and the clusters of $\Psi_{ij}$
Construct the directed graph $\overrightarrow{\mathcal{G}_{ij}}$ using the rules (G1)–(G3)
if $\overrightarrow{\mathcal{G}_{ij}}$ contains a directed (G3)-cycle
    then return "not"
For each cell $C$ construct the lists $\Omega_1(C)$, $\Omega_2(C)$, and $\Omega_3(C)$ of (G2)-arcs
Construct the 2-SAT formula $\Phi_{ij}$ using the rules (F1)–(F3)
if $\Phi_{ij}$ is not satisfiable
    then return "not"
    else if $A$ is a true assignment of $\Phi_{ij}$
        Set $X^-_{ij} = \{C : A(x_C) = 0\}$, $X^+_{ij} = \{C : A(x_C) = 1\}$
Find a topological order $C_{j_1} \prec C_{j_2} \prec \cdots \prec C_{j_p}$ on the cells of $X^+_{ij}$
Find a dual topological order $C_{i_q} \prec C_{i_{q-1}} \prec \cdots \prec C_{i_1}$ on the cells of $X^-_{ij}$
Set $a_i \prec C_{i_q} \prec C_{i_{q-1}} \prec \cdots \prec C_{i_1} \prec a_{i+1}$                    (2)
Set $a_{j-1} \prec C_{j_1} \prec C_{j_2} \prec \cdots \prec C_{j_p} \prec a_j$                        (3)
For each cell $C$ of $X^+_{ij}$ and all $x, y \in C$ such that $x \twoheadrightarrow y$ set $x \preccurlyeq^+_{ij} y$, and update $\preccurlyeq^+_{ij}$
    using the rules (PO1)–(PO6)
For each cell $C$ of $X^-_{ij}$ and all $x, y \in C$ such that $x \twoheadrightarrow y$ set $y \preccurlyeq^-_{ij} x$, and update $\preccurlyeq^-_{ij}$
    using the rules (PO1)–(PO6)
return $X^+_{ij}$, $X^-_{ij}$ and $\preccurlyeq^-_{ij}$, $\preccurlyeq^+_{ij}$
```

according to the topological order of the acyclic graph $\overrightarrow{\mathcal{G}^+_{ij}}$. Analogously, it locates the cells from X^-_{ij} according to the dual topological order of the acyclic graph $\overrightarrow{\mathcal{G}^-_{ij}}$. Note that if at some stage one of the procedures Refine or Partition_and_Sort returns the answer "not", then there does not exist any ϵ-compatible total order on X and the current value of ϵ is too small.

It can be easily shown that the algorithm is polynomial. Indeed, the construction of the list Δ, of the canonical partial order $\preccurlyeq$, and of the sets $AH(x), x \in X^\circ$ can be done in $O(n^4)$ time each (recall that $n = |X|$). Let $n_{ij} = |X_{ij}|$. The graph $\mathcal{L}_{ij}$ is constructible in $O(n^2_{ij})$ time and its blocks can be derived within the same time bounds. Implementing the conditions (A1)–(A2) and (L1)–(L2) in the most direct way, the second graph $\overrightarrow{\mathcal{L}_{ij}}$ can be constructed in $O(n^3_{ij})$ time. Since the cells are the strongly connected components of $\overrightarrow{\mathcal{L}_{ij}}$, they can be computed in $O(n^2_{ij})$ time. The graph Ψ_{ij} can be constructed in time proportional to the square of the number of blocks (see the conditions (S1)–(S2)), i.e., in $O(n^2_{ij})$ time. Finally, the graph $\overrightarrow{\mathcal{G}_{ij}}$ can be constructed in $O(n^3_{ij})$ time because of condition (G3) (the conditions (G1) and (G2) can be implemented in $O(n^2_{ij})$ time). That $\overrightarrow{\mathcal{G}_{ij}}$ does not contain directed (G3)-cycles can be tested in linear time in the size of this graph. Since the clusters define a partition of X_{ij} and finding a shortest path between two vertices in an acyclic graph can be done in linear time in the size of the graph, for a cell C, we can compute the lists $\Omega_1(C)$, $\Omega_2(C)$, and $\Omega_3(C)$ in total $O(n_{ij}^4)$ time. Consequently, the 2-SAT formula Φ_{ij} can be constructed in $O(n_{ij}^5)$ time and its length is $O(n_{ij}^2)$. A topological ordering of an acyclic graph can be found in linear time in the number of arcs. We conclude that the complexity of the procedure Partition_and_Sort on set X_{ij} is $O(n_{ij}^5)$. Since the sets X_{ij} define a partition of the set X°, we have $\sum_{i<j} n_{ij} < n$, thus the overall complexity of Partition_and_Sort on all sets of this partition is $O(n^5)$. Now, taking into account the recursive calls to the cells and the fact that the cells also define a partition of X°, we conclude that the complexity of the procedures Refine and Partition_and_Sort for a given ϵ is $O(n^6)$ (we skipped the analysis of several easy low-complexity steps of both procedures). Now, if we run a binary search on the list Δ of $O(n^4)$ possible values of ϵ, we deduce that the complexity of the algorithm l_∞-Fitting_by_Robinson is $O(n^6 \log n)$.

6 Performance Guarantee

We will establish here that the algorithm described above is a constant factor approximation algorithm for the problem l_∞-FITTING-BY-ROBINSON.

Theorem 6.1 *For a given $\epsilon \in \Delta$, if the algorithm returns the answer "not", then the dissimilarity d is not ϵ-Robinson, else if the algorithm returns a total order $\prec$ on X, then $\prec$ is 16ϵ-compatible with d. In particular, the algorithm is a factor 16 approximation algorithm for the problem l_∞-*FITTING-BY-ROBINSON.

Proof First we investigate the cases when the algorithm returns the answer "not". We will show that in all such cases, no ϵ-admissible order exist. This conclusion holds if some separability graph Ψ_{ij} is not bipartite, because for any location of X_{ij} in the holes of H_{ij}, two adjacent in Ψ_{ij} cells will be placed in the same bounding hole. If some directed graph $\overrightarrow{\mathcal{G}_{ij}}$ contain a (G3)-cycle, then the negative answer is justified by Lemma 4.11. Analogously, if some 2-SAT formula Φ_{ij} is not satisfiable, then the answer "not" is confirmed by Proposition 4.16. If the canonical order $\preccurlyeq$ on X constructed by the procedure Refine is not a partial order, then clearly no ϵ-compatible order exist because any such total order is an extension of $\preccurlyeq$. Finally, suppose that the answer "not" is returned at the recursive call of the algorithm to a cell C and to one of the partial orders $\preccurlyeq^+$ or $\preccurlyeq^-$, say the first. By the induction assumption, no ϵ-compatible total order on C extending $\preccurlyeq^+$ (and therefore its dual $\preccurlyeq^-$) exist. But then, from Lemma 4.19 we infer that no ϵ-compatible order on X exist as well.

Now suppose that the algorithm returns a total order $\prec$. Suppose by induction assumption that $\prec$ is 16ϵ-compatible on each cell to which we perform a recursive call. On the chain P, the total order $\prec$ coincides with $\preccurlyeq$, therefore $\prec$ is ϵ-compatible on P. Moreover, $\prec$ is ϵ-compatible on $P \cup \{x\}$ for any $x \in X^\circ$, because every element x is located in a bounding hole of $H(x)$ which is x-admissible. Finally notice that $\prec$ is 12ϵ-compatible on $P \cup \{x, y\}$ for any $x, y \in X^\circ$ because by the results of Sect. 3 any bounding hole of $H(x)$ and any bounding hole of $H(y)$ define a $(x, y, 12)$-admissible pair (see Proposition 3.6). In order to establish that $\prec$ is 16ϵ-compatible on the whole set X, by Lemma 2.2 it suffices to show that $d(x,z) \gtrsim_{16} \max\{d(x,y), d(y,z)\}$ for any three elements $x, y, z \in X$ such that $x \prec y \prec z$. From previous discussion we conclude that it suffices to suppose that $x, y, z \in X^\circ$. We continue with several easy but helpful properties.

Claim 1: *If* $x, y \preccurlyeq z$ *and* $x?y$ *or* $x \preccurlyeq y$*, then* $d(x,z) \gtrsim_2 \max\{d(x,y), d(y,z)\}$*. If* $y \preccurlyeq x \preccurlyeq z$ *and* $(I^\circ(x) \cap I^\circ(y)) \setminus I^\circ(z) \neq \emptyset$*, then* $d(x,y) \gtrsim_6 \max\{d(x,y), d(y,z)\}$*.*

Proof of Claim 1 Indeed, if $x \preccurlyeq y \preccurlyeq z$, then $d(x,z) \gtrsim_2 \max\{d(x,y), d(y,z)\}$ because the partial order $\preccurlyeq$ is ϵ-compatible. On the other hand, if $x?y$, then from Lemma 2.7 we infer that $d(x,z) \approx_2 d(y,z)$ and $d(x,y) \lesssim_2 \min\{d(x,z), d(y,z)\}$, and we are done. Finally, if $a \in (I^\circ(x) \cap I^\circ(y)) \setminus I^\circ(z)$, then we have $x, y, a \preccurlyeq z$ and $x?a$, $y?a$. Again, from Lemma 2.7 we infer that $d(x,z) \approx_2 d(a,z)$ and $d(y,z) \approx_2 d(a,z)$, i.e., $d(x,z) \approx_4 d(y,z)$. On the other hand, since $y \preccurlyeq x \preccurlyeq z$, we deduce that $d(y,z) \gtrsim_2 d(x,y)$, thus $d(x,z) \gtrsim_4 d(y,z)$, yielding $d(x,z) \gtrsim_6 d(x,y)$. This establishes Claim 1.

Claim 2: *If* $d(x,z) \gtrsim_c \max\{d_x, d_z\}$*,* $H(y) \circ H(x)$ *and* $H(y) \circ H(z)$*, then we have* $d(x,z) \gtrsim_{(c+3)} \max\{d(x,y), d(y,z)\}$*. In particular, if the elements* x *and* z *are not linked, then* $d(x,z) \gtrsim_{10} \max\{d(x,y), d(y,z)\}$*.*

Proof of Claim 2 Since $H(y) \circ H(x)$ and $H(y) \circ H(z)$, Lemma 3.12 implies that $d(x,y) \approx_3 d_x, d_y$ and $d(y,z) \approx_3 d_y, d_z$. Hence $d(x,z) \gtrsim_c d_x$, yielding $d(x,z) \gtrsim_{(c+3)} d(x,y)$. Analogously $d(x,z) \gtrsim_c d_z$, yielding $d(x,z) \gtrsim_{(c+3)} d(y,z)$. If x and z are not linked, then from the results of Sect. 3 we conclude that $d(x,z) \gtrsim_7 \max\{d_x, d_z\}$, and we can apply the first assertion with $c = 7$. This establishes Claim 2.

Claim 3: *If* $d(x,z) \gtrsim_c \max\{d_x, d_z\}$ *and either* $H(x) \Subset H(y)$*, or* $H(y) \circ H(z)$*, or* $H(y) \circ H(x)$ *and* $H(z) \Subset H(y)$ *holds, then* $d(x,z) \gtrsim_{(c+9)} \max\{d(x,y), d(y,z)\}$*. In particular, if the elements* x *and* y *are not linked, then* $d(x,z) \gtrsim_{16} \max\{d(x,y), d(y,z)\}$*.*

Proof of Claim 3 Suppose that $H(x) \Subset H(y)$ and $H(y) \circ H(z)$ (the another case is similar). Then $d(y,z) \approx_3 d_y, d_z$ by Lemma 3.12. On the other hand, from Lemma 3.16 we infer that $d(x,y) \approx_3 d_y$, whence $d(y,z) \approx_6 d(x,y)$. In conclusion, $d(x,z) \gtrsim_{(c+3)} d(y,z) - (c+3)$ and $d(x,z) \gtrsim_{(c+9)} d(x,y)$. Analogously to Claim 2, if x and z are not linked, then we can apply the first assertion of Claim 3 with $c = 7$.

Claim 4: *If* $d(x,z) \gtrsim_c d(y,z)$ *and one of the bounding holes of* $H(z)$ *is to the left or to the right of* $H(x) \cup H(y)$*, then* $d(x,z) \gtrsim (c+2) d(x,y)$*.*

Proof of Claim 4 Notice that if $d(x,y) \gg_2 \max\{d(x,z), d(z,y)\}$, then $z \in [x,y]$ in the canonical betweenness. Therefore in all ϵ-compatible orders z must be located between x and y. Hence, if the rightmost hole of $H(z)$ is to the right of $H(x) \cup H(y)$, then this hole cannot be z-admissible, which is impossible. Thus

$d(x,y) \lesssim_2 \max\{d(x,z), d(z,y)\} \lesssim_c d(x,z)$ and therefore $d(x,z) \gtrsim_{(c+2)} d(x,y)$. This establishes Claim 4.

We distinguish 6 cases in function of the mutual location of the segments $H(x)$ and $H(z)$. In Cases 1 and 2, we will assume without loss of generality that, if x and z are comparable in the canonical partial order $\preccurlyeq$, then always $x \preccurlyeq z$ holds.

Case 1: $H(x) \cap H(z) = \emptyset$.

Proof of Case 1 If $H(y)$ is disjoint from $H(x)$ and $H(z)$, then $H(y)$ lies between $H(x)$ and $H(z)$. Consequently, $x \preccurlyeq y \preccurlyeq z$, and this case is covered by Claim 1. Now assume that $H(y)$ intersects only one of the segments $H(x)$ and $H(z)$, say $H(x)$. Then $y \preccurlyeq z$ and $x \preccurlyeq z$. We can apply Claim 1, unless $y \preccurlyeq x$ and $I^\circ(x) \cap I^\circ(y) = \emptyset$. But this case occurs only when $H(x)$ and $H(y)$ have exactly one hole in common. Since $y \preccurlyeq x$, this hole is necessarily the rightmost hole of $H(y)$ (and respectively, the leftmost hole of $H(x)$). According to the algorithm, y will be placed to the left of x, contrary to the assumption that $x \prec y$.

Now suppose that $H(y)$ intersects both segments $H(x)$ and $H(z)$. By Lemma 3.9 we have $d(x,z) \gtrsim_5 \max\{d_x, d_z\}$, therefore we can use the Claims 2 and 3 with $c = 5$. In particular, if $H(y) \circ H(x)$ and $H(z) \circ H(y)$, by Claim 2 we obtain $d(x,z) \gtrsim_8 \max\{d(x,y), d(y,z)\}$. On the other hand, if $H(x) \circ H(y)$, $H(z) \Subset H(y)$ or $H(x) \Subset H(y)$, $H(y) \circ H(z)$, from Claim 3 we infer that $d(x,z) \gtrsim_{14} \max\{d(x,y), d(y,z)\}$. If $H(x) \Subset H(y)$ and $H(z) \Subset H(y)$, then the constraint $x \prec y \prec z$ and the fact that y is placed by the algorithm in a bounding hole of $H(y)$, we conclude that one of the elements x, z, say x, is located by the algorithm in the same hole as y. Necessarily, this is the leftmost hole of $H(x)$ and $H(y)$. Since $H(x) \Subset H(y)$, the algorithm will locate y to the left of x, contrary to the assumption that $x \prec y$.

Case 2: $H(x) * H(z)$.

Proof of Case 2 Suppose $H(x) \cap H(z) = H_i$. Since the algorithm locate x to the left of z, necessarily H_i is the rightmost hole of $H(x)$ and the leftmost hole of $H(z)$. Then either $x?z$ or $x \preccurlyeq z$. From Lemma 3.14 we know that $d(x,z) \approx_4 \delta_i$ and $d(x,z) \gtrsim_3 \max\{d_x, d_z\}$.

First suppose that $H(y)$ intersects only one of the segments $H(x)$ and $H(z)$, say $H(y) \cap H(z) = \emptyset$. Thus $y \preccurlyeq z$. If $x \preccurlyeq z$, then we can use Claim 1 except the case when $y \preccurlyeq x$ and $I^\circ(x) \cap I^\circ(y) = \emptyset$ (notice that $I^\circ(x) \cap I^\circ(y)$ is always disjoint from $I^\circ(z)$ because $H(y) \cap H(z) = \emptyset$). As we noticed already, this is possible only if $H(y) * H(x)$, namely if $H(y) \cap H(x)$ is the rightmost hole of $H(y)$ (and the leftmost hole of $H(x)$). But then the algorithm will locate y to the left of x, contrary to $x \prec y$. Therefore, further we can assume that $x?z$ and that either $H(x) \circ H(y)$ or $H(y) \Subset H(x)$. Since $x?z$, Lemma 3.14 implies that $d(x,z) \approx_3 d_x, d_z$. On the other hand, since $H(x) \circ H(y)$ or $H(y) \Subset H(x)$, from Lemmata 3.12 and 3.16 we infer that $d(x,y) \approx_3 d_x$, yielding $d(x,z) \gtrsim_6 d(x,y)$. It remains to compare $d(x,z)$ to $d(y,z)$. Pick an element $a_j \in I^\circ(x) \cap I^\circ(y)$. Since $y?a_j$ and $y, a_j \preccurlyeq z$, we conclude that $d(y,z) \approx_2 d(a_j, z)$. On the other hand, since the location of x in the leftmost hole of $H(x)$ and the location of z in the rightmost hole of $H(z)$ is $(x, z, 6)$-admissible by Lemma 3.14 and a_j is between x and z, we conclude that $d(x,z) \gtrsim_6 d(a_j, z) \gtrsim_2 d(y,z)$. Thus, if $H(y) \cap H(z) = \emptyset$ and $H(y) \cap H(x) \neq \emptyset$, then $d(x,z) \gtrsim_8 \max\{d(x,y), d(y,z)\}$.

Now suppose that $H(y)$ intersects both segments $H(x)$ and $H(z)$, i.e., $H(y)$ contains the hole $H_i = H(x) \cap H(z)$. Since $x \prec y \prec z$ and y is located in a bounding hole of $H(y)$, necessarily y is located in a hole of $H(x) \cup H(z)$. Suppose without loss of generality that y is located in a hole H of $H(x)$. Then either H is the leftmost hole of $H(y)$ or $H = H_i$ is the rightmost hole of $H(y)$. Recall that $d(x,z) \gtrsim_3 \max\{d_x, d_z\}$. Therefore, if $H(x) \circ H(y)$ and $H(y) \circ H(z)$, then we can use Claim 2 with $c = 3$. Analogously, if $H(x) \Subset H(y)$ and $H(y) \circ H(z)$ or if $H(z) \Subset H(y)$ and $H(y) \circ H(x)$, the result follows from Claim 3 with $c = 3$. Now, if $H(y)$ contains both $H(x)$ and $H(y)$, the algorithm will locate y either to the left of x or to the right of z, contrary to $x \prec y \prec z$. It remains to consider the cases when either $H(y) * H(z)$ or $H(y) * H(x)$. If $H(y) * H(z)$, by Lemma 3.14 we infer that $d(y,z) \approx_4 \delta_i$. Hence $d(x,z) \approx_8 d(y,z)$, i.e. $d(x,z) \gtrsim_8 d(y,z)$. Now, if $d(x,y) \gg_2 \max\{d(x,z), d(z,y)\}$, then $z \in [x,y]$ in the canonical betweenness. Since $H(z) \cap H(x) = H(z) \cap H(y) = H_i$, the element z must be located in the hole H_i contrary to the assumption that $z \in X^\circ$. Thus $d(x,y) \lesssim_2 \max\{d(x,z), d(z,y)\} \lesssim_8 d(x,z)$. Hence $d(x,z) \gtrsim_{10} d(x,y)$. The case $H(y) * H(x)$ is completely analogous. The analysis of case $H(x) * H(z)$ is now complete.

Case 3: $H(x) \circ H(z)$.

Proof of Case 3 Then $d(x,z) \approx_3 d_x, d_z$ in view of Lemma 3.12. First suppose that $H(y)$ intersects both segments $H(x)$ and $H(z)$. If $H(x) \circ H(y)$ and $H(y) \circ H(z)$, then we can use Claim 2 with $c = 3$. Analogously, if $H(x) \circ H(y)$, $H(z) \Subset H(y)$ or $H(x) \Subset H(y)$, $H(y) \circ H(z)$, then we can use Claim 3 with $c = 3$. If $H(y)$ coincides with one of the segments, say with $H(x)$, then $H(y) \circ H(z)$, therefore $d(y,z) \approx_3 d_y, d_z$, and we infer that $d(x,z) \approx_6 d(y,z)$. Since the rightmost hole of $H(z)$ does not belong to $H(x) = H(y)$, the result follows from Claim 4 with $c = 6$. If $H(x) \Subset H(y)$ and $H(z) \Subset H(y)$, then the algorithm will locate y either to the left of x or to the right of z, contrary to $x \prec y \prec z$. Finally, if $H(y) \Subset H(z)$, then $d(y,z) \approx_3 d_z$ by Lemma 3.16, yielding $d(x,z) \approx_6 d(y,z)$. Notice that in this case either $H(y)$ is properly contained in $H(x)$ or $H(y)$ overlaps $H(x)$. In both cases, the leftmost hole of $H(x)$ is located to the left of $H(y) \cup H(z)$, thus we can use Claim 4 with $c = 6$.

Now suppose that $H(y)$ intersects only one of the segments $H(x)$ and $H(z)$, say $H(y) \cap H(z) = \emptyset$. Then $y \preccurlyeq z$. If $H(y) * H(x)$, then they share the leftmost hole of $H(x)$, therefore the algorithm will locate y to the left of x, contrary with $x \prec y$. On the other hand, if $H(y) \Subset H(x)$ or $H(y) \circ H(x)$, then we conclude that $d(x,y) \approx_3 d_x$, and therefore $d(x,z) \approx_6 d(x,y)$. If $d(y,z) \lesssim_2 \max\{d(x,y), d(x,z)\}$, then we are done because $d(x,z) \gtrsim_8 d(y,z)$. Otherwise, if $d(y,z) \gg_2 \max\{d(x,y), d(x,z)\}$, then $x \in [y,z]$. Since $y \preccurlyeq z$, we obtain that $y \preccurlyeq x \preccurlyeq z$ and the result follows from the second assertion of Claim 1 because obviously $I^\circ(x) \cap I^\circ(y) \neq \emptyset$. This concludes the analysis of case $H(x) \circ H(z)$.

Case 4: $H(x) \Subset H(z)$.

Proof of Case 4 According to Lemma 3.16, we have $d(x,z) \approx_3 d_z$ and $d(x,z) \gtrsim_7 d_x$. Therefore, if $H(y) \circ H(x)$, $H(y) \circ H(z)$ or $H(x) \Subset H(y)$, $H(y) \circ H(z)$ we can use Claims 2 and 3 with $c = 7$. If $H(z) \Subset H(y)$, then the algorithm will locate y either to the left of x or to the right of z, contrary to $x \prec y \prec z$. A similar contradiction is obtained if $H(x) * H(y)$ and $H(x) \cap H(y)$ is the leftmost hole of $H(x)$. If $H(x) \Subset H(y) \Subset H(z)$, then $d(x,y) \approx_3 d_y$ and $d(y,z) \approx_3 d_z$ by Lemma 3.16. Since $d(x,z) \approx_3 d_z$ and $d(y,z) \gtrsim_7 d_y$, we conclude that $d(x,z) \approx_6 d(y,z)$ and $d(x,z) \gtrsim_6 d(y,z) \gtrsim_7 d_y \gtrsim_3 d(x,y)$, yielding $d(x,z) \gtrsim_{16} d(x,y)$. A similar proof holds when $H(y) \Subset H(x)$. Now suppose that $H(x) * H(y)$ and $H_i = H(x) \cap H(y)$ is the rightmost hole of $H(x)$. By Lemma 3.14, $d(x,y) \approx_4 \delta_i$. If $H(y) \Subset H(z)$ or $H(y) \circ H(z)$, then $d(y,z) \approx_3 d_z$, thus $d(x,z) \approx_6 d(y,z)$. On the other hand, H_i is an inner hole of $H(z)$, therefore $\delta_i \lesssim_3 d_z$ by Lemma 3.3, thus $d(x,y) \lesssim_7 d_z$, i.e., $d(x,z) \approx_{10} d(x,y)$. So, assume that $H(y) * H(z)$, i.e. $H(x)$ and $H(z)$ have the same rightmost hole H_i which is the leftmost hole of $H(y)$. But in this case the algorithm will locate y to the right of x and z, contrary to $y \prec z$.

If $H(x) = H(y)$, then $H(y) \Subset H(z)$ and $d(y,z) \approx_3 d_z$ by Lemma 3.16, yielding $d(x,z) \approx_6 d(y,z)$. Then we can use Claim 4 with $c = 6$, because one of the bounding holes of $H(z)$ is necessarily located outside the segment $H(x) = H(y)$. Now suppose that $H(y) = H(z)$. If the algorithm locates x in the leftmost hole of $H(x)$, since $x \prec y$, necessarily y is located in the rightmost hole of $H(y)$, yielding that z is also located in that hole. The analogous conclusion holds if x is located in the rightmost hole of $H(x)$. Since in both cases y and z are located in a common bounding hole, these elements are not separated, i.e., $d(y,z) \lesssim_3 \max\{d_y, d_z\}$. By Lemma 3.16, $d(x,z) \approx_3 d_z$ and $d(x,y) \approx_3 d_y$. If $d_y \lesssim_4 d_z$, then the result is immediate because $d(y,z) \lesssim_7 d_z \lesssim_3 d(x,z)$ and $d(x,y) \lesssim_3 d_y \lesssim_4 d_z \lesssim_3 d(x,z)$. However, if $d_y \gg_4 d_z$, then the algorithm will place y to the right of z, contrary to $y \prec z$.

Finally, consider the case when $H(x)$ and $H(y)$ are disjoint, i.e., $H(y)$ is to the right of $H(x)$. Then $d(x,z) \approx_3 d_z$ and $d(y,z) \approx_3 d_z$ by Lemma 3.16. By Lemma 3.10 $d(x,y) \approx_8 d(a_i, a_j)$, where $a_i \in I^\circ(x)$ and $a_j \in I^\circ(y)$. Since $a_i, a_j \in I^\circ(z)$, from Lemma 3.3 we infer that $d(a_i, a_j) \lesssim_3 d_z$, thus $d(x,y) \lesssim_8 d(a_i, a_j) \lesssim_3 d_z \lesssim_3 d(x,z)$, yielding $d(x,y) \lesssim_{14} d(x,z)$. This completes the analysis of Case $H(x) \Subset H(z)$.

Case 5: $H(z) \Subset H(x)$.

Proof of Case 5 Analogous to Case 4.

Case 6: $H(x) = H(z)$.

Proof of Case 6 Let $H(x) = H_{ij} = H(z)$. If $H(y) * H_{ij}$ or $H(y) \cap H_{ij} = \emptyset$, then the algorithm will locate y either to the left of x or to the right of z, contrary to $x \prec y \prec z$. If $H(y) \circ H_{ij}$, then from Lemma 3.12 we infer that $d(x,y) \approx_3 d_x, d_y$ and $d(y,z) \approx d_y, d_z$. If the elements x and z are not linked, then $d(x,z) \gtrsim_3 \max\{d_x, d_z\}$, thus $d(x,z) \gtrsim_6 d(x,y)$ and $d(x,z) \gtrsim_6 d(y,z)$. On the other hand, if x and z are linked, then all x and z are located in a common bounding hole of H_{ij}, say in the hole

H_i. But then y must be located in H_i as well because $x \prec y \prec z$. This is impossible because H_i can be only an inner hole of $H(y)$ and the algorithm locates y in a bounding hole of $H(y)$. Analogously, if $H(y) \Subset H_{ij}$, then $d(x, y) \approx_3 d_x$ and $d(y, z) \approx_3 d_z$. If x and z are not linked, then the proof is similar to the case $H_{ij} \circ H(y)$. If x and z are linked, then they will be placed by the algorithm in a common bounding hole, say in H_i. Since $x \prec y \prec z$, H_i is the leftmost bounding hole of $H(y)$ and y must be placed in H_i. However, since $H(y) \Subset H(x)$ and $H(y) \Subset H(z)$, the algorithm locates x and z to the left of y, contrary to our assumption $x \prec y \prec z$. Analogously, if $H_{ij} \Subset H(y)$, since $x \prec y \prec z$ and all three elements x, y, z are located in their bounding holes, we conclude that H_{ij} and $H(y)$ must share a common bounding hole, say the hole H_i, into which both x and y are located. Since $H(x) \Subset H(y)$, the algorithm must locate y to the left of x, contrary to $x \prec y$.

It remains to consider the case $H(y) = H_{ij}$, i.e., $x, y, z \in X_{ij}$. If x, y, and z belong to the same cell C, then we can apply the induction hypothesis to C. Analogously, if x, y, and z belong to three distinct cells of X_{ij}, then the result follows from Lemma 4.18. The remaining case, in which the elements x, y, z belong to two cells of X_{ij} is covered by Lemma 4.20. This settles the analysis of Case 6 and conclude the proof of the theorem. □

7 Approximating Minimum Additive Relaxation

In this section, we prove that the minimum l_∞-error ϵ^* and the minimum additive relaxation α^* for ordinal embedding of a dissimilarity d on X into the cone $\mathcal{R}$ of all Robinsonian dissimilarities on X satisfy the equality $\alpha^* = 2\epsilon^*$. As a consequence, the Robinsonian dissimilarity returned by the algorithm l_∞-Fitting_by_Robinson provides an ordinal embedding of d into $\mathcal{R}$ with additive relaxation at most $16\alpha^*$.

First notice that for any target set $\mathcal{D}'$ of dissimilarities, we have $\alpha^* \leq 2\epsilon^*$. Pick $d' \in \mathcal{D}'$ such that $\|d - d'\|_\infty = \epsilon^*$ and any elements $x, y, u, v \in X$ such that $d(x, y) + 2\epsilon^* < d(u, v)$. If $d'(x, y) \geq d'(u, v)$, then $d(x, y) + \epsilon^* \geq d'(x, y) \geq d'(u, v) \geq d(u, v) - \epsilon^*$, whence $d(x, y) + 2\epsilon^* \geq d(u, v)$, and we obtain a contradiction. Thus d' provides an ordinal embedding of d into $\mathcal{D}'$ with additive relaxation at most $2\epsilon^*$.

To establish the converse inequality $\alpha^* \geq 2\epsilon^*$ for $\mathcal{R}$, pick an arbitrary total order $\prec$ on X and denote by $\tilde{\alpha}_\prec$ the minimum additive relaxation of an ordinal embedding of d into $\mathcal{R}_\prec$. We assert that $\tilde{\alpha}_\prec \geq 2\tilde{\epsilon}_\prec = \|d - \check{d}_\prec\|_\infty$, where $\check{d}_\prec$ is the super-dominant of d in $\mathcal{R}_\prec$ defined in Lemma 2.3. Suppose by way of contradiction that $\tilde{\alpha}_\prec < 2\tilde{\epsilon}_\prec$ and let $d'_\prec$ be a Robinsonian dissimilarity providing an ordinal embedding of d into $\mathcal{R}_\prec$ with relaxation $\tilde{\alpha}_\prec$. Pick the elements $x \prec u \prec v \prec y$ such that $d(u, v) - d(x, y) = 2\tilde{\epsilon}_\prec$. Since $\tilde{\alpha}_\prec < 2\tilde{\epsilon}_\prec$, we conclude that $d(x, y) + \tilde{\alpha}_\prec < d(u, v)$, thus $d'_\prec(x, y) < d'_\prec(u, v)$ by the definition of an ordinal embedding. However, since $x \prec u \prec v \prec y$, we obtain a contradiction with the assumption that $d'_\prec \in \mathcal{R}_\prec$. This shows that $\tilde{\alpha}_\prec \geq 2\tilde{\epsilon}_\prec$. From Lemma 2.4 and its proof we conclude that $\tilde{d}_\prec$ is a Robinsonian dissimilarity that minimizes $\|d - d'\|_\infty$ for $d' \in \mathcal{R}_\prec$ and that $\tilde{\epsilon}_\prec = \|d - \tilde{d}_\prec\|_\infty$. Setting $\mathcal{D}' := \mathcal{R}_\prec$ in the first part of the proof, we deduce that $\tilde{d}_\prec$ provides an ordinal embedding of d into $\mathcal{R}_\prec$ with relaxation at most $2\tilde{\epsilon}_\prec$. Hence $\tilde{\alpha}_\prec = 2\tilde{\epsilon}_\prec$. Since this equality holds for any total order on X, we conclude that $\alpha^* = 2\epsilon^*$. Finally, since our algorithm l_∞-Fitting_by_Robinson returns a total order $\prec$ and the Robinsonian dissimilarity $\tilde{d}_\prec$ which is $\tilde{\epsilon}_\prec$-compatible with $\prec$ such that $\tilde{\epsilon}_\prec \leq 16\epsilon^*$, from our proof we conclude that $\tilde{d}_\prec$ provides an ordinal embedding of d into $\mathcal{R}$ with additive relaxation $2\tilde{\epsilon}_\prec \leq 32\epsilon^* = 16\alpha^*$. Summarizing, we obtain the following result:

Proposition 7.1 *For a dissimilarity d on X and the set $\mathcal{R}$ of all Robinsonian dissimilarities on X the equality $\alpha^* = 2\epsilon^*$ holds. The algorithm l_∞-*Fitting_by_Robinson *returns a Robinsonian dissimilarity $\tilde{d}_\prec$ which provides an ordinal embedding of d into $\mathcal{R}$ with additive relaxation at most* $16\alpha^*$. *In particular, the algorithm is a factor* 16 *approximation algorithm for the* ORDINAL ADDITIVE RELAXATION *problem in the case of Robinsonian dissimilarities.*

8 Glossary

For reader convenience, we recall here the main terms used in this paper.

Robinsonian dissimilarity: A dissimilarity d on a set X which has a total order $\prec$ (called *compatible order*) such that $x \prec z \prec y$ implies that $d(x, y) \geq \max\{d(x, z), d(z, y)\}$.

ϵ-Compatible order: a total order $\prec$ on X such that $u \prec x \prec y \prec v$ implies $d(u, v) + 2\epsilon \geq d(x, y)$.

ϵ-Robinsonian dissimilarity: a dissimilarity d on X admitting an ϵ-compatible order.

Problem l_∞-FITTING-BY-ROBINSON: Given a dissimilarity d on X find a Robinson dissimilarity d_R on X minimizing the l_∞-error $\epsilon^* = \|d - d_R\|_\infty$.

Compact list of breakpoints: $\Delta = \{|d(x, y) - d(x', y')|/2 : x, y, x', y' \in X\}$.

Canonical partial order $\preccurlyeq$: a partial order such that any ϵ-compatible order refines $\preccurlyeq$ or its dual.

Path $P = (a_1, a_2, \ldots, a_{p-1}, a_p)$: a maximal by inclusion path of the canonical partial order.

Bounding holes: the holes $[a_i, a_{i+1}]$ and $[a_{j-1}, a_j]$ of $[a_i, a_j]$.

x-Admissible hole: a hole H_i such that the total order on $P \cup \{x\}$ obtained by adding $a_i \preccurlyeq x \preccurlyeq a_{i+1}$ to $\preccurlyeq$ is ϵ-compatible.

(x, y, c)-Admissible holes: a pair $\{H_i, H_j\}$ of holes such that H_i is x-admissible, H_j is y-admissible, and the total order on $P \cup \{x, y\}$ obtained by adding $a_i \preccurlyeq x \preccurlyeq a_{i+1}$ and $a_j \preccurlyeq y \preccurlyeq a_{j+1}$ to $\preccurlyeq$ is $c\epsilon$-compatible.

$H(x)$: the set of all x-admissible holes H_i such that for any y there exists a y-admissible hole H_j such that H_i, H_j is a $(x, y, 1)$-admissible pair.

X_{ij}: the set of all elements x such that $H(x) = [a_i, a_j]$.

Linked elements: two elements $x, y \in X_{ij}$ which must be located in the same hole in any ϵ-compatible order on $P \cup \{x, y\}$.

Separated elements: two elements $x, y \in X_{ij}$ which must be located in distinct bounding holes in any ϵ-compatible order on $P \cup \{x, y\}$.

Graph $\mathcal{L}_{ij}$: X_{ij} is the set of vertices and $L_{ij} = \{xy : x, y \in X_{ij} \text{ and } d(x, y) \ll_3 \max\{d_x, d_y\}\}$ is the set of edges.

Block: a connected component of the graph $\mathcal{L}_{ij}$.

$x \twoheadrightarrow y$: if either $d_x \ll_4 d_y$ or $d_x \gtrsim_4 d_y$ and there exists an element $z \in X_{ij}$ such that $xz, yz \notin L_{ij}$ and $d(x, z) \ll_{16} d(y, z)$.

Graph $\overrightarrow{\mathcal{L}_{ij}}$: its vertices are the elements of X_{ij} and there is an arc $x \rightarrow y$ iff either (L1) $x \twoheadrightarrow y$ and x, y belong to a common block of $\mathcal{L}_{ij}$ or (L2) $d(x, y) \ll_5 \max\{d_x, d_y\}$.

Cell: a strongly connected component of the graph $\overrightarrow{\mathcal{L}_{ij}}$.

Graph Ψ_{ij}: its vertices are the cells of $\overrightarrow{\mathcal{L}_{ij}}$ and there is an edge between two cells C', C'' iff either (S1) there exist $x \in X_{ij}$ belonging to the same block as C' and $y \in X_{ij}$ belonging to the same block as C'' such that $d(x, y) \gg_3 \max\{d_x, d_y\}$ or (S2) there exist $x, x' \in X_{ij}$ belonging to the same block as C' and $y, y' \in X_{ij}$ belonging to the same block as C'' such that the pairs xx' and yy' are strongly linked, and $x \twoheadrightarrow y, y' \twoheadrightarrow x'$.

Twin clusters: the two halves of a connected component of the (bipartite) graph Ψ_{ij}.

Graph $\overrightarrow{\mathcal{G}_{ij}}$: its vertices are the cells of the graph $\overrightarrow{\mathcal{L}_{ij}}$, and there exists an arc $C' \rightarrowtail C$ iff (G1) either C' and C belong to twin clusters of Ψ_{ij} or (G2) C' and C are not connected by (G1)-arcs and there exist $x \in C$ and $x' \in C'$ such that $d_{x'} \ll_4 d_x$ or (G3) C' and C are not connected by (G1) and (G2)-arcs and there exist $x \in C, x' \in C'$, and $z \in X_{ij}$ such that $xz, x'z \notin L_{ij}$ and $d(x', z) \ll_{16} d(x, z)$.

Notations:
$\alpha \approx_c \beta$: $|\alpha - \beta| \le c\epsilon$;
$\beta \gtrsim_c \alpha : \beta \ge \alpha - c\epsilon$; $\beta \gg_c \alpha$ if $\beta > \alpha + c\epsilon$;
d_x: the mean value of $\min\{d(x, a_k) : a_k \in H(x) \setminus \{a_i, a_j\}\}$ if $H(x) = [a_i, a_j]$.

Acknowledgements We wish to thank an anonymous referee for a careful reading of the first version of the manuscript, his useful suggestions, and for pointing us to the reference [3] which leads to the result of Sect. 7. We are also grateful to Bernard Fichet for numerous insightful discussions and for his encouragement during the work on this paper. This research was partly supported by the ANR grant BLAN06-1-138894 (project OPTICOMB).

References

1. Agarwala, R., Bafna, V., Farach, M., Narayanan, B., Paterson, M., Thorup, M.: On the approximability of numerical taxonomy (fitting distances by tree metrics). SIAM J. Comput. **17**, 1073–1085 (1999)
2. Ailon, N., Charikar, M.: Fitting tree metrics: hierarchical clustering and phyloleny. In: FOCS 2005
3. Alon, N., Bădoiu, M., Demaine, E.D., Farach-Colton, M., Hajiaghayi, M., Sidiropoulos, A.: Ordinal embeddings of minimum relaxation: general properties, trees, and ultrametrics. In: SODA 2005
4. Aspval, B., Plass, M.F., Tarjan, R.E.: A linear time algorithm for testing the truth of certain quantified boolean formulas. Inf. Process. Lett. **8**, 121–123 (1979)

5. Atkins, J.E., Boman, E.G., Hendrickson, B.: A spectral algorithm for seriation and the consecutive ones problem. SIAM J. Comput. **28**, 297–310 (1998)
6. Bădoiu, M.: Approximation algorithm for embedding metrics into a two-dimensional space. In: SODA 2003
7. Bădoiu, M., Chuzhoy, J., Indyk, P., Sidiropoulos, A.: Low-distortion embeddings of general metrics into the line. In: STOC 2005
8. Bădoiu, M., Gupta, A., Dhamdhere, K., Rabinovich, Y., Räcke, H., Ravi, R., Chuzhoy, J., Sidiropoulos, A.: Approximation algorithms for low-distortion ebeddings into low-dimensional spaces. In: SODA 2005
9. Barthelemy, J.-P., Brucker, F.: NP-hard approximation problems in overlapping clustering. J. Classif. **18**, 159–183 (2001)
10. Benzer, S.: The fine structure of the gene. Sci. Am. **206**, 70–84 (1962)
11. Berry, M.W., Hendrickson, B., Raghavan, P.: Sparse matrix reordering schemes for browsing hypertext. In: Renegar, J., Shub, M., Smale, S. (eds.) Lectures in Applied Mathematics. The Mathematics of Numerical Analysis, vol. 32. American Mathematical Society, Providence (1996)
12. Bertin, J.: Graphics and Graphic Information Processing. de Gruyter, Berlin (1981)
13. Bertrand, P.: Etude de la représentation pyramidale. Thèse de 3ème cycle, Université Paris IX-Dauphine (1986)
14. Booth, K.S., Lueker, G.E.: Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. J. Comput. Syst. Sci. **13**, 335–379 (1976)
15. Brito, P., Bertrand, P., Cucumel, G., de Carvalho, F. (eds.): Selected Contributions in Data Analysis and Classification. Studies in Classification, Data Analysis, and Knowledge Organization. Springer, Berlin (2007)
16. Brusco, M.J.: A branch-and-bound algorithm for fitting anti-Robinson structures to symmetric dissimilarity matrices. Psychometrica **67**, 459–471 (2002)
17. Caraux, G., Pinloche, S.: PermutMatrix: a graphical environment to arrange gene expression profiles in optimal linear order. Bioinformatics **21**, 1280–1281 (2005)
18. Chen, C.H.: Generalized association plots for information visualization: the applications of the convergence of iteratively formed correlation matrices. Stat. Sin. **12**, 1–23 (2002)
19. Chen, C.H., Hwu, H.G., Jang, W.J., Kao, C.H., Tien, Y.J., Tzeng, S., Wu, H.M.: Matrix visualization and information mining. In: COMPSTAT 2004, pp. 85–100
20. Chepoi, V., Fichet, B.: Recognition of Robinsonian dissimilarities. J. Class. **14**, 311–325 (1997)
21. Chepoi, V., Fichet, B.: l_∞-approximation via subdominants. J. Math. Psychol. **44**, 600–616 (2000)
22. Chepoi, V., Fichet, B., Seston, M.: Seriation in the presence of errors: NP-hardness of l_∞-fitting Robinson structures to dissimilarity matrices. J. Class. (to appear)
23. Chor, B., Sudan, M.: A geometric approach to betweenness. SIAM J. Discrete Math. **11**, 511–523 (1998)
24. Critchley, F., Fichet, B.: The partial order by inclusion of the principal classes of dissimilarity on a finite set, and some properties of their basic properties. In: van Cutsen, B. (ed.) Classification and Dissimilarity Analysis. Lecture Notes in Statistics, pp. 5–65. Springer, Berlin (1994)
25. Cunningham, J.P., Shepard, R.N.: Monotone mapping of similarities into a general metric space. J. Math. Psychol. **11**, 335–364 (1974)
26. Dhamdhere, K.: Approximating additive distortion of line embeddings. In: APPROX-RANDOM 2004
27. Diday, E.: Orders and overlapping clusters by pyramids. In: de Leeuw, J., Heiser, W., Meulman, J., Critchley, F. (eds.) Multidimensional Data Analysis, pp. 201–234. DSWO, Leiden (1986)
28. Durand, C.: Ordres et graphes pseudo-hiérarchiques: théorie et optimisation algorithmique. Thèse de l'Université de Provence, Marseille (1989)
29. Durand, C., Fichet, B.: One-to-one correspondences in pyramidal representation: a unified approach. In: Bock, H.H. (ed.) Classification and Related Methods of Data Analysis, pp. 85–90. North-Holland, Amsterdam (1988)
30. Eisen, M.B., Spellman, P.T., Browndagger, P.O., Botstein, D.: Cluster analysis and display of genome-wide expression patterns. Proc. Natl. Acad. Sci. USA **95**, 14863–14868 (1998)
31. Farach, M., Kannan, S., Warnow, T.: A robust model for finding optimal evolutionary trees. Algorithmica **13**, 155–179 (1995)
32. Kruskal, J.B.: Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. Psychometrika **29**, 1–28 (1964)
33. Kruskal, J.B.: Non-metric multidimensional scaling. Psychometrika **29**, 115–129 (1964)

34. Fichet, B.: Data analysis: geometric and algebraic structures. In: Prohorov, Yu.A., Sazonov, V.V. (eds.) First World Congress of Bernoulli Society Proceedings, Tashkent, USSR, 1986, vol. 2. VNU Science Press, Utrecht (1986)
35. Fulkerson, D.R., Gross, O.A.: Incidence matrices and interval graphs. Pac. J. Math. **15**, 835–855 (1965)
36. Hahsler, M., Hornik, K., Buchta, C.: Getting things in order: an introduction to the R package seriation. J. Stat. Softw. **25**, 1–34 (2008)
37. Halperin, D.: Musical chronology by seriation. Comput. Humanit. **28**, 13–18 (1994)
38. Håstad, J., Ivansson, L., Lagergren, J.: Fitting points on the real line and its application to RH mapping. J. Algorithms **49**(1), 42–62 (2003)
39. Hubert, L.J.: Some applications of graph theory and related nonmetric techniques to problems of approximate seriation: The case of symmetric proximity measures. Br. J. Math. Stat. Psychol. **27**, 133–153 (1974)
40. Hubert, L., Arabie, P., Meulman, J.: The Structural Representation of Proximity Matrices with Matlab. ASA-SIAM Series on Statistics and Applied Probability. SIAM, Philadelphia (2006)
41. Indyk, P., Matousek, J.: Low-distortion embeddings of finite metric spaces. In: Goodman, J.E., O'Rourke, J. (eds.) Handbook of Discrete and Computational Geometry, 2nd edn., pp. 177–196. CRC Press, Boca Raton (2004)
42. Johnson, D.S.: The NP-completeness column: an outgoing guide. J. Algorithms **3**, 182–195 (1982)
43. Kendall, D.G.: Incidence matrices, interval graphs and seriation in archeology. Pac. J. Math. **28**, 565–570 (1969)
44. Kendall, D.G.: Seriation from abundance matrices. In: Hodson, F.R., Kendall, D.G., Tautu, P. (eds.) Mathematics in the Archaeological and Historical Sciences, pp. 215–252. Edinburgh University Press, Edinburgh (1971)
45. Klinz, B., Rudolf, R., Woeginger, G.: On the recognition of permutted bottleneck Monge matrices. Discrete Appl. Math. **63**, 43–74 (1995)
46. Ling, R.L.: A computer generated aid for cluster analysis. Commun. ACM **16**, 355–361 (1973)
47. Marquardt, W.: Advances in archaeological seriation. In: Schiffer, M. (ed.) Advances in Archaeological Method and Theory 1, pp. 257–314. Academic Press, San Diego (1978)
48. Miklos, I., Somodi, I., Podani, I.: Rearrangement of ecological data matrices via Markov chain Monte Carlo simulation. Ecology **86**, 3398–3410 (2005)
49. Mirkin, B., Rodin, S.: Graphs and Genes. Springer, Berlin (1984)
50. Robinson, W.S.: A method for chronologically ordering archaeological deposits. Am. Antiq. **16**, 293–301 (1951)
51. Saxe, J.B.: Embeddability of weighted graphs in k-space is strongly NP-hard. In: Proceedings of the 17th Allerton Conference on Communications, Control, and Computing, pp. 480–489 (1979)
52. Semple, C., Steel, M.: Phylogenetics. Oxford University Press, London (2003)
53. Sneath, P.H.A., Sokal, R.R.: Numerical Taxonomy. Freeman and Company, San Francisco (1973)
54. Strehl, A., Ghosh, J.: Relationship-based clustering and visualization for high-dimensional data mining. INFORMS J. Comput. **15**, 208–230 (2003)
55. Torgerson, W.S.: Multidimensional scaling I: theory and method. Psychometrika **17**, 401–414 (1952)
56. Van de Vel, M.: Theory of Convex Structures. Elsevier, Amsterdam (1993)
57. Vazirani, V.: Approximation Algorithms. Springer, Berlin (2001)
58. Whittaker, R.H. (ed.): Ordination of Plant Communities. Junk, The Hague (1978)