Hardness and modifications of the weak graph distance

Maike Buchin[1] and Wolf Kißler[1]

1 **Faculty of Computer Science, Ruhr-Universität Bochum**
`{maike.buchin, wolf.kissler}@rub.de`

Abstract

The weak graph distance is a distance measure for immersed graphs. We extend previous NP-hardness results for deciding this distance. Also, we present variations that we conjecture to be fixed-parameter tractable under regularity assumptions when parameterized in the number of crossings.

1 Introduction

Embedded and immersed graphs are widely used natural representations for many kinds of geometric networks. Given multiple models of the same network or representations of related networks, one is typically interested in comparing the models. In recent years, many different distance measures for embedded and immersed graphs have been proposed, cf. [2].

Two such distance measures are the strong and weak graph distance proposed by Akitaya et al. [1], which are based on the strong and weak Fréchet distance for polygonal curves, respectively. Both distance measures are metrics, cf. [2]. A key advantage of these measures is that they capture both geometric and topological (dis)similarity. As discussed in [1], first experiments on reconstructions of real road networks showed promising results.

The strong graph distance is NP-complete to approximate within a 1.10566 ratio even on plane graphs. The best known exact algorithm due to [1] runs in an XP-like time bound when parameterized in the number of faces.

For the weak graph distance, there is a quadratic-time decision algorithm on spike-free (i.e., cycles are embedded in a nice way) plane graphs. Akitaya et al. also showed that when both graphs are immersed in $\mathbb{R}^2$, the weak graph distance is NP-complete to decide.

Hence we are interested in whether (a variant of) the weak graph distance is tractable on realistic networks, in particular those with few edge crossings. For this, we first extend the hardness result of [1] in showing that deciding the directed weak graph distance remains NP-complete even if the source graph is plane. Moreover, we show that deciding the directed distance is NP-complete for G_1, G_2 embedded in $\mathbb{R}^d$ for $d \geq 3$. In both scenarios, constant-factor approximation is NP-complete as well.

Then we suggest the family of crossing-rigid weak graph distances as alternative distance measures. Under reasonable regularity assumptions, we conjecture that these measures admit fixed-parameter tractable algorithms when parameterized in the numbers of crossings.

1.1 The weak graph distance

Here, we introduce relevant notation from [1]. First, we recall the weak Fréchet distance:

▶ **Definition 1.1.** Let $s_1, s_2 \colon [0,1] \to \mathbb{R}^2$ be curves. Define their *weak Fréchet distance* by

$$\delta_{wF}(s_1, s_2) := \inf_{\alpha,\beta} \max_{t\in[0,1]} d(s_1(\alpha(t)), s_2(\beta(t))),$$

40th European Workshop on Computational Geometry, Ioannina, Greece, March 13–15, 2024.
This is an extended abstract of a presentation given at EuroCG'24. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

where α, β range over all continuous self-surjections of $[0,1]$ that keep the endpoints fixed and d is the standard Euclidean metric.

The weak graph distance is defined for embedded and immersed graphs. We use the terms embedded and immersed in the topological sense, that is, an embedding is (essentially) a crossing-free drawing in $\mathbb{R}^d$ and an immersion is any drawing that may also contain crossings. Moreover, we use the term plane graph for embeddings of planar graphs in $\mathbb{R}^2$.

In the following, let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be graphs immersed in $\mathbb{R}^2$ using straight-line immersions. We will slightly abuse notation and refer to the immersions of graphs, edges and vertices using the same notation as for the abstract graphs.

To define the weak graph distance, we first define graph mappings:

▶ **Definition 1.2.** A *graph mapping* $s\colon G_1 \to G_2$ is a map that maps

1. each vertex $v \in V_1$ to a point $s(v)$ on an edge of G_2 and
2. each edge $\{u, v\} \in E_1$ to a simple path from $s(u)$ to $s(v)$ in G_2.

The weak graph distance is now defined as the maximum weak Fréchet distance between an edge and its image under a (globally) optimal graph mapping:

▶ **Definition 1.3.** For immersed graphs G_1, G_2, define the *directed weak graph distance* via

$$\vec{\delta}_{wG}(G_1, G_2) := \min_{s:G_1 \to G_2} \max_{e \in E_1} \delta_{wF}(e, s(e)),$$

where s ranges over all graph mappings and e and $s(e)$ refer to the corresponding immersions as curves in $\mathbb{R}^2$. The *(undirected) weak graph distance* between G_1 and G_2 is defined as

$$\delta_{wG}(G_1, G_2) := \max(\vec{\delta}_{wG}(G_1, G_2), \vec{\delta}_{wG}(G_2, G_1)).$$

Lastly, we outline the general decision alg. described in [1]. For that, we define placements:

▶ **Definition 1.4.** An ε-placement of a vertex v is a connected component (w.r.t. the canonical topologization of G_2 as a simplicial complex) of $G_2 \cap B_\varepsilon(v)$. A *weak edge placement* of an edge $e = \{u, v\} \in E_1$ is a path P in G_2 that connects placements of u and v, respectively, such that $\delta_{wF}(e, P) \leq \varepsilon$. A *weak ε-placement* of G_1 is a graph mapping $s\colon G_1 \to G_2$ that maps each edge to a weak ε-placement.

Furthermore, we call a vertex placement C_v *weakly valid* if each adjacent vertex u has a placement C_u such that C_v and C_u are connected by a weak ε-placement of $\{u, v\}$. Otherwise, we call the placement *weakly invalid.*

The general decision algorithm now proceeds as follows:

Algorithm 1 General Decision Algorithm [1]

1: Compute vertex placements.
2: Compute mutual reachability information for vertex placements.
3: Prune invalid placements.
4: Decide if there exists a placement for the whole graph G_1.

As described in [1], steps 1-3 can be performed in quadratic time for general immersed graphs. However, existence of a weakly valid placement for each vertex does not imply existence of a weak placement of the whole graph, cf. Fig. 1. Thus, step 4 is non-trivial in general. In [1], it is shown that step 4 is in fact trivial if both graphs are plane and the embedding of G_1 meets the following regularity condition:

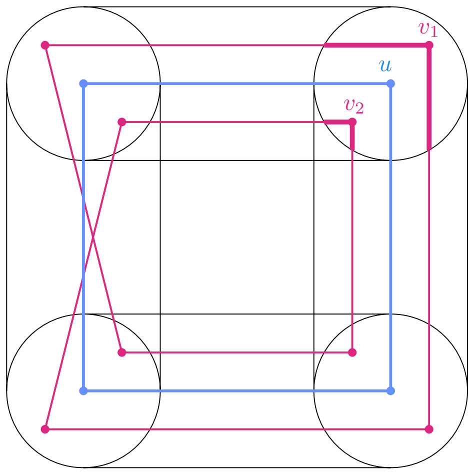

Figure 1 Graphs G_1 (blue) and G_2 (red) such that each vertex of G_1 has exactly two weakly valid placements, but no weak placement for G_1 exists.

▶ **Definition 1.5.** An immersed graph G is called *spike-free* if each cycle C of G is 2ε-thick and, for each three consecutive vertices $u, v, w \in C$, the ε-ball around u does not intersect the ε-tube around the edge $\{v, w\}$.

2 Hardness of deciding the weak graph distance

Here, we extend the hardness result from [1] to show that deciding the directed weak graph distance remains NP-hard even if the source graph is plane. Our proof idea resembles their original proof idea. However, their reduction is from binary CSP, where we cannot guarantee planarity of the resulting graphs. Instead, we reduce from 3-colorability of planar graphs.

▶ **Theorem 2.1.** *Deciding whether $\vec{\delta}_{wG}(G_1, G_2) \leq \varepsilon$ is NP-complete even if G_1 is plane.*

Proof idea. A more detailed version of this proof will be included in the full version due to space restrictions. We reduce from planar 3-coloring, which is NP-complete due to [6]. Given a planar graph $G = (V, E)$, we construct an embedding of G on a grid in linear time, cf. [3]. We construct a graph G_c by placing vertices $w_{v,i}$ in ε-balls around each $v \in V$ for $i \in \{1, 2, 3\}$. For adjacent vertices u, v in G, G_c has an edge $e_{u,v,i,j}$ connecting $w_{u,i}$ to $w_{v,j}$ iff $i \neq j$ and $u \prec v$ for some fixed linear order $\preceq$ on V.

By choosing ε sufficiently small, we can achieve that each $u \in V$ has exactly three placements in G_c corresponding to the $w_{u,i}$. Our idea is that placing u onto the placement corresponding to $w_{u,i}$ is comparable to coloring u with color i. However, since our distance measure is based on the weak Fréchet distance, multiple edges might be used to connect same colored placements of adjacent vertices.

This can be prevented by inserting a vertex in the middle of each edge of G. Denote the resulting immersed graph by G_s. Then, all edges of G must be placed essentially through some $e_{u,v,i,j}$, which exists iff $i \neq j$. Thus, a consistent ε-placement of G_s onto G_c must use globally consistent $e_{u,v,i,j}$, implying that G is 3-colorable. See Fig. 2 for an illustration. ◀

Starting with 4-regular planar graphs instead, the problem remains NP-complete due to [4] and we obtain slightly stronger results. Those observations will be included in the full version due to space restrictions, as well as detailed proofs of the following results:

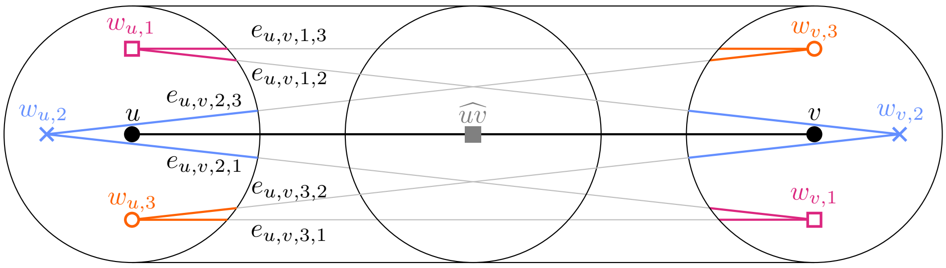

Figure 2 Illustration of the reduction in Thm. 2.1 on a single edge $\{u, v\}$ of the input graph.

▶ **Corollary 2.2.** *It is NP-hard to approximate $\vec{\delta}_{wG}(G_1, G_2)$ within any constant ratio $c \geq 1$ even if G_1 is embedded in $\mathbb{R}^2$.*

Proof idea. In the above construction, place the vertices of G_c within $\frac{\varepsilon}{c}$-balls instead. ◀

▶ **Theorem 2.3.** *The (directed) weak graph distance is NP-hard to approximate within any constant ratio $c \geq 1$ for graphs G_1, G_2 embedded in $\mathbb{R}^d$ for all $d \geq 3$.*

Proof idea. In a similar construction, embed the vertices of G, G_c on the moment curve. [5, Lemma 5.4.2] implies planarity. Approximation hardness is analogous to Cor. 2.2. ◀

3 Crossing-rigid weak graph distances

3.1 Definitions and properties

As seen in the previous results, the directed weak graph distance is NP-hard to decide if we do not restrict the crossings of G_2. Although the search space can be shown to have subexponential size in various cases, it is currently unknown whether there exists an FPT decision algorithm parameterized in the number of crossings of G_2 for the general case.

Hence, we propose modifying the distance measure by requiring crossings to be mapped onto crossings. This allows us to design FPT algorithms, and also captures the intuition that if two immersed graphs describe similar networks, they should have similar crossings.

In our notation, a crossing is a tuple (e, p) of an edge e and a point $p \in \mathbb{R}^2$ in which the immersion of e crosses the immersion of (at least) one other edge. First, we formalize the notion of mapping crossings onto crossings:

▶ **Definition 3.1.** Let $s\colon G_1 \to G_2$ be a graph mapping. We say that s is *loosely crossing-rigid* if s maps each edge $e = \{u, v\}$ that has crossings in points $p_1, \ldots, p_n$ to a sequence of (possibly constant) paths $P_0, \ldots, P_n$ in G_2 such that

1. all initial and terminal points of the P_i are $s(u), s(v)$ or crossings of G_2,
2. the P_i visit no crossings except for their initial and terminal points and
3. the concatenated path $P_0 P_1 \ldots P_n$ is defined and is a simple path from $s(u)$ to $s(v)$

s is *crossing-rigid* if there exists such a sequence such that P_0 and P_n are not constant. s is *strictly crossing-rigid* if there exists such a sequence such that none of the P_i are constant.

In other words: For edges without crossings, nothing is changed. For an edge e that has $n \geq 1$ crossings, the image of e

- under a loosely crossing-rigid graph mapping has at most n crossings,
- under a crossing-rigid graph mapping has at least one and at most n crossings,
- under a strictly crossing-rigid graph mapping has exactly n crossings.

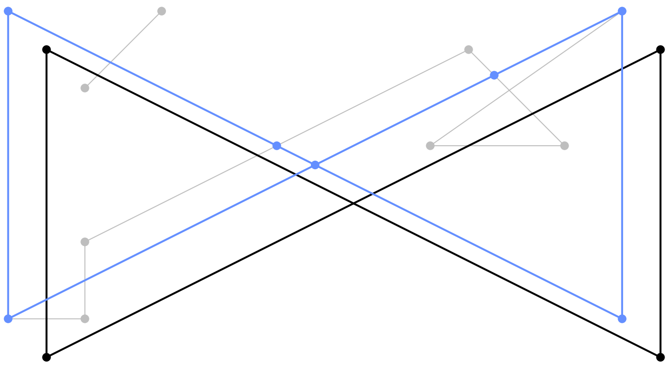

Figure 3 Graphs G_1 (black) and G_2 (blue and grey). The blue part is a valid image for G_1 under a crossing-rigid graph mapping. The diagonal grey path may not be an image of an edge of G_1 since it passes 3 crossings (the upper right crossing is passed through two edges).

As such, crossing-rigidity is of purely combinatorial rather than geometric nature.

▶ **Definition 3.2.** For immersed graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, the *directed crossing-rigid weak graph distance* is defined as

$$\vec{\delta}_{crwG}(G_1, G_2) = \inf_{s\colon G_1 \to G_2} \max_{e \in E_1} \delta_{wF}(e, s(e))$$

where s ranges over all crossing-rigid graph mappings $s\colon G_1 \to G_2$ and e and $s(e)$ are interpreted as the corresponding polygonal curves.

The *directed loosely crossing-rigid weak graph distance* $\vec{\delta}^{l}_{crwG}(G_1, G_2)$ and the *directed strictly crossing-rigid weak graph distance* $\vec{\delta}^{s}_{crwG}(G_1, G_2)$ are defined analogous.

Respective undirected versions $\delta_{crwG}, \delta^{l}_{crwG}, \delta^{s}_{crwG}$ can be defined as in Def. 1.3.

Note that when G_1 has crossings and G_2 is plane, there exist no (strictly) crossing-rigid graph mappings and as such, $\vec{\delta}_{crwG}(G_1, G_2) = \vec{\delta}^{s}_{crwG}(G_1, G_2) = \infty$. Moreover, since placements are no longer compact, we might have $\vec{\delta}_{crwG}(G_1, G_2) = \varepsilon$ (analogous for the loosely or strictly crossing-rigid versions) even if no ε-placement exists.

▶ **Observation 3.3.** *Without further restrictions, the above distance measures have several counterintuitive properties:*

1. *s may map crossing edges e_1 and e_2 such that $s(e_1)$ does not cross $s(e_2)$. Even if $s(e_1)$ and $s(e_2)$ cross, they need not cross in the corresponding crossings from Def. 3.1.*
2. *The implicit mapping $(e, p) \mapsto (e', p')$ of crossings of G_1 onto crossings of G_2 need not be one-to-one even for the strict distance.*
3. *A crossing p of an edge e may be mapped onto an edge crossing p' such that $d(p, p') > \varepsilon$.*

See Figs. 3 and 4 for illustrations.

However, our FPT approach is to assign a crossing of G_2 to each crossing of G_1 and for a fixed assignment, the above properties may be decided efficiently. Thus, we may require the mappings to not have any subset of the above properties without impacting tractability.

▶ **Lemma 3.4.** *Let G_1, G_2 be graphs immersed in $\mathbb{R}^2$. It holds that*

$$\vec{\delta}_{wG}(G_1, G_2) \le \vec{\delta}^{l}_{crwG}(G_1, G_2) \le \vec{\delta}_{crwG}(G_1, G_2) \le \vec{\delta}^{s}_{crwG}(G_1, G_2). \quad (1)$$

If G_1 and G_2 are plane, $\vec{\delta}_{wG}(G_1, G_2) = \vec{\delta}^{l}_{crwG}(G_1, G_2) = \vec{\delta}_{crwG}(G_1, G_2) = \vec{\delta}^{s}_{crwG}(G_1, G_2)$.

▶ Remark. None of the ratios between the terms of eq. 1 are bounded.

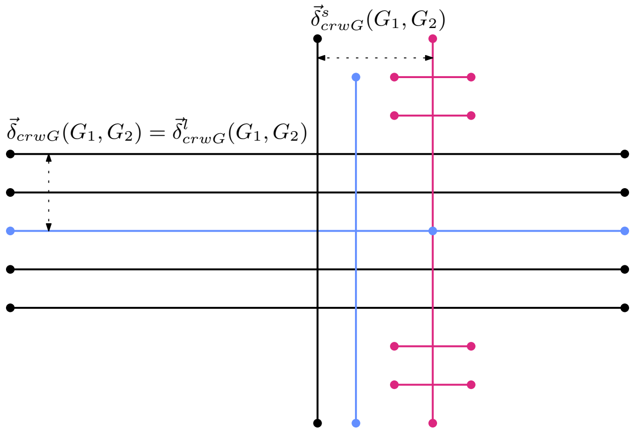

Figure 4 Graphs G_1 (black) and G_2 (blue and red). For the (loosely) crossing-rigid distance, G_1 can be mapped onto the blue edges. For the strict distance, the vertical edge needs to be mapped onto the red vertical edge, resulting in a larger distance.

3.2 Decision algorithm

Note that we have to amend the definitions in Def. 1.4: Placements of a vertex v are now connected components of $G_2 \backslash C(G_2)$ within $B_\varepsilon(v)$, where $C(G_2)$ is the set of points in $\mathbb{R}^2$ in which G_2 has crossings. Edge placements now have to adhere to Def. 3.1. We propose the following algorithmic approach for the crossing-rigid distances:

Algorithm 2 Decision algorithm for the existence of a crossing-rigid weak ε-placement

1: If not given, compute where the immersion of G_1 resp. G_2 has crossings.
2: Compute vertex placements.
3: **for each** valid assignment of crossings of G_2 to crossings of G_1
4: Compute reachability information for vertex placements under current assignment.
5: Prune invalid placements.
6: **if** there exists a placement for the whole graph G_1 **then return** true.
7: **return** false.

The definition of "validity" of an assignment in step 3 depends on whether we consider the loosely crossing-rigid, crossing-rigid or strictly crossing-rigid distance. Less assignments will be valid if we demand that the graph mapping does not have some of the properties from Obs. 3.3. There are at most $(k_2 + 1)^{k_1}$ such assignments. Restricting the assignments in the sense of Def. 3.1 or Obs. 3.3 takes polynomial time per assignment.

Steps 1, 2, 4 and 5 can be performed in polynomial time. Regarding step 6, we conjecture:

▶ **Conjecture 3.5.** *For the crossing-rigid and strictly crossing-rigid weak graph distance, step 6 can be performed in polynomial time if G_1 is spike-free. For the loosely crossing-rigid weak graph distance, step 6 can be performed in polynomial time if G_1 satisfies some slightly stronger regularity condition.*

Essentially, the idea is to have a similar situation to the plane case from [1, Lemma 7] for edges without crossings. For edges with crossings, after assigning crossings, all consistent weakly valid placements of the incident vertices are mutually reachable.

Verifying the above conjecture and developing a computation algorithm are natural next steps. Additionally, as mentioned above, it is currently unknown whether the weak graph distance admits an FPT algorithm when parameterized in the respective number of crossings, which is also an interesting question. Lastly, more experimental work would give insight into the use of our distance measures for comparing realistic networks.

References

1 Hugo A. Akitaya, Maike Buchin, Bernhard Kilgus, Stef Sijben, and Carola Wenk. Distance measures for embedded graphs. *Computational Geometry*, 95:101743, 2021.

2 Maike Buchin, Erin Chambers, Pan Fang, Brittany Terese Fasy, Ellen Gasparovic, Elizabeth Munch, and Carola Wenk. Distances between immersed graphs: Metric properties. *La Matematica*, pages 1–26, 2023.

3 Marek Chrobak and Thomas H. Payne. A linear-time algorithm for drawing a planar graph on a grid. *Information Processing Letters*, 54(4):241–246, 1995.

4 David P. Dailey. Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete. *Discrete Mathematics*, 30(3):289–293, 1980.

5 Jiri Matousek. *Lectures on discrete geometry*, volume 212. Springer, 2013.

6 Larry Stockmeyer. Planar 3-colorability is polynomial complete. *ACM Sigact News*, 5(3):19–25, 1973.