The Complexity of the Partial Order Dimension Problem – Closing the Gap

Stefan Felsner*
Institut für Mathematik
Technische Universität Berlin
Strasse des 17. Juni 136
D-10623 Berlin, Germany

Irina Mustaţă
Berlin Mathematical School
Technische Universität Berlin
Strasse des 17. Juni 136
D-10623 Berlin, Germany

Martin Pergel†
Department of Software and
Computer Science Education
Charles University
Malostranské nám. 25
118 00 Praha 1, Czech Republic

Abstract

The dimension of a partial order P is the minimum number of linear orders whose intersection is P. There are efficient algorithms to test if a partial order has dimension at most 2. In 1982 Yannakakis [25] showed that for $k \geq 3$ to test if a partial order has dimension $\leq k$ is NP-complete. The height of a partial order P is the maximum size of a chain in P. Yannakakis also showed that for $k \geq 4$ to test if a partial order of height 2 has dimension $\leq k$ is NP-complete. The complexity of deciding whether an order of height 2 has dimension 3 was left open. This question became one of the best known open problems in dimension theory for partial orders. We show that the problem is NP-complete.

Technically, we show that the decision problem (3DH2) for dimension is equivalent to deciding for the existence of bipartite triangle containment representations (BTCon). This problem then allows a reduction from a class of planar satisfiability problems (P-3-CON-3-SAT(4)) which is known to be NP-hard.

Mathematics Subject Classifications (2010) 06A07, 68Q25, 05C62,
Keywords : NP-complete, triangle containment, intersection graph.

1 Introduction

Let $P = (X, \leq_P)$ be a partial order. A linear order $L = (X, \leq_L)$ on X is a *linear extension* of P when $x \leq_P y$ implies $x \leq_L y$. A family $\mathcal{R}$ of linear extensions of P is a *realizer* of P if $P = \bigcap_i L_i$, i.e., $x \leq_P y$ if and only if $x \leq_L y$ for every $L \in \mathcal{R}$. The *dimension* of P, denoted $\dim(P)$, is the minimum size of a realizer of P. This notion of dimension for partial orders was defined by Dushnik and Miller [5]. The dimension of P can, alternatively, be defined as the minimum t such that P admits an order preserving embedding into the product order on $\mathbb{R}^t$, i.e., elements $x \in X$ have associated vectors $\hat{x} = (x_1, \ldots, x_t)$ with real entries, such that $x \leq_P y$ if and only if $x_i \leq y_i$ for all i. In the sequel, we denote this by $\hat{x} \leq \hat{y}$. The concept of dimension plays a role which in

*Partially supported by DFG grant FE-340/7-2 and ESF EuroGIGA projects Compose and GraDR.
†Partially supported by a Czech research grant GAČR GA14-10799S.

many instances is analogous to the chromatic number for graphs. It has fostered a large variety of research about partial orders, see e.g. Trotter's monograph [23].

The first edition of Garey and Johnson [11] listed the decision problem, whether a partial order has dimension at most k, in their selection of twelve important problems which were not known to be polynomially solvable or NP-complete. The complexity status was resolved by Yannakakis [25], who used a reduction from 3-colorability to show that the problem is NP-complete for every fixed $k \geq 3$. He also showed that the problem remains NP-complete for every fixed $k \geq 4$ if the partial order is of height 2. The recognition of partial orders of dimension ≤ 2 is easy. Efficient algorithms for the problem have been known since the early 1970's (e.g. [6]) a linear time algorithm is given in [17]. The gap that remained in the complexity landscape was that for partial orders of height 2, the complexity of deciding if the dimension is at most 3 was unknown. This was noted by Yannakakis [25], but also listed as Problem 1 in Spinrad's account on dimension and algorithms [21]. It was also mentioned at several other places, e.g. [24]. In this paper we prove NP-completeness for this case.

Schnyder's characterization of planar graphs in terms of order dimension promoted interest in this gap in the complexity landscape. With a finite graph $G = (V, E)$ we associate the *incidence order* $P_{VE}(G)$ of vertices and edges. The ground set of $P_{VE}(G)$ is $V \cup E$. The order relation is defined by $x < e$ in P_{VE} if $x \in V$, $e \in E$ and $x \in e$. The incidence order of a graph is an order of height two. Schnyder [20] proved: A graph G is planar if and only if the dimension of its incidence order is at most 3. This characterizes a reasonably large class of orders of height 2 and dimension 3. Efficient planarity tests yield polynomial-time recognition algorithms for orders in this class. Motivated by Schnyder's result, Trotter [24] asked about the complexity of deciding if $\dim(P) \leq 3$ for the class of orders P of height 2 with the property that every maximal element covers at most k minimal elements. As a by-product of our proof, we provide an answer: the problem remains NP-complete even if each element is comparable to at most 5 other elements.

In [22] Trotter stated the following interesting problem:

a) For fixed $t \geq 4$, is it NP-complete to determine whether the dimension of the incidence order of a graph is at most t?

Dimension seems to be a particularly hard NP-complete problem. This is indicated by the fact that we have no heuristics or approximation algorithms to produce realizers of partial orders that have reasonable size. A hardness result for approximations was first obtained by Hegde and Jain [12] and recently strengthened by Chalermsook et al. [4]. They show that unless NP = ZPP there exists no polynomial-time algorithm to approximate the dimension of a partial order with a factor of $O(n^{1-\varepsilon})$ for any $\varepsilon > 0$, where n is the number of elements of the input order. The reduction indeed shows more: 1) the same hardness result holds for partial orders of height 2. 2) Approximating the fractional dimension (c.f. [2, 9]) of a partial order is also hard. Another class of partial orders where computing the dimension was shown to be NP-complete [13] are the N-free partial orders, i.e., orders whose cover graphs have no induced path of length four. A motivation for the investigation was that an algorithm for computing the dimension of N-free partial orders would have implied a constant factor approximation algorithm for general orders. As far as approximation is concerned, the best known result is a factor $O(n\frac{\sqrt{\log\log n}}{\sqrt{\log n}})$ approximation which comes from an approximation algorithm with the same factor for boxicity, see [1].

The following two problems about computational aspects of dimension are open:

b) For fixed $t \geq 3$, is it NP-complete to determine whether the dimension of an interval order is at most t?

c) For fixed $w \geq 3$, is it NP-complete to decide whether the dimension of order of width w is precisely w?

1.1 Outline of the paper

In the next section we provide some technical background on dimension of orders. In particular we show that 3-dimensional orders are characterized by having a containment representation with homothetic triangles. In Subsection 2.1 we prove the *disjoint paths lemma* and the *rotor lemma*, two crucial lemmas for the reduction.

In Section 3 we formally introduce the decision problem P-3-CON-3-SAT(4), a special planar version of 3-SAT and present the reduction from P-3-CON-3-SAT(4) to the bipartite triangle containment problem BTCon. We explain the *clause gadget* and the *variable gadget* and show how to build the graph G_Φ for a given P-3-CON-3-SAT(4) instance Φ. It is then shown that a triangle containment representation of G_Φ exists iff there is a satisfying assignment for Φ. Theorem 2 summarizes the result.

In Section 4 we modify the construction to prove a 'sandwich result', Theorem 3. The theorem asserts hardness of recognition for every class $\mathcal{C}$ between BTCon and PUTCon (point unit-triangle containment). For the proof we introduce a new variable gadget and construct a graph H_Φ such that if Φ is satisfiable, then H_Φ has a triangle containment representation using triangles of only two different sidelengths, while if Φ is not satisfiable then H_Φ has no triangle containment representation. As an application we show that recognition of bipartite graphs that admit a special representation with triangle intersection (BTInt) is again NP-hard. We then show that PUTCon and BTInt are in NP. We conclude with two open questions and remarks on extensions.

2 Dimension Three and Triangle Containment

In the introduction we have mentioned the two definitions of dimension for partial orders. For completeness, we include a proof of the equivalence.

Proposition 1 *A partial order $P = (X, \leq_P)$ has a realizer of size t if and only if it admits an order preserving embedding into $(\mathbb{R}^t, \leq)$.*

Proof. Let $L_1, \ldots, L_t$ be a realizer of P. For $x \in X$ and $i \in [t]$ let x_i be the position of $x \in L_i$, i.e., $x_i = |\{y : y \leq_{L_i} x\}|$. The map $x \to (x_1, \ldots, x_t)$ is an order preserving embedding into $(\mathbb{N}^t, \leq)$.

If an order preserving embedding of P into $\mathbb{R}^t$ is given we consider the projection orders W_i defined by $x <_{W_i} y$ if and only if $x_i < y_i$. Elements of P which share the value of the i-th coordinate form antichains of W_i, i.e., being a linear order of antichains W_i is a *weak-order*. If L_i is taken as an arbitrary linear extension of W_i that is also a linear extension of P this yields a realizer $L_1, \ldots, L_t$ of P. □

The definition of dimension via realizers is related to the notion of *reversible sets of incomparable pairs* and *alternating cycles* which have been applied throughout the literature on dimension of partial orders. The definition of dimension via embeddings into product orders is used less frequently. The most notable application may be in the context of a generalization of Schnyder's dimension theorem due to Brightwell and Trotter.

Theorem 1 (Brightwell-Trotter Theorem [3]) *The incidence order $P_{VEF}(G)$ of the vertices, edges and bounded faces of a 3-connected plane graph G has dimension three.*

Miller [18] gave a proof of the theorem by constructing a rigid orthogonal surface. Based on the connection between Schnyder woods and rigid orthogonal surfaces, simpler proofs of Theorem 1 were obtained in [7] and [10]. We omit the details here, but Figure 1 shows an example of a rigid orthogonal surfaces with the corresponding embedding of $P_{VEF}(G)$.

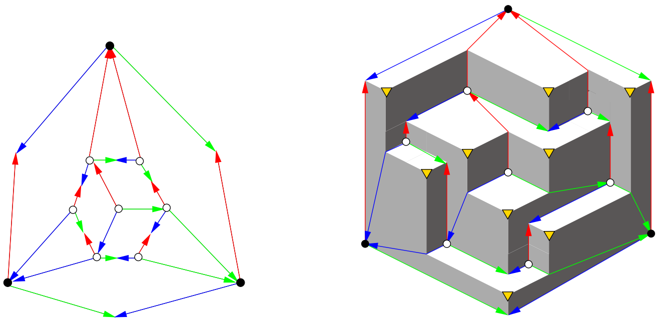

Figure 1: Left: a 3-connected plane graph G with a Schnyder wood. Right: a rigid orthogonal surface corresponding to the Schnyder wood. The vertices (circles), faces (triangles) and bend points of edges show an order preserving embedding of $P_{VEF}(G)$ into $\mathbb{R}^3$.

The key to our reduction is the characterization of 3-dimensional orders which is given in the following proposition. A *homothet* of a geometric object O is an object O' that can be obtained by scaling and translating O. With any family $\mathcal{F}$ of sets we associate the *containment order* $(\mathcal{F}, \subseteq)$.

Proposition 2 *The dimension of a partial order $P = (X, \leq_P)$ is at most t if and only if P is isomorphic to the containment order of a family of homothetic simplices in $\mathbb{R}^{t-1}$.*

Proof. "$\Rightarrow$" Let $x \to \hat{x}$ be an order preserving embedding of P to $\mathbb{R}^t$. With a point $\hat{x}$ in $\mathbb{R}^t$ associate the (infinite) orthogonal cone $C(\hat{x}) = \{p \in \mathbb{R}^t \ : \ p \leq \hat{x}\}$. Note that $x \leq_P y$ if and only if $C(\hat{x}) \subseteq C(\hat{y})$. Consider an oriented hyperplane H such that for all $x \in X$ the point $\hat{x}$ is in the positive halfspace H^+ of H, moreover, $\lambda\mathbb{1} \in H^+$ and $-\lambda\mathbb{1} \in H^-$ for some positive λ and the all-ones vector $\mathbb{1}$. For $x \in X$, the intersections of the cone $C(\hat{x})$ with H is a $(t-1)$-dimensional polytope $\Delta(x)$ with t vertices, i.e., a simplex. The simplices for different elements are homothetic and $\Delta(x) \subseteq \Delta(y)$ iff $C(\hat{x}) \subseteq C(\hat{y})$ iff $x \leq_P y$. Hence, P is isomorphic to the containment order of the family $\{\Delta(x) \ : \ x \in X\}$ of homothetic simplices in $H \cong \mathbb{R}^{t-1}$.
"$\Leftarrow$" Now suppose that there is a containment order of a family $\mathcal{F}$ of homothetic simplices in $\mathbb{R}^{t-1}$ that is order isomorphic to P. Let $\Delta(x)$ be the simplex corresponding to x in the order isomorphism. Apply an affine transformation to get a family $\mathcal{F}' = \{\Delta'(x) \ : \ x \in X\}$ of regular simplices in $H = \mathbb{R}^{t-1}$ with the same containment order. Embed H into $\mathbb{R}^t$ with normal vector $\mathbb{1}$. For each $\Delta'(x)$ there is a unique point $\hat{x} \in \mathbb{R}^t$ such that $C(\hat{x}) \cap H = \Delta'(x)$. Since the containment orders of $\{C(\hat{x}) \ : \ x \in X\}$ and $\{\Delta'(x) \ : \ x \in X\}$ are isomorphic we identify $x \to \hat{x}$ as an order preserving embedding of P into $(\mathbb{R}^t, \leq)$. $\square$

An instance of this proposition, that has been mentioned frequently in the literature, is the equivalence between 2-dimensional orders and containment orders of intervals.

2.1 Lemmas for triangle containment

From now on we will restrict the attention to partial orders of height 2. Note that these orders have a bipartite comparability graph. Conversely, any bipartite graph with black and white vertices can be seen as a height 2 order: define $u < w$ whenever u is white, w is black and (u, w) is an edge. Hence, partial orders of height 2 and bipartite graphs are essentially the same.

Given a triangle containment representation of a bipartite graph $G = (V, E)$, let $B(V)$ be the set of barycenters of the triangles (it can be assumed that all barycenters are different). Define the β-graph $\beta(G)$ as the straight line drawing of G with vertices placed at their corresponding points of $B(V)$. The following lemma allows some control on the crossings of edges in $\beta(G)$.

Two edges are called *strongly independent* if they share no vertex (i.e., they are independent) and they are the only edges induced on their four vertices. The lemma is closely related to the easy direction of Schnyder's theorem [20, Theorem 4.1].

Lemma 1 (disjoint paths lemma) *In $\beta(G)$ there is no crossing between two strongly independent edges.*

Proof. Map the triangle representation via an affine transformation to a plane H with normal $\mathbb{1}$ in $\mathbb{R}^3$, such that the triangles on H are equilateral. As in the proof of Proposition 2, we obtain an embedding $v \to \hat{v}$ of the vertices of V to $\mathbb{R}^3$ such that the intersection of the cone $C(v)$ with H is the triangle T_v associated with v. The containment relation of triangles T_v on H and the comparability of points $\hat{v}$ in $\leq$ both realize G.

Let $e = (u, v)$ and $f = (x, y)$ be strongly independent edges in G such that $T_u \subset T_v$ and $T_x \subset T_y$. Now suppose that in $\beta(G)$ the two edges intersect in a point p. The line $\ell(p) = p + \lambda\mathbb{1}$ shares a point p_e with the segments $\hat{e} = [\hat{u}, \hat{v}]$ and a point p_f with the segment $\hat{f} = [\hat{x}, \hat{y}]$. Without loss of generality we may assume that p_f separates p_e and p on $\ell(p)$. Now we look at the apices of the cones. From triangle containment we obtain $\hat{u} < p_e < \hat{v}$ and $\hat{x} < p_f < \hat{y}$ and along $\ell(p)$ we get $p < p_f < p_e$ (see Fig. 2). Combining this and using transitivity we get $\hat{x} < \hat{v}$. This implies $T_x \subset T_v$ and $(x, v) \in E$, contradiction. □

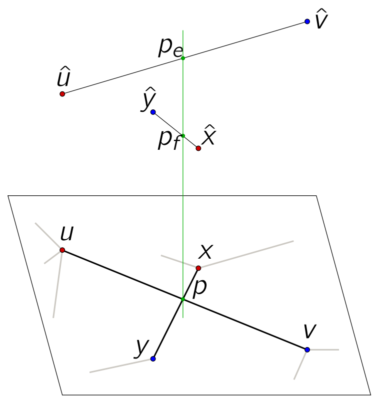

Figure 2: Illustrating the proof of Lemma 1. A crossing pair of strongly independent edges in the β-graph and the implied comparability $\hat{x} < \hat{v}$.

For the reduction we will construct a bipartite graph with an embedding in the plane which has only few crossings. Most of these crossings between edges occur locally in subgraphs that are named *rotor*. A rotor has a center, which is an adjacent pair u, v of vertices. Additionally there are some non-crossing paths p_i. Each path p_i is connected to the center either with an edge (x_i, u) or with an edge (x_i, v) where x_i is an end-vertex of the path. The interesting case of a rotor is an *alternating rotor*. In this case it is possible to add a simple closed curve γ to the picture such that (1) γ contains the center and (2) there is a collection of six paths intersecting the curve γ cyclically so that paths leading to u and paths leading to v alternate. Figure 3(b) shows a triangle containment representation of an alternating rotor.

We will show that in the β-graph of a triangle containment representation of an alternating rotor the six paths $p_1, \ldots, p_6$ which do the alternation have to connect to the center in a very specific way. We assume that p_1, p_3, p_5 connect to u and p_2, p_4, p_6 connect to v.

Let u, v be the center of a rotor with $T_u \subset T_v$. Define the *shadow intervals* as the intervals between the intersections of the sides of T_v with the supporting lines of T_u, see Figure 3(a). As *tip* *regions* of the triangle T_v, we define the three parallelograms determined by the sides of T_v and the support lines of T_u. We have the following

Claim 1. In a representation of an alternating rotor

- each of $T_{x_1}, T_{x_3}, T_{x_5}$ contains exactly one of the shadow intervals and each of the shadow intervals is contained in one of the triangles, and
- each of $T_{x_2}, T_{x_4}, T_{x_6}$ intersects exactly one of the three tip regions of T_v and each of the tip regions is intersected by one of the triangles.

Proof. From Lemma 1 we deduce that in the β-graph the paths p_i are pairwise non-intersecting, moreover the edges (u, x_i), with odd i, are disjoint from paths p_j with even j. Hence, the three paths emanating from u define three regions and each of these regions contains one of the remaining three paths, see Figure 3(c). Now each of $T_{x_1}, T_{x_3}, T_{x_5}$ contains T_u but it is not contained in T_v, hence, it has to contain a shadow interval. If two of them, say T_{x_1} and T_{x_3}, contain the same shadow interval then there is no space for T_{x_2} in T_v such that in the β-graph p_2 is between p_1 and p_3. Since T_{x_j}, with even j, is not contained in any T_{x_i}, with odd i, it has to intersect one of the tip regions of T_v. Again the alternation of the paths connecting to u and v implies the one-to-one correspondence between $T_{x_2}, T_{x_4}, T_{x_6}$ and the tip regions intersected by them. △

The essence of the claim is that at an alternating rotor there are six *ports* where paths can attach, these ports alternatingly belong to the center vertices u and v. This property is captured by the schematic picture of an alternating rotor given in Figure 3(d).

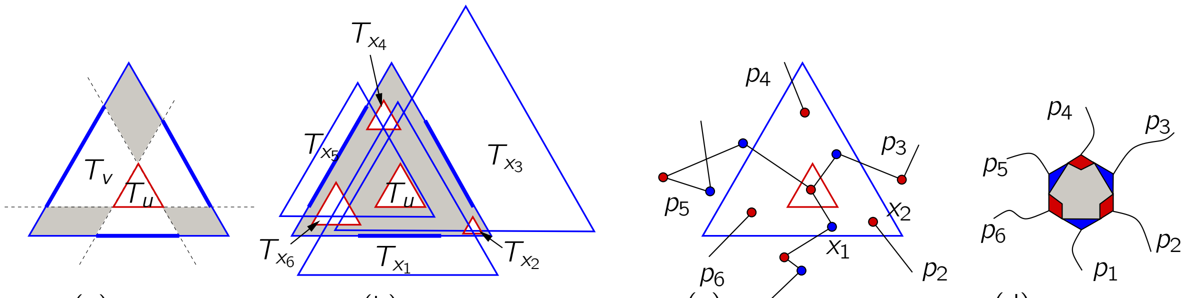

Figure 3: Part (a) shows the shadow intervals induced by T_u on the sides of T_v and in gray the tip regions. Part (b) shows a triangle containment representation of a rotor. Part (c) illustrates how the paths p_1, p_3 and p_5 together with u in the β-graph partition the interior of T_v in three regions, where the paths p_2, p_4 and p_6 start. Part (d) shows a schematized picture of an alternating rotor.

An *alternating 8-rotor* is a rotor with an 8 alternation, i.e., there are eight paths $p_1, \ldots, p_8$ intersecting a simple closed curve γ that contains the center cyclically in the order of indices such that paths leading to u and paths leading to v alternate.

Lemma 2 *There is no triangle containment representation of an alternating 8-rotor such that in the β-graph of the representation the eight paths $p_1, \ldots, p_8$ intersect a simple closed curve around the center of the rotor in the order of the indices.*

Proof. Assume that there is a representation. By considering the representation of the two alternating rotors with paths $p_1, \ldots, p_6$ and $p_3, \ldots, p_8$, respectively, we conclude from Claim 1 that T_{x_1} and T_{x_7} contain the same shadow interval on T_v. This contradicts the representation of the alternating rotor with paths $p_1, p_2, p_3, p_4, p_7, p_8$. □

3 The Reduction

3.1 Decision problems

Dimension 3 for Height 2 Orders (3DH2)
Instance: A partial order $P = (X, <)$ of height 2.
Question: Is the dimension of P at most 3?

Bipartite Triangle Containment Representation (BTCon)
Instance: A bipartite graph $G = (V, E)$.
Question: Does G admit a containment representation with a family of homothetic triangles?

From the $t = 3$ case of Proposition 2 it follows that the decision problems 3DH2 and BTCon are polynomially equivalent. In this section we design a reduction from P-3-CON-3-SAT(4) to show that BTCon and, hence, also 3DH2 is NP-complete.

In order to define this special version of 3-SAT, we have to recall the notion of the *incidence graph* of a SAT instance Φ. This is a bipartite graph whose vertices are in correspondence to the clauses on one side of the partition, and to the variables of Φ on the other side. Edges of the incidence graph correspond to membership of a variable in a clause. Lichtenstein [16] showed that the satisfiability problem for CNF-formulae with a planar incidence graph is NP-complete. Based on this, more restricted Planar-SAT variants have been shown to be hard. The NP-hardness of P-3-CON-3-SAT(4) was shown by Kratochvíl [14]. One of the first applications was to show the hardness of recognizing grid intersection graphs.

P-3-Con-3-SAT(4)
Instance: A SAT instance Φ with the additional properties:

- The incidence graph of Φ is 3-connected and planar.
- Each clause consists of 3 literals.
- Each variable contributes to at most 4 clauses.

Question: Does Φ admit a satisfying truth assignment?

The advantage of working with 3-connected planar graphs, instead of just planar graphs, is that in the 3-connected case the planar embedding is unique up to the choice of the outer face (Whitney's Theorem). Our reduction is inspired by the reduction from [14], and even more so by the recent NP-completeness proof for the recognition of unit grid intersection graphs with arbitrary girth [19].

3.2 The idea for the reduction

Let Φ be an instance of P-3-CON-3-SAT(4) with Φ we assume a fixed embedding of the incidence graph I_Φ in the plane. Such an embedding can be part of the input, otherwise it can be computed efficiently. Our aim is to construct a bipartite graph G_Φ such that G_Φ has a triangle containment representation if and only if Φ is satisfiable. The construction of G_Φ is done by replacing the constituents of I_Φ by appropriate gadgets. The connections between the gadgets are given by pairs of paths that correspond to the edges of the 3-connected planar graph I_Φ.

The β-graph of a triangle containment representation of G_Φ is a drawing of G_Φ. Because of Lemma 1 all the crossings in this drawing are local. In particular there is no interference (crossings) between different gadgets, no interference between a gadget and a connecting paths (except when the corresponding vertex and edge of I_Φ are incident), and no interference between two of the connecting paths. Except for crossings along a path, see e.g. the path p_5 in Figure 3(c), we can find crossings of the drawing only within the gadgets. In particular the global structure of the drawing, i.e., the cyclic order in which pairs of paths leave a clause, is as prescribed by I_Φ.

In G_Φ every edge of the incidence graph I_Φ is replaced by a pair of paths. These two paths join a variable gadget and a clause gadget. Such a pair of paths is called an *incidence strand*. At their clause ends, the two paths of an incidence strand are two of the p_i paths of a rotor. At this rotor the two paths are incident to different center vertices, hence, they are distinguishable and we may think of one of them as the green path and the other as the yellow path. Assuming a triangle containment representation we look at the two paths of an incidence strand in the β-graph. Since they do not cross each other they define a strip. Looking along this strip in the direction from the variable gadget to the clause gadget we either see the green path on the left and the yellow path on the right boundary or the other way round. This yields an 'orientation' of the incidence strand. The orientation is used to transmit the truth assignment from the variable to the clause.

The notion of oriented strands is crucial for the design of the clause gadgets and variable gadgets described in the following two subsections.

3.3 The clause gadget

The clause gadget consists of a rotor surrounded by a cycle and two paths that fix the rotor in the interior of the cycle (magenta). The paths of the three incidence strands which lead to the clause are also connected to the rotor. Three vertices of the cycle have two extra edges connecting to the two paths of an incidence strand. This construction, we call it *gate*, enables the incidence strands to enter the interior of the cycle. Figure 4 shows the clause gadget as a graph and in a more schematic view. The lengths of paths shown in the left figure are chosen at will. For a precise description of the gadget the lenght of the paths and of the enclosing circle have to be chosen so that a triangle containment representation is possible in all cases of a satisfied clause, i.e., in all cases where the paths are non-crossing. It is quite obvious that if these lenghts are given by sufficiently large constants, then the triangle containment representation exists. A reasonable choice could be 314 for the enclosing circle and 50 for each of the eight paths from the rotor to the circle. Another detail that has to be verified is that the gates can be realized via triangle containment. This is shown in the right part of Figure 8.

In Figures 4 and 5 the two paths connecting the rotor to the cycle are shown in magenta. The gates are shown as blue elements (triangles) adjacent to a red vertex of each of the two paths of the incidence strand associated to the gate. The coloring of the paths of the incidence strand has been chosen such that at the rotor the yellow path ends at the red central element u and the green path at the blue central element v.

The orientation of the incidence strand corresponding to a literal transmits TRUE if one of the two paths of the strand can share a port with one of the magenta paths, it transmits FALSE if the two paths together with an adjacent magenta path have to use three different ports of the rotor. To exemplify this rule: the literal ℓ_1 is true if the corresponding incidence strand has yellow left of green when seen from the rotor, whereas the other two literals are true when on their strand green is left of yellow. Note that due to this asymmetry we have to represent a clause as an ordered 3-tuple of literals.

In the right part of Figure 4, we see that if the literals evaluate to (F, T, F), then the clause gadget can be drawn with noncrossing paths.

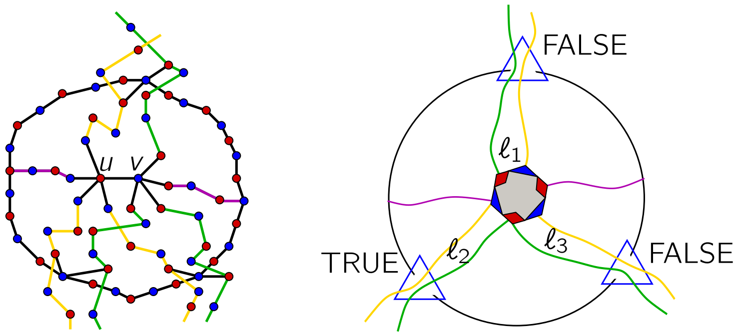

Figure 4: Left: The graph of a clause gadget. Right: A schematic view for the case where literals evaluate to (F, T, F).

Figure 5 shows that in all other cases where the clause evaluates to true, the schematic clause gadget can as well be drawn without crossing paths, thus warranting a corresponding triangle containment representation.

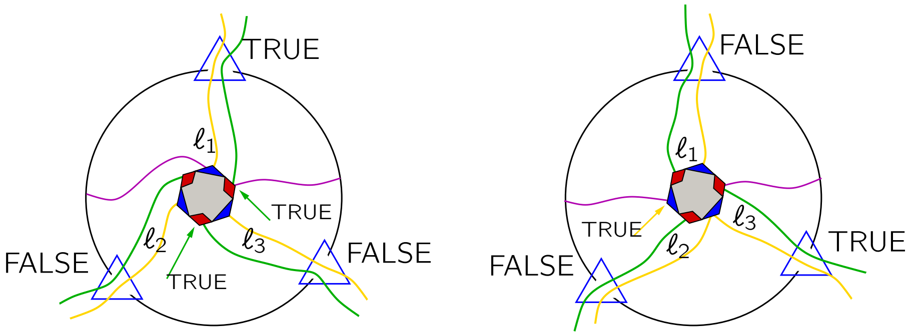

Figure 5: The left part of the figure shows that if the literals evaluate to (T, F, F), then the clause gadget can be drawn with noncrossing paths. In fact, by attaching the green paths of ℓ_2 and ℓ_3 where the arrows show, we get non-crossing representations for $(T, *, *)$. The right part shows the same for $(F, *, T)$.

Finally, we have to show that in the (F, F, F) case there is no triangle containment representation for the clause gadget. In this case the six paths of the three incidence strands together with the two (magenta) connecting paths of the rotor form an alternating 8-rotor, i.e, in the cyclic order the eight paths have to be connected to the center vertices u and v alternatingly. From Lemma 2 we know that such a configuration has no corresponding triangle containment representation. We summarize the result in a proposition.

Proposition 3 *The clause gadget has a triangle containment representation if and only if the two paths of the incidence strand of at least one literal are in the orientation representing TRUE.*

3.4 The variable gadget

The variable gadget corresponding to a variable x depends on the number of occurrences of the variable in clauses.

We begin with the simplest case. This is when the variable x has only two occurrences. The basic construction is extremely easy in this case. The variable x corresponds to one of the incidence strands at each clause that contains an occurrence of x. Let C_1 and C_2 be the two clauses containing an occurrence of x. Connect the four paths of the two incidence strands emanating from clauses C_1 and C_2 to a single pair of paths. As shown in Figure 6 there are two options, either we pairwise connect paths of the same color (color preserving) or we connect paths of different colors (color mixing). Which of the two options is the right one to synchronize the choice of truth values for the

Figure 6: Variable gadget for two occurrences. The left part of the figure shows a color preserving combination of the paths, the right part a color mixing combination.

literals depends on the position of the literal containing x in the ordered clauses C_1 and C_2 and it depends on whether the occurrence is negated or not. The right pairing for the connection can be determined through an easy 'calculation'. The result for the case of two positive or two negative literals is shown in the left table below, where P corresponds to a color preserving combination of the paths and M for a color mixing combination. The right table shows how to connect the paths if exactly one literal is negated. In this case with respect to the first table the entries P and M are simply exchanged.

	ℓ_1	ℓ_2	ℓ_3
ℓ_1	M	P	P
ℓ_2	P	M	M
ℓ_3	P	M	M

	ℓ_1	ℓ_2	ℓ_3
ℓ_1	P	M	M
ℓ_2	M	P	P
ℓ_3	M	P	P

We make an example of how to determine one of the entries of the tables: Consider the entry (ℓ_1, ℓ_2) of the second table. The variable x is the first literal of a clause C and $\neg x$ is the second literal of a clause C'. For C the literal ℓ_1 is T if at the rotor the port of the green path of ℓ_1 is the clockwise successor of the port of the yellow path. If x is T, then the literal ℓ_2 of C' has to be F, i.e., at the rotor of C' the port of the green path of ℓ_2 has to be a clockwise successor of the port of the yellow path. Hence, if x is T, then reading the combined strand in clockwise order starting at the rotor of C is yellow(ℓ_1, C), green(ℓ_2, C'), yellow(ℓ_2, C'), and green(ℓ_1, C). This shows that in this case the paths have to be combined so that the colors are mixed.

Now assume that the variable x has three occurrences. For two of the occurrences we can combine the paths of the incidence strands as in the previous case. To transmit the truth information to a third incidence strand we use a gadget involving an alternating rotor, the construction is shown in Figure 7. The colors are changed to avoid confusion with between the variable-ends and the clause-ends of strands.

In the Figure the combined incidence strands of two occurrences are shown in cyan and pink. The adjacencies of vertices on these paths to the center vertices of the rotor between them make an alternating rotor. The duty of the black path is to shield two of the ports of the rotor from outer access. The gate allows access for one of the incidence strands. There remain two ports of the rotor where the paths of the third occurrence (magenta ends) can connect to the center of the rotor. Switching the orientation of the first two occurrences also makes the strand of the third occurrence switch orientation, i.e., the truth values transmitted by the three strands are synchronized. Each connection between a strand of the variable gadget and the respective strand of a clause gadget has again to be adjusted so that the correct truth value is received at the clause.

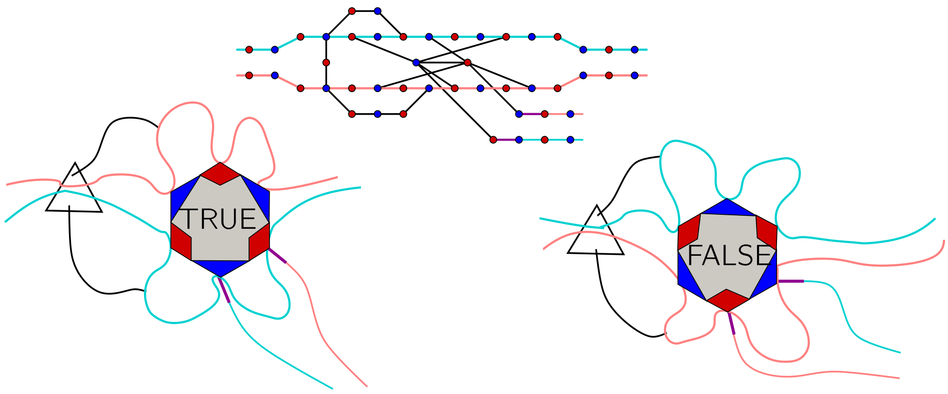

Figure 7: The variable gadget for a variable with three occurrences. The upper picture shows the variable gadget as a bipartite graph. The schematic drawings below show the two possible states of the variable gadget which correspond to the two possible values of the variable.

Such an adjustment consists in deciding whether the pairing is (cyan-green,pink-yellow) or (cyan-yellow,pink-green). Which of the pairings has to be chosen depends on (1) which strand of the variable gadget is in question, (2) the position of the literal in the clause, and (3) whether there is a negation involved. We refrain from listing the possible cases.

Finally, if the variable x has four occurrences, then we take two of the gadgets used for variables with three occurrences and connect two paths of a strand of each of them. This synchronizes the truth values transmitted by the remaining four strands.

3.5 Wrap-up

We have described the gadgets needed to construct the instance G_Φ for the decision problem BTCon based on an instance Φ for P-3-Con-3-SAT(4). The discussion of the clause gadget lead to the insight that a clause gadget has a triangle containment representation if and only if at least one of its literals evaluates to true (Proposition 3).

Regarding the variable gadget, we claim that a triangle containment representation exists whenever all the involved paths are given by a large enough constant. To see this we need the representation of an alternating rotor, and we have to verify that it is possible to realize the ports where a pink or cyan path is bypassing a magenta path (colors as in Figure 7). The realizability of the configuration is shown in Figure 8. Regarding the length of the path we propose to use 20 for each segment of a path between two vertices of degree three. (The value 20 is rather arbitrary. We chose it because it allows a realization of the variable gadget which is clearly arranged.)

The graph G_Φ is a subdivision of some core graph $G_\Phi^\bullet$ with minimum degree 3. To actually build the graph G_Φ, the length of the paths replacing edges of $G_\Phi^\bullet$ have to be specified. For the clause gadget and for variable gadgets we identify subdivisions so that they are representable in all cases except the unsatisfied clause. The corresponding path lengths are copied to all edges of $G_\Phi^\bullet$ that are internal to a gadget. For the incidence strands we have initial contributions to the path length that come from the gadgets. For the final adjustment of the path lengths of incidence strands, we consider a fixed straight line drawing of I_Φ on a small integer grid. The length of the path of an incidence strand can be taken as a multiple of the edge length of the corresponding edge in the drawing of I_Φ. The proportionality factor should be large enough to ensure that the paths can be represented with a sequence of small triangles so that unwanted interference between different paths can be avoided. If the enclosing circle of the clause gadget has length 314 as suggested in

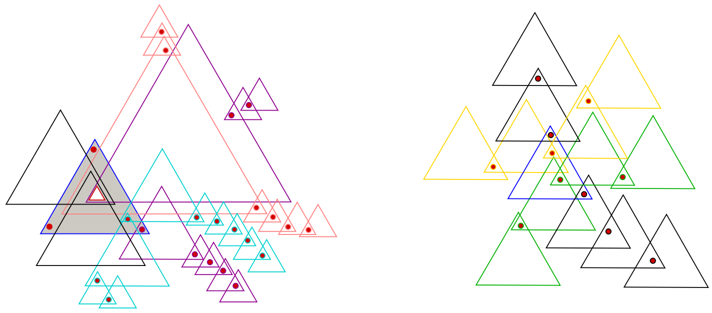

Figure 8: Left: The rotor of a variable gadget with a pink path bypassing a magenta path at a blue port and a cyan path bypassing a magenta path at an adjacent red port.
Right: A realization of a gate.

Subsection 3.3, then we may take the proportionality factor to be 500. Again the choice of 500 is rather arbitrary, but we think of a representation where the small triangles are points or almost points and the large triangles have sidelength close to one. With this assumption the clause and variable gadgets can live in disks of radius 25, hence they are small compared to the distance that can be spanned with a path, resp. strand, connecting two of the gadgets.

In summary: Given an instance Φ for P-3-Con-3-SAT(4) we can construct a bipartite graph G_Φ such that G_Φ admits a triangle containment representation if and only if Φ is satisfiable.

Since P-3-Con-3-SAT(4) is NP-complete [14], this shows that BTCon is NP-hard. To show membership of BTCon in the class NP we recall that BTCon is equivalent to 3DH2. Membership of 3DH2 in the class NP is quite immediate, a realizer can serve as certificate. The certificate can be checked efficiently.

The theorem below claims that BTCon remains NP-complete for bipartite input graphs with maximum degree at most 5 and arbitrarily large constant girth. Regarding the degree, it suffices to check that all the vertices involved in clause and variable gadgets have degree at most 5 (see Figure 4 and 7). Indeed, only center vertices of rotors are of degree 5 and degree 4 only occurs at gates in clause gadgets. Regarding the girth, we observe that every cycle of $G_\Phi^\bullet$ contains an edge that is subject to subdivisions. Hence the girth can be made arbitrarily large.

Theorem 2 *The recognition problems BTCon and 3DH2 are NP-complete. NP-hardness of the recognition is preserved for input graphs of arbitrarily large (constant) girth and maximum degree 5.*

4 A Sandwich Theorem

The proof of Theorem 2 in the previous section was based on 3 properties of triangle containment representations.

1) Validity of Lemma 1 implying that independent paths are non-crossing in any representation.

2) The restricted representability of an alternating rotor and the non-representability of an alternating 8-rotors (Lemma 2).

3) If G_Φ admits a schematized drawing without crossings and the path lengths are assigned properly, then G_Φ admits a triangle containment representation.

These three properties also hold for restricted classes of bipartite triangle containment representations. For example, we may require that the representing triangles are pairwise homothetic and of only two sizes, i.e., form two equivalence classes under translation. This is equivalent to the condition that the triangles representing one color class of the bipartite graph are points and triangles representing the other class are translates of a fixed triangle. This fixed triangle can be assumed to be equilateral with unit sidelength.

Point Unit-Triangle Containment (PUTCon)

Instance: A bipartite graph $G = (V, E)$.

Question: Does G admit a triangle containment representation with points and translates of an equilateral unit triangle.

The following theorem states that the decision problem PUTCon is hard. Even stronger: recognition of every class of graphs that is sandwiched between PUTCon and BTCon is also NP-hard.

Theorem 3 (Sandwich Theorem) *The recognition problems PUTCon and BTCon are NP-complete. Moreover:*

- *Recognition of every class $\mathcal{C}$ of bipartite graphs such that $\mathcal{C}$ contains all yes instances of PUTCon and $\mathcal{C}$ is contained in the set of all yes instances of BTCon is also NP-hard.*
- *NP-hardness of the recognition of $\mathcal{C}$ is preserved for input graphs of arbitrarily large (constant) girth and maximum degree at most 7.*

Proof. The idea for the proof is to encode a P-3-Con-3-SAT(4) instance Φ in a graph H_Φ such that:

(F) If there is no satisfying assignment for Φ, then there is no BTCon representation for H_Φ. Since $\mathcal{C} \subset$ BTCon this implies that $H_\Phi \notin \mathcal{C}$.

(T) If there is a satisfying assignment for Φ, then there is a PUTCon representation for H_Φ. Since PUTCon $\subset \mathcal{C}$ this implies that $H_\Phi \in \mathcal{C}$.

Property (F) is satisfied by the graph G_Φ from the previous section. Unfortunately the graph G_Φ has no PUTCon representation, the problem is that in the variable gadget we need that a path can bypass another path at a port of a rotor, see Figure 8(Left). To realize a bypass at each of two consecutive ports, we need triangles of at least three different sizes.

Clause gadgets with sufficiently long paths can be realized in the PUTCon model provided that one of the literals represents TRUE. The realization of the gate with points and unit triangles is shown in Figure 8(Right).

To make the approach work, we define H_Φ on the basis of our clause gadgets from Subsection 3.3 but with a new variable gadget. Since Property (F) depends on the clause gadget, it is also valid for H_Φ.

The new variable gadget has space to accommodate six occurrences of the variable in clauses. Since in our instances variables only have three or four occurrences, we use a *waste gadget* (see below) to get rid of some of them. The variable gadget superimposes two alternating rotors on the same two center vertices u and v. Let $r_1, .., r_6$ and $l_1, .., l_6$ be, in cyclical order, the six vertices of the first rotor and second rotor, respectively, that are connected to the center. However, for each i one of r_i and l_i is connected to u and the other to v. The paths starting at r_i, l_i form an incidence strand or they form a waste strand. The ordering of the strands around the variable is in accordance with the plane embedding of I_Φ.

For $i \neq j$ there are no crossings between any paths starting at l_i and r_j, hence, the two paths of an incidence strand have to connect adjacent ports of the rotor. Therefore there are just two possible ways in which the two rotors can interlace, see Figure 9. These two configurations correspond to the assignment of the values TRUE or FALSE to the variable. As in the case of the variable gadget from Subsection 3.4, the strands of the variable have to be connected to the strands of the clause by a proper choice of the pairing (cyan-green,pink-yellow) or (cyan-yellow,pink-green).

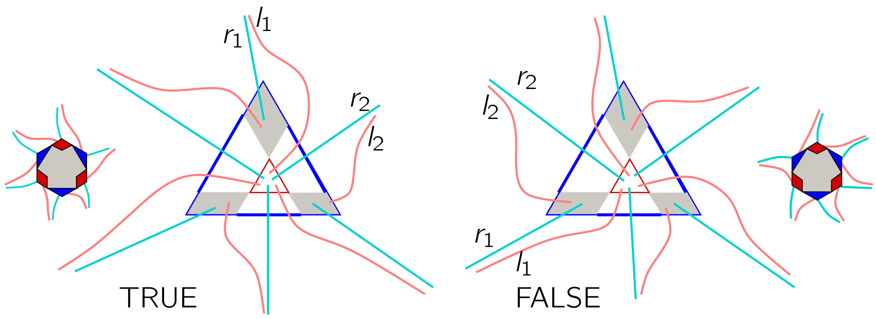

Figure 9: Two pictures for each of the two possible ways of interlacing the rotors of a variable gadget, as defined by the truth value. The small pictures use schematized rotors, while the larger pictures show the triangles of the center vertices of the rotors.

The *waste gadget* is designed to keep wasted strands in place between two incidence strands. For this purpose, they are connected to the neighboring incidence strands as shown in Figure 10.

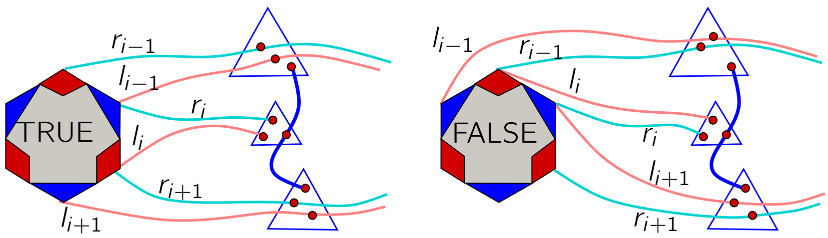

Figure 10: Two instances of a waste gadget. The hexagon belongs to a variable gadget.

The new variable gadget as well as the waste gadgets have PUTCon representations, hence, Property (T) is also true for H_Φ. The claim regarding the girth follows, because each cycle contains a path that can be subdivided accordingly, whereas for the maximum degree we only have to note that the maximum degree in the new variable and waste gadgets is 7. □

4.1 An application: bipartite triangle intersection representations

We use Theorem 3 to show hardness of recognizing bipartite graphs that admit a special type of triangle intersection representation.

Bipartite Triangle Intersection Representations (BTIntR)

Instance: A bipartite graph $G = (V, E)$ with bipartition X, Y.

Question: Are there families $\Delta_X = \{T_x : x \in X\}$ and $\Delta_Y = \{T^*_y : y \in Y\}$ of triangles such that

- Triangles in each of the families are pairwise homothetic, and there is a point reflection transforming T_x into T^*_y.
- $(x, y) \in E$ if and only if $T_x \cap T^*_y \neq \emptyset$.
- Within each of the two families Δ_X and Δ_Y there is no containment.

As a consequence of Proposition 4, we obtain that the problem BTIntR is equivalent to the question of whether a partial order P of height 2 admits an order preserving embedding $P \to \mathbb{R}^3$ such that there is a plane H in $\mathbb{R}^3$ that separates the minimal and the maximal elements of P. In particular, P has to be 3-dimensional. The partial orders corresponding to bipartite graphs in PUTCon clearly admit such a 'separated' embedding in $\mathbb{R}^3$, hence, yes instances of BTIntR form a class of graphs sandwiched between PUTCon and BTCon. With Theorem 3 this implies hardness. In Proposition 5 we show that BTIntR is in NP. Together this proves the following:

Theorem 4 *The decision problem BTIntR is NP-complete.*

The proof of Proposition 4 is based on representations of t-dimensional orders using simplices in $\mathbb{R}^{t-1}$. In Proposition 2 we have shown that t-dimensional orders are containment orders of homothetic simplices in $\mathbb{R}^{t-1}$.

Now consider a representation $\hat{X} = \{\hat{x} : x \in X\}$ of $P = (X, <_P)$ in $(\mathbb{R}^t, \leq)$ and let H be a hyperplane with normal vector $\mathbb{1}$ and an element $h_0 \in H$. Define the half-spaces $H^+ = \{y : \langle y - h_0, \mathbb{1}\rangle \geq 0\}$ and $H^- = \{y : \langle y - h_0, \mathbb{1}\rangle \leq 0\}$ and two subsets $X^+ = \{x \in X : \hat{x} \in H^+\}$ and $X^- = \{x \in X : \hat{x} \in H^-\}$ of X.

An element $x \in X^+$ is represented on H by the simplex $\Delta(x) = H \cap C(\hat{x})$, where $C(\hat{x})$ is the cone $\{p \in \mathbb{R}^t : p \leq \hat{x}\}$. The containment relation on $\{\Delta(x) : x \in X^+\}$ represents the induced order $P[X^+]$.

For elements of X^- we consider the dual cone $C^*(\hat{y}) = \{p \in \mathbb{R}^t : p \geq \hat{y}\}$. Note that the containment of dual cones represents the dual $P^*[X^-]$ of the suborder $P[X^-]$ of P induced by X^-, i.e., $C^*(\hat{x}) \subseteq C^*(\hat{y})$ if and only if $y \leq_P x$. The intersections of H with the dual cones yields a family homothetic regular simplices $\{\Delta^*(y) : y \in X^-\}$ whose containment order represents $P^*[X^-]$.

If $x \in X^+$ and $y \in X^-$, then $\Delta(x)$ and $\Delta^*(y)$ are both regular simplices, however, they are not homothetic, a point reflection is needed to get from one to the other. Comparabilities between $x \in X^+$ and $y \in X^-$ also have a nice description in the family $\{\Delta(x) : x \in X^+\} \cup \{\Delta^*(y) : y \in X^-\}$:

Claim. $x \in X^+$ and $y \in X^-$ are comparable if and only if $\Delta(x) \cap \Delta^*(y) \neq \emptyset$.

A comparability has to be of the form $y <_P x$. Hence, if $\overline{xy}$ is the segment with endpoints $\hat{x}$ and $\hat{y}$, then $\overline{xy} \subset C(\hat{x}) \cap C^*(\hat{y})$. It follows that $\overline{xy} \cap H$ is a point in the intersection $\Delta(x) \cap \Delta^*(y)$.

Conversely, if $p \in \Delta(x) \cap \Delta^*(y)$, then $\hat{y} \leq p \leq \hat{x}$, i.e., $y <_P x$. △

Figure 11 shows an order P with the triangle configurations obtained from different choices of H in an embedding of P into $\mathbb{R}^3$.

Based on these considerations the following certificate for dimension at most t is easily obtained.

Proposition 4 *Let $P = (X, \leq_P)$ be a partial order with a partition $X = X^- \cup X^+$ such that if $x \in X^-$ and $y \in X^+$ are a comparable pair, then $X < y$, i.e., X^- is a downset and X^+ is an upset in P. If there exist families of full-dimensional regular simplices $\{\Delta(x) : x \in X^+\}$ and $\{\Delta^*(y) : y \in X^-\}$ in $\mathbb{R}^{t-1}$, where the simplices in each of the families are pairwise homothetic and there is a point reflection transforming $\Delta(x)$ into $\Delta^*(y)$ such that*

- *$x \leftrightarrow \Delta(x)$ is an isomorphism between $P[X^+]$ and the containment order on $\{\Delta(x) : x \in X^+\}$*
- *$y \leftrightarrow \Delta^*(y)$ is an isomorphism between $P^*[X^-]$ and the containment order on $\{\Delta^*(y) : y \in X^-\}$,*
- *$x \in X^+$ and $y \in X^-$ are comparable if and only if $\Delta(x) \cap \Delta^*(y) \neq \emptyset$,*

then $\dim(P) \leq t$.

Proposition 5 *BTIntR and PUTCon are in NP.*

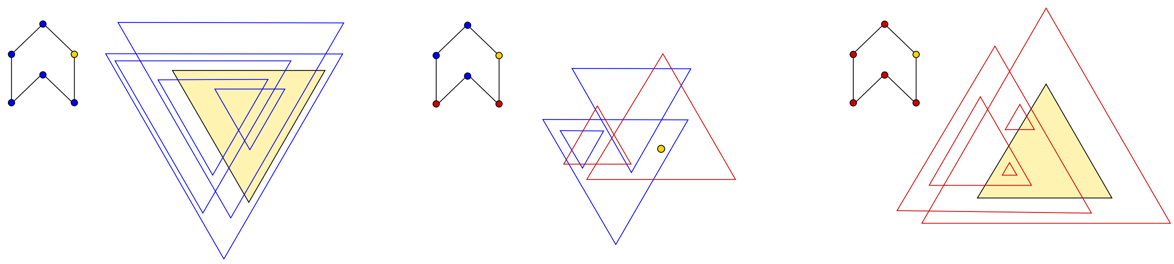

Figure 11: Three representations of the six element partial order P (the chevron) with triangles. The representations are based on the same embedding in $\mathbb{R}^3$ with different choices of the hyperplane H. Left: H is below all points $\hat{x}$ and we obtain a containment representation with triangles. Center: H contains $\hat{x_0}$, where x_0 is the element that is distinguished in all three representations. The simplex of x_0 is degenerated to a point elements of X^+ are blue and elements of X^- are red. Right: H is above all points $\hat{x}$ and the containment of triangles represents the dual order P^*.

Proof. Kratochvíl and Matoušek [15] showed that the class of segment intersection graphs where the segments of the representation are restricted to the elements of a fixed set K of directions with $|K| = k$ is in NP for every k. Our proof is based on their ideas. As certificate of a yes instance for the BTIntR problem we propose a combinatorial description of the intersection pattern of the triangles. In this description each vertex comes with three lists, each of them representing the order of intersections along a side of the triangle. To check whether the certificate corresponds to a BTInt/PUTCon representation, we can use a linear program to test whether the combinatorial description can be "stretched" to a set of homothetic triangles with the same intersection pattern.

The requirement that the intersection of side s with q immediately precedes its intersection with r is encoded in a linear inequality: Consider the equations $a_s x + b_s = y$, $a_r x + b_r = y$, and $a_q x + b_q = y$ of the supporting lines and note that the a_i are slopes, i.e., they are constants and can be taken from the fixed 3-element set $\{-1, 0, 1\}$, the b_i are the variables. The x coordinate of the intersection of s and r is $\frac{b_s - b_r}{a_s - a_r}$. The condition that this is smaller than the x coordinate of $s \cap q$ is captured by the strict inequality $(b_s - b_r)(a_s - a_q) < (b_s - b_q)(a_s - a_r)$. Introducing some constant c for the min distance, we obtain the inequality $(b_s - b_r)(a_s - a_q) + c \leq (b_s - b_q)(a_s - a_r)$ in the variables b_s, b_r, b_q. In the case of PUTCon we additionally need that the distance between two corners of a triangle realizes a prescribed value. Since corners also are intersections of segments the same technique applies.

From the system of equations we make linear program by adding the objective max c. The decision whether the linear program has a solution with objective $c > 0$ can be obtained in polynomial time. □

5 Open questions and extensions

We have shown that maximum degree at least 5 is enough to make BTCon hard. From Schnyder's work [20] we know that bipartite graphs where the degree in one of the color classes is 2 are yes instances for BTCon if and only if they are incidence orders of planar graphs. What about maximum degree 3?

- What is the complexity of deciding whether a bipartite graph of maximum degree 3 admits a BTCon representation?

We can also restrict the class of inputs to planar bipartite graphs. It is known that the incidence order of vertices and faces of a 3-connected planar graph is of dimension 4, see [3], moreover there are outerplanar graphs whose incidence order of vertices and faces is of dimension 4, see [8]. Is it hard to decide whether a planar bipartite graph is of dimension 3? Or in terms of triangle containments:

- What is the complexity of deciding whether a planar bipartite graph admits a BTCon representation?

From Schnyder's proof we know that incidence orders of planar graphs have BTInt representations.

- What is the complexity of deciding whether an incidence orders of planar graphs (a subdivision of a planar graph) admits a PUTCon representation?

Constructions with rotors with larger or smaller numbers of corners may yield other NP-hardness proofs. An example is given by the proof of hardness of recognition of unit grid intersection graphs (UGIG) in [19] where rotors have four corners. In fact, the paper shows hardness of recognition for all classes that are sandwiched between UGIG and pseudosegment intersection graphs.

We invite the reader to adapt the rotors and the gadgets to show that intersection graphs of polygonal regions with k corners are hard to recognize for every fixed k. This has already been shown by Kratochvíl [14], however, his proof does not show hardness for inputs of large girth.

References

[1] A. Adiga, J. Babu, and S. L. Chandran, *Polynomial time and parameterized approximation algorithms for boxicity*, in Parameterized and Exact Computation, LNCS, vol. 7535, Springer, 2012, 135–146.

[2] G. Brightwell and E. Scheinerman, *Fractional dimension of partial orders*, Order, 9 (1992), 139–158.

[3] G. Brightwell and W. T. Trotter, *The order dimension of convex polytopes*, SIAM J. Discrete Math., 6 (1993), 230–245.

[4] P. Chalermsook, B. Laekhanukit, and D. Nanongkai, *Graph products revisited: Tight approximation hardness of induced matching, poset dimension and more*, in SODA, 2013.

[5] B. Dushnik and E. Miller, *Partially ordered sets*, American Journal of Mathematics, 63 (1941), 600–610.

[6] S. Even, A. Pnueli, and A. Lempel, *Permutation graphs and transitive graphs*, J. of the ACM, 19 (1972), 400–410.

[7] S. Felsner, *Geodesic embeddings and planar graphs*, Order, 20 (2003), 135–150.

[8] S. Felsner and J. Nilsson, *On the dimension of outerplanar maps*, Order, 28 (2011), 415–435.

[9] S. Felsner and W. T. Trotter, *On the fractional dimension of partially ordered sets*, Discrete Mathematics, 136 (1994), 101–117.

[10] S. Felsner and F. Zickfeld, *Schnyder woods and orthogonal surfaces*, Discrete & Computational Geometry, 40 (2008), 103–126.

[11] M. R. Garey and D. S. Johnson, *Computers and Intractability, a Guide to the Theory of NP-Completeness*, W. H. Freeman, San Francisco, 1979.

[12] R. Hegde and K. Jain, *The hardness of approximating poset dimension*, Electronic Notes in Discrete Mathematics, 29 (2007), 435–443.

[13] H. Kierstead and S. Penrice, *Computing the dimension of N-free ordered sets is NP-complete*, Order, 6 (1989), 133–135.

[14] J. Kratochvíl, *A special planar satisfiability problem and a consequence of its NP-completeness*, Discrete Applied Mathematics, 52 (1994), 233–252.

[15] J. Kratochvíl and J. Matoušek, *Intersection graphs of segments*, Journal of Combinatorial Theory, Series B, 62 (1994), 289 – 315.

[16] D. Lichtenstein, *Planar formulae and their uses*, SIAM J. Comput., 11 (1982), 329–343.

[17] R. M. McConnell and J. Spinrad, *Linear-time transitive orientation*, in SODA, 1997, 19–25.

[18] E. Miller, *Planar graphs as minimal resolutions of trivariate monomial ideals*, Documenta Math., 7 (2002), 43–90.

[19] I. Mustaţă and M. Pergel, *Unit grid intersection graphs: Recognition and properties*, CoRR, abs/1306.1855 (2013).

[20] W. Schnyder, *Planar graphs and poset dimension*, Order, 5 (1989), 323–343.

[21] J. Spinrad, *Dimension and algorithms*, in Orders, Algorithms, and Applications, LNCS, vol. 831, Springer, 1994, 33–52.

[22] W. T. Trotter, *Progress and new directions in dimension theory for finite partially ordered sets*, in Extremal problems for finite sets, Bolyai Society, 1991, 457–477.

[23] W. T. Trotter, *Combinatorics and Partially Ordered Sets: Dimension Theory*, Johns Hopkins Series in the Mathematical Sciences, The Johns Hopkins University Press, 1992.

[24] W. T. Trotter, *Graphs and partially ordered sets: Recent results and new directions*, Congr. Numerantium, 116 (1996), 253–278.

[25] M. Yannakakis, *The complexity of the partial order dimension problem*, SIAM J. Algebraic Discrete Methods, 3 (1982), 351–358.