Recognizing Weak Embeddings of Graphs

HUGO A. AKITAYA, Tufts University, USA

RADOSLAV FULEK, Institute of Science and Technology, Austria

CSABA D. TÓTH, California State University Northridge, Tufts University, USA

We present an efficient algorithm for a problem in the interface between clustering and graph embeddings. An **embedding** $\varphi : G \rightarrow M$ of a graph G into a 2-manifold M maps the vertices in $V(G)$ to distinct points and the edges in $E(G)$ to interior-disjoint Jordan arcs between the corresponding vertices. In applications in clustering, cartography, and visualization, nearby vertices and edges are often bundled to the same point or overlapping arcs due to data compression or low resolution. This raises the computational problem of deciding whether a given map $\varphi : G \rightarrow M$ comes from an embedding. A map $\varphi : G \rightarrow M$ is a **weak embedding** if it can be perturbed into an embedding $\psi_\varepsilon : G \rightarrow M$ with $\|\varphi - \psi_\varepsilon\| < \varepsilon$ for every $\varepsilon > 0$, where $\|\cdot\|$ is the uniform norm.

A polynomial-time algorithm for recognizing weak embeddings has recently been found by Fulek and Kynčl. It reduces the problem to solving a system of linear equations over $\mathbb{Z}_2$. It runs in $O(n^{2\omega}) \leq O(n^{4.75})$ time, where $\omega \in [2, 2.373]$ is the matrix multiplication exponent and n is the number of vertices and edges of G . We improve the running time to $O(n \log n)$. Our algorithm is also conceptually simpler: We perform a sequence of *local operations* that gradually “untangles” the image $\varphi(G)$ into an embedding $\psi(G)$ or reports that φ is not a weak embedding. It combines local constraints on the orientation of subgraphs directly, thereby eliminating the need for solving large systems of linear equations.

CCS Concepts: • **Mathematics of computing** → **Graphs and surfaces**; • **Theory of computation** → **Computational geometry**;

Additional Key Words and Phrases: Graph embedding, graph operations

ACM Reference format:

Hugo A. Akitaya, Radoslav Fulek, and Csaba D. Tóth. 2019. Recognizing Weak Embeddings of Graphs. *ACM Trans. Algorithms* 15, 4, Article 50 (October 2019), 27 pages.

<https://doi.org/10.1145/3344549>

1 INTRODUCTION

Given a graph G and a 2-dimensional manifold M , one can decide in linear time whether G can be embedded into M [19], although finding the smallest genus of a surface into which G embeds is NP-hard [24]. An **embedding** $\psi : G \rightarrow M$ is a continuous piecewise linear injective map where the graph G is considered as a 1-dimensional simplicial complex. Equivalently, an embedding

Research supported in part by the NSF awards CCF-1422311 and CCF-1423615, and the Science Without Borders program. The second author gratefully acknowledges support from Austrian Science Fund (FWF): M2281-N35.

Authors’ addresses: H. A. Akitaya, Tufts University, 161 College Ave, Medford, MA, 02155, USA; email: hugo.alves_akitaya@tufts.edu; R. Fulek, Institute of Science and Technology, Am Campus 1, 3400 Klosterneuburg, Austria; email: radoslav.fulek@gmail.com; C. D. Tóth, California State University Northridge, 18111 Nordhoff St, Los Angeles, CA 91330-8313, USA; email: cdtoth@acm.org.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1549-6325/2019/10-ART50 \$15.00

<https://doi.org/10.1145/3344549>

Fig. 1. (a) An embedding of $H = K_5$ in the torus. (b) Strip system $\mathcal{H}$ of the embedding of H . (c) A weak embedding where G is disconnected and $H = C_4$. (d) A negative instance where $G = C_8$ and $H = C_3$.

maps the vertices into distinct points and the edges into interior-disjoint Jordan arcs between the corresponding vertices that do not pass through the images of vertices. We would like to decide whether a given map $\varphi : G \rightarrow M$ can be “perturbed” into an embedding $\psi : G \rightarrow M$. Let M be a 2-dimensional manifold equipped with a metric. A continuous piecewise linear map $\varphi : G \rightarrow M$ is a **weak embedding** if, for every $\varepsilon > 0$, there is an embedding $\psi_\varepsilon : G \rightarrow M$ with $\|\varphi - \psi_\varepsilon\| < \varepsilon$, where $\|\cdot\|$ is the uniform norm (i.e., sup norm).

In some cases, it is easy to tell whether $\varphi : G \rightarrow M$ is a weak embedding: Every embedding is a weak embedding; and if φ maps two edges into Jordan arcs that cross transversely, then φ is not a weak embedding. The problem becomes challenging when φ maps several vertices (edges) to the same point (Jordan arc), although no two edges cross transversely. This scenario arises in applications in clustering, cartography, and visualization, where nearby vertices and edges are often bundled to the same point or overlapping arcs due to data compression, graph semantics, or low resolution. A *cluster* in this context is a subgraph of G mapped by φ to the single point in M .

The recognition of weak embeddings turns out to be a purely combinatorial problem independent of the global topology of the manifold M (as noted in Reference [6]). The key observation here is that we are looking for an embedding in a small neighborhood of the image $\varphi(G)$, which can be considered as the embedding of some graph H . As such, we can replace M with a neighborhood of an embedded graph in the formulation of the problem.

Problem Statement and Results. In Sections 2–5, we make two simplifying assumptions: We assume that M is an **orientable** 2-manifold, and $\varphi : G \rightarrow H$ is a *simplicial map* for some graph H embedded in M . A map $\varphi : G \rightarrow H$ is **simplicial** if it maps the vertices of G to vertices of H and the edges of G to edges or vertices of H such that incidence relations are preserved (in particular, the image of an edge cannot pass through the image of a vertex). We explain below how to drop the assumption that $\varphi : G \rightarrow H$ is simplicial at the expense of increasing the running time of our algorithms. We extend our results to nonorientable surfaces in Section 6.

An **embedded graph** H in an orientable 2-manifold M is an abstract graph together with a **rotation system** that specifies, for each vertex of H , the ccw cyclic order of incident edges. The **strip system** $\mathcal{H}$ of H (a.k.a. the *thickening* of H) is a 2-manifold with boundary constructed as follows (see Figure 1(a)–(b)): For every $u \in V(H)$, create a topological disk D_u , and for every edge $uv \in E(H)$, create a rectangle R_{uv} . For every D_u and R_{uv} , fix an arbitrary orientation of the boundaries ∂D_u and ∂R_{uv} , respectively. Partition ∂D_u into $\deg(u)$ arcs, and label them by $A_{u,v}$, for all $uv \in E(H)$, in the cyclic order around ∂D_u determined by the rotation of u in the rotation system of H . Finally, the manifold $\mathcal{H}$ is obtained by identifying two opposite sides of every rectangle R_{uv} with $A_{u,v}$ and $A_{v,u}$ via an orientation-preserving homeomorphism (i.e., consistently with the chosen orientations of ∂R_{uv} , ∂D_u , and ∂D_v). If $\varphi : G \rightarrow M$ is a continuous piecewise linear map, then $\varphi(G)$ is the embedding of some graph H in M , and an ε -neighborhood

of $\varphi(G)$ in M is homeomorphic to $\mathcal{H}$ for a sufficiently small $\varepsilon > 0$. Consequently, $\varphi : G \rightarrow M$ is a weak embedding if and only if $\varphi : G \rightarrow \mathcal{H}$ is a weak embedding. We may further assume that ψ_ε maps every edge uv to a Jordan arc that crosses the boundaries of D_u (resp., D_v) precisely once (cf. Reference [6, Lemma B2] and Reference [15, Section 4]).

We can now formulate an instance of the weak embeddability problem as a simplicial map $\varphi : G \rightarrow H$ (for short, φ), where G is an abstract graph and H is an embedded graph. The simplicial map $\varphi : G \rightarrow H$ is a **weak embedding** if there exists an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$ that maps each vertex $v \in V$ to a point in $D_{\varphi(v)}$, and each edge $uv \in E(G)$ to a Jordan arc in $D_{\varphi(u)} \cup R_{\varphi(u)\varphi(v)} \cup D_{\varphi(v)}$ that has a connected intersection with each of $D_{\varphi(u)}$, $R_{\varphi(u)\varphi(v)}$, and $D_{\varphi(v)}$, and $R_{\varphi(u)\varphi(v)} = \emptyset$ if $u = v$. We say that the embedding ψ_φ **approximates** φ . Our main results are the following.

THEOREM 1.1.

- (i) *Given an abstract graph G with m edges, an embedded graph H , and a simplicial map $\varphi : G \rightarrow H$, we can decide in $O(m \log m)$ time whether φ is a weak embedding.*
- (ii) *If $\varphi : G \rightarrow H$ is a weak embedding, then for every $\varepsilon > 0$, we can also find an embedding $\psi_\varepsilon : G \rightarrow M$ with $\|\varphi - \psi_\varepsilon\| < \varepsilon$ in $O(m \log m)$ time.*

Throughout the article, we assume that G has n vertices and m edges. In the plane (i.e., $M = \mathbb{R}^2$), only planar graphs admit weak embeddings, hence $m = O(n)$, but our techniques work for 2-manifolds of arbitrary genus, and G may be a dense graph. Our result improves the running time of the previous algorithm [15] from $O(m^{2\omega}) \leq O(m^{4.75})$ to $O(m \log m)$, where $\omega \in [2, 2.373]$ is the matrix multiplication exponent. It also improves the running times of several recent polynomial-time algorithms in special cases, e.g., when the embedding of G is restricted to a given isotopy class [14] and H is a path [3] (see below).

Extension to Nonsimplicial Maps and Nonorientable Manifolds. If $\varphi : G \rightarrow H$ is a continuous map (not necessarily simplicial) that is injective on the edges (each edge is a Jordan arc), we may assume that $\varphi(V(G)) \subseteq V(H)$ by subdividing the edges in $E(H)$ with at most $n = |V(G)|$ new vertices in $V(H)$ if necessary. Then φ maps every edge $e \in G$ to a path of length $O(n)$ in H . By subdividing the edges $e \in E(G)$ at all vertices in $V(H)$ along $\varphi(e)$, we reduce the recognition problem to the regime of simplicial maps (Theorem 1.1). The total number of vertices of G may increase to $O(mn)$ and the running time to $O(mn \log(mn)) = O(mn \log m)$.

COROLLARY 1.2. *Given an abstract graph G with m edges and n vertices, an embedded graph H and a piecewise linear continuous map $\varphi : G \rightarrow H$ that is injective on the interior of every edge in $E(G)$, we can decide in $O(mn \log(mn))$ time whether φ is a weak embedding.*

For example, this applies to straight-line drawings in $\mathbb{R}^2$ if the edges may pass through vertices.

COROLLARY 1.3. *Given an abstract graph G with n vertices and a map $\varphi : G \rightarrow \mathbb{R}^2$ where every edge is mapped to a straight-line segment, we can decide in $O(n^2 \log n)$ time whether φ is a weak embedding.*

In Section 6, we extend Theorem 1.1 and Corollary 1.2 to **nonorientable** surfaces with minor changes in the combinatorial representations, using a signature $\lambda : E(H) \rightarrow \{-1, 1\}$ to indicate whether an edge uv (and R_{uv}) is orientation-preserving or -reversing (similarly to Reference [5]).

Related Previous Work. The study of weak embeddings lies at the interface of several independent lines of research in mathematics and computer science. In topology, the study of weak embeddings and its higher dimensional analogs were initiated by Sieklucki [21] in the 1960s. One of his main results [21, Theorem 2.1] implies the following: Given a graph G and an embedded path H , every simplicial map $\varphi : G \rightarrow H$ is a weak embedding if and only if every connected component

of G is a subcubic graph with at most one vertex of degree three. However, even if H is a cycle or a 3-star, it is easy to construct simplicial maps $\varphi : G \rightarrow H$, where G has maximum degree 2, that are not weak embeddings. In this case, a series of recent papers on weakly simple *polygons* [1, 6, 7] show that weak embeddings can be recognized in $O(n \log n)$ time (the same time suffices when $\varphi : C_n \rightarrow H$ is a simplicial map [6] or a continuous map [1]).

Finding efficient algorithms for the recognition of weak embeddings $\varphi : G \rightarrow H$, where G is an arbitrary graph, was posed as an open problem in References [1, 6, 7]. The first polynomial-time solution for the general version follows from a recent variant [15] of the Hanani-Tutte theorem [16, 25], which was conjectured by M. Skopenkov [23] in 2003 and in a slightly weaker form already by Repovš and A. Skopenkov [20] in 1998. However, this algorithm reduces the problem to a system of $O(m)$ linear equations over $\mathbb{Z}_2$, where $m = |E(G)|$. The running time is dominated by solving this system in $O(m^{2\omega}) \leq O(m^{4.75})$ time, where $\omega \in [2, 2.373)$ is the matrix multiplication exponent; cf. Reference [26].

Weak embeddings of graphs also generalize various graph visualization models such as the recently introduced *strip planarity* [3] and *level planarity* [18]; and can be seen as a special case [2] of the notoriously difficult *clustered-planarity* (for short, *c-planarity*) [8, 12, 13], whose tractability remains elusive despite many attempts by leading researchers.

Outline. Our results rely on ideas from References [3, 6, 7] and Reference [15]. To distinguish the graphs G and H , we use the convention that G has **vertices** and **edges**, and H has **clusters** and **pipes**. A cluster $u \in V(H)$ corresponds to a subgraph $\varphi^{-1}[u]$ of G , and a pipe $uv \in E(H)$ corresponds to a set of edges $\varphi^{-1}[uv] \subseteq E(G)$.

The main tool in our algorithm is a local operation, called “cluster expansion,” which generalizes similar operations introduced previously for the case that G is a cycle [7]. Given an instance $\varphi : G \rightarrow H$ and a cluster $u \in V(H)$, it modifies u and its neighborhood (by replacing u with several new clusters and pipes) such that weak embeddability is invariant under the operation in the sense that the resulting new instance $\varphi' : G' \rightarrow H'$ is a weak embedding if and only if $\varphi : G \rightarrow H$ is a weak embedding. Our operation increases the number of clusters and pipes, but it decreases the number of “ambiguous” edges (i.e., multiple edges in the same pipe). The proof of termination and the running time analysis use potential functions.

In a preprocessing phase, we perform a cluster expansion operation at each cluster $u \in V(H)$. The main loop of the algorithm applies another operation, “pipe expansion,” for two adjacent clusters $u, v \in V(H)$ under certain conditions. It merges the clusters u and v , and the pipe $uv \in E(H)$ between them, and then invokes cluster expansion. If any of these operations finds a local configuration incompatible with an embedding, then the algorithm halts and reports that φ is not a weak embedding (this always corresponds to some nonplanar subconfiguration, since the neighborhood of a single cluster or pipe is homeomorphic to a disk). We show that after $O(m)$ successive operations, we obtain an irreducible instance for which our problem is easily solvable in $O(m)$ time. Ideally, we end up with $G = H$ (one vertex per cluster and one edge per pipe), and $\varphi = \text{id}$ is clearly an embedding. Alternatively, G and H may each be a cycle (possibly G winds around H multiple times), and we can decide whether φ is a weak embedding in $O(m)$ time by a simple traversal of G . If G is disconnected, then each component falls into one of the above two cases (Figure 1(c)–(d)), i.e., the case when $\varphi = \text{id}$ or the case when $\varphi \neq \text{id}$.

The main challenge was to generalize previous local operations (that worked well for cycles [6, 7, 23]) to arbitrary graphs. Our expansion operation for a cluster $u \in V(H)$ simplifies each component of the subgraph $\varphi^{-1}[u]$ of G independently. Each component is planar (otherwise it cannot be perturbed into an embedding in a disk D_u). However, a planar (sub)graph with k vertices may have $2^{\Omega(k \log k)}$ combinatorially different embeddings: some of these may or may not be compatible

Fig. 2. Two adjacent clusters, u and v , that each contain two components of pipe-degree 2 (left). These components merge into a single component in $D_u \cup R_{uv} \cup D_v$. In every embedding, the order of the pipe-edges a, b determines the order of the pipe-edges c, d . The operation $\text{pipeExpansion}(uv)$ transforms the component on the left to two wheels connected by three edges (a so-called thick edge) shown on the right.

with adjacent clusters. The embedding of a (simplified) component C of $\varphi^{-1}[u]$ depends on, among other things, the edges that connect C to adjacent clusters. The **pipe-degree** of C is the number of pipes that contain the edges incident to C . If the pipe-degree of C is 3 or higher, then the rotation system of H constrains the embedding of C . If the pipe-degree is 2, however, then the embedding of C can only be determined up to a reflection, unless C is connected by *two* independent edges to a component in $\varphi^{-1}[v]$ whose orientation is already fixed; see Figure 2.

We need to maintain the feasible embeddings of the components in all clusters efficiently. In Reference [15], this problem was resolved by introducing 0–1 variables for the components and aggregating the constraints into a system of linear equations over $\mathbb{Z}_2$, which was eventually resolved in $O(m^{2\omega}) \leq O(m^{4.75})$ time. We improve the running time to $O(m \log m)$ by maintaining the feasible embeddings simultaneously with our local operations.

Another challenge comes from the simplest components in a cluster $\varphi^{-1}[u]$. Long chains of degree-2 vertices, with one vertex per cluster, are resilient to our local operations. Their length may decrease by only one (and cycles are irreducible). We need additional data structures to handle these “slowly evolving” components efficiently. We use a dynamic heavy-path decomposition data structure and a suitable potential function to bound the time spent on such components.

Organization. In Section 2, we introduce additional terminology for an instance $\varphi : G \rightarrow H$ and show how to modify the subgraphs of G within each cluster to reduce the problem to an instance in “normal form,” a “simplified form,” and introduce a combinatorial representation of weak embeddings that we use in our algorithm. The simplification step relies on the concept of SPQR-trees, developed in Reference [10] for the efficient representation of combinatorial embeddings of a graph, which we also review in this section. In Section 3, we present the cluster expansion and pipe expansion operations and prove that weak embeddability is invariant under both operations. We use these operations repeatedly in Section 4 to decide whether a simplicial map $\varphi : G \rightarrow H$ is a weak embedding. Section 5 discusses how to reverse a sequence of operations to perturb a weak embedding into an embedding. The adaptation of our results to nonorientable surfaces M is discussed in Section 6. We conclude with open problems in Section 7.

2 PRELIMINARIES

In this section, we describe modifications within the clusters of a simplicial map $\varphi : G \rightarrow H$ to bring it to “normal form” (properties (P1)–(P2)) and “simplified form” (properties (P3)–(P4)). These properties allow for a purely combinatorial representation of weak embeddability (in terms of permutations), which we use in the proof of correctness of the algorithms in Sections 3 and 4.

Definitions. Two instances $\varphi : G \rightarrow H$ and $\varphi' : G' \rightarrow H'$ are called **equivalent** if φ is a weak embedding if and only if φ' is a weak embedding. We call an edge $e \in E(G)$ a **pipe-edge** if $\varphi(e) \in E(H)$, or a **cluster-edge** if $\varphi(e) \in V(H)$. For every cluster $u \in V(H)$, let G_u be the subgraph of G induced by $\varphi^{-1}[u]$. For every pipe $uv \in E(H)$, $\varphi^{-1}[uv]$ stands for the set of pipe-edges mapped to uv by φ .

Fig. 3. Changes in a cluster caused by normalize and simplify. (a) Input, (b) after normalize, (c) after first part of step 1, (d) after step 1, and (e) after step 2 of subroutine simplify. Dashed lines, green dots, green lines, and blue lines represent pipe-edges, pipe-vertices, edges in $\bar{E}_C$, and virtual edges, respectively.

The **pipe-degree** of a connected component C of G_u , denoted $\text{pipe-deg}(C)$, is the number of pipes that contain some edge of G incident to C . A vertex v of G_u is called a **terminal** if it is incident to a pipe-edge. For an integer $k \geq 3$, the k -vertex **wheel** graph W_k is a join of a **center** vertex c and a cycle of $k - 1$ **external** vertices. Refer to Reference [11] for standard graph theoretic terminology (e.g., cut vertex, 2-cuts, biconnectivity).

2.1 Normal Form

An instance $\varphi : G \rightarrow H$ is in **normal form** if every cluster $u \in V(H)$ satisfies:

- (P1) Every terminal in G_u is incident to exactly one cluster-edge and one pipe-edge.
- (P2) There are no degree-2 vertices in G_u .

We now describe subroutine $\text{normalize}(u)$ that, for a given instance $\varphi : G \rightarrow H$ and a cluster $u \in V(H)$, returns an equivalent instance $\varphi' : G' \rightarrow H$ such that u satisfies (P1)–(P2); refer to Figure 3(a)–(b).

normalize(u). Input: an instance $\varphi : G \rightarrow H$ and a cluster $u \in V(H)$.

Subdivide every pipe-edge pq where $\varphi(p) = u$ into a path (p, p', q) such that $\varphi'(p) = \varphi'(p') = u$, and $\varphi'(q) = \varphi(q)$. Note that the new vertex p' is a terminal in G and a leaf in G_u (i.e., $\deg_{G_u}(p') = 1$). Successively suppress every vertex p of G_u with $\deg_{G_u}(p) = 2$ by merging its incident edges. If this creates a loop, delete the loop.

LEMMA 2.1. *Given an instance $\varphi : G \rightarrow H$ and a cluster $u \in V(H)$, the instance $\varphi' = \text{normalize}(u)$ and φ are equivalent and u satisfies (P1)–(P2) in φ' . The subroutine runs in $O(\sum_{p \in V(G_u)} \deg_G(p))$ time. By successively applying normalize to all clusters in $V(H)$, we obtain an equivalent instance in normal form in $O(|E(G)|)$ time.*

PROOF. The instances φ and φ' are clearly equivalent, since (i) we can always replace the embedding of an edge by a path and vice versa, and (ii) a loop can always be deleted or added to a vertex in an embedding. There are $O(\sum_{p \in V(G_u)} \deg_G(p))$ pipe-edges incident to a vertex in u . Hence, the subroutine performs $O(\sum_{p \in V(G_u)} \deg_G(p))$ subdivisions. There are at most $|E(G_u)| = O(\sum_{p \in V(G_u)} \deg_G(p))$ degree-2 vertices in G_u . By construction, the resulting graph G'_u satisfies

(P1)–(P2). All changes are local and applying normalize to u does not change properties (P1)–(P2) in other clusters. Therefore, we can obtain the normal form of φ in $O(|E(G)|)$ time. $\square$

For every cluster $u \in V(H)$, the graph G_u may have several components. For each component C , we define a multigraph $\bar{C}$ that represents the interactions of C with vertices in adjacent clusters. Specifically, for each component C of G_u of a cluster $u \in V(H)$ satisfying (P1)–(P2), we define the multigraph $\bar{C}$ in two steps as follows (refer to Figure 3(b)):

- (1) For every pipe $uv \in E(H)$ incident to u , create a new vertex v' , called **pipe-vertex**, and identify all terminal vertices of C incident to some edge in $\varphi^{-1}[uv]$ with v' (this may create multiple edges incident to v').
- (2) If $\text{pipe-deg}(C) = 2$, connect the two pipe-vertices with an edge e and let $\bar{E}_C = \{e\}$. If $\text{pipe-deg}(C) \geq 3$, connect all pipe-vertices by a cycle in the order determined by the rotation of u and let $\bar{E}_C$ be the set of edges of this cycle. Hence, if $\text{pipe-deg}(C) \leq 1$, then $\bar{E}_C = \emptyset$.

Let $\bar{C} = (V(\bar{C}), E(\bar{C}))$, where $V(\bar{C})$ consists of nonterminal vertices in $V(C)$ and pipe-vertices, and $E(\bar{C})$ consists of the edges induced by nonterminal vertices in $V(C)$, (multi) edges created in step 1 (each of which corresponds to an edge in $E(C)$) and edges in $\bar{E}_C$. It is clear that every embedding of $\bar{C}$ can be converted into an embedding of C such that the rotation of the pipe vertices in $\bar{C}$ determines the cyclic order of terminals along the facial walk of the outer face of C .

2.2 SPQR-Trees

SPQR-trees were introduced by Di Battista and Tamassia [10] for an efficient representation of all combinatorial plane embeddings of a graph. Let G be a biconnected planar graph. The SPQR-tree T_G of G represents a recursive decomposition of G defined by its (vertex) 2-cuts. A deletion of a 2-cut $\{u, v\}$ disconnects G into two or more components $C_1, \dots, C_i$, $i \geq 2$. A **split component** of $\{u, v\}$ is either an edge uv (which is not one of the components C_j) or the subgraphs of G induced by $V(C_j) \cup \{u, v\}$ for $j = 1, \dots, i$. The tree T_G captures the recursive decomposition of G into split components defined by 2-cuts of G . A node μ of T_G is associated with a multigraph called $\text{skeleton}(\mu)$ on a subset of $V(G)$ and has a **type** in $\{S, P, R\}$. If the type of μ is S, then $\text{skeleton}(\mu)$ is a cycle of 3 or more vertices. If the type of μ is P, then $\text{skeleton}(\mu)$ consists of 3 or more parallel edges between a pair of vertices. If the type of μ is R, then $\text{skeleton}(\mu)$ is a 3-connected graph on 4 or more vertices. An edge in $\text{skeleton}(\mu)$ is **real** if it is an edge in G , or **virtual** otherwise. A virtual edge connects the two vertices of a 2-cut, u and v , and represents a subgraph of G obtained in the recursive decomposition, containing a uv -path in G that does not contain any edge in $\text{skeleton}(\mu)$. Two nodes μ_1 and μ_2 of T_G are adjacent if $\text{skeleton}(\mu_1)$ and $\text{skeleton}(\mu_2)$ share exactly two vertices, u and v , that form a 2-cut in G . Each virtual edge in $\text{skeleton}(\mu)$ corresponds bijectively to a pair of adjacent nodes in T_G . No two S nodes (resp., no two P nodes) are adjacent. Every edge in $E(G)$ appears in the skeleton of exactly one node. The tree T_G has $O(|E(G)|)$ nodes and it can be computed in $O(|E(G)|)$ time [10].

It is also known that T_G represents all combinatorial embeddings of G in $\mathbb{R}^2$ in the following manner [10]: Choose a root node for T_G and an embedding of its skeleton. Then successively replace each virtual edge uv by the skeleton of the corresponding node μ minus the virtual edge uv in $\text{skeleton}(\mu)$. In each step of the recursion, if μ is of type R, $\text{skeleton}(\mu)$ can be flipped (reflected) around u and v , and if μ is of type P, the parallel edges between u and v can be permuted arbitrarily.

2.3 Combinatorial Representation of Weak Embeddings

Given an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$, where φ is in normal form, we define a combinatorial representation $\pi = \pi(\psi_\varphi)$ as the set of total orders of edges in R_{uv} , for all pipes $uv \in E(H)$. Specifically, for every pipe $uv \in E(H)$, fix an orientation of the boundary of R_{uv} (e.g., the one used in the construction of the strip system $\mathcal{H}$). Record the order in which the pipe-edges in the embedding ψ_φ intersect a fixed side of ∂R_{uv} (say, the side $\partial R_{uv} \cap \partial D_u$) when we traverse it in the given orientation. Let $\Pi(\varphi)$ be the set of combinatorial representations $\pi(\psi_\varphi)$ of all embeddings $\psi_\varphi : G \rightarrow \mathcal{H}$. Note that $\varphi : G \rightarrow H$ is a weak embedding if and only if $\Pi(\varphi) \neq \emptyset$.

Conversely, let π be a set of total orders of edges in rectangles R_{uv} , for all pipes $uv \in E(H)$. We show (in Lemma 2.2 below) how to use the normal form and SPQR-trees to decide whether $\pi \notin \Pi(\varphi)$. We say that two components C_1 and C_2 of G_u **cross** with respect to π if and only if their terminals interleave in the cyclic order around ∂D_u (i.e., there exists no cut in the cyclic order in which all terminals of C_1 appear before all terminals of C_2). If $\pi \in \Pi(\varphi)$, then π cannot induce two crossing components in G_u , for any $u \in V(H)$.

LEMMA 2.2. *Given a set of total orders π for a map $\varphi : G \rightarrow H$, we can decide whether $\pi \notin \Pi(\varphi)$ in $O(m)$ time. If $\pi \in \Pi(\varphi)$, then we can also find an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$, $\pi = \pi(\psi_\varphi)$, in $O(m)$ time.*

PROOF. Let $\mathcal{H}$ be the strip system for $\varphi : G \rightarrow H$. For each pipe $uv \in E(H)$, draw $|\varphi^{-1}[uv]|$ parallel Jordan arcs in R_{uv} connecting ∂D_u and ∂D_v . For each cluster $u \in V(H)$, π defines a ccw cyclic order of terminals around ∂D_u . Let t_C denote the number of terminals in C . Create a graph $\bar{C}$ by the union of C and a wheel W_{t_C+1} whose center is a new vertex and whose external vertices are the terminals of C in the order defined by π . An embedding of C in which the terminals appear in the outer face in the same cyclic order as the one defined by π exists if and only if $\bar{C}$ is planar. If $\bar{C}$ is not planar, report that $\pi \notin \Pi(\varphi)$. Else, embed C inside D_u given the position of the already embedded terminals on ∂D_u . This subdivides D_u into faces. If a component connects terminals in different faces, then two components in G_u cross, hence $\pi \notin \Pi(\varphi)$. If no two components cross, we can incrementally embed all components of G_u , as there is always a face of D_u that contains all terminals of each remaining component. $\square$

2.4 Simplified Form

Given an instance $\varphi : G \rightarrow H$ in normal form, we simplify the graph G by removing parts of G_u , for all $u \in V(H)$, that are locally “irrelevant” for the embedding, such as 0-, 1-, and 2-connected components that are not adjacent to edges in any pipe incident to u . Formally, for each component C of G_u , we call a split component defined by a 2-cut $\{p, q\}$ of $\bar{C}$ **irrelevant** if it contains no pipe-vertices. An instance is in **simplified form** if it is in normal form and every $u \in V(H)$ satisfies properties (P3)–(P4) below.

(P3) For every component C of G_u , $\bar{C}$ is biconnected and every 2-cut of $\bar{C}$ contains at least one pipe-vertex.

Assuming that a cluster u satisfies (P3), we define T_C as the SPQR tree of $\bar{C}$ where C is a component of G_u . Given a node μ of T_C , let the *core* of μ , denoted $\text{core}(\mu)$, be the subgraph obtained from $\text{skeleton}(\mu)$ by deleting all pipe-vertices. Property (P4) below will allow us to bound the number of vertices of G_u in terms of its number of terminals (cf. Lemma 2.4).

(P4) For every component C of G_u , and every R node μ of T_C , $\text{core}(\mu)$ is isomorphic to a wheel W_k , for some $k \geq 4$, whose external vertices have degree 4 in G_u .

We now describe subroutine $\text{simplify}(u)$ that, for a given instance $\varphi : G \rightarrow H$ in normal form and a cluster $u \in V(H)$, returns an instance $\varphi' : G' \rightarrow H$ such that u satisfies (P1)–(P4). We break the subroutine into two steps.

simplify(u). Input: an instance $\varphi : G \rightarrow H$ in normal form and a cluster $u \in V(H)$.

For every component C of G_u , do the following.

(1) If C is not planar, report that φ is not a weak embedding and halt. If $\text{pipe-deg}(C) = 0$, then delete C . Else compute $\bar{C}$ and find the maximal biconnected component $\hat{C}$ of $\bar{C}$ that contains all pipe-vertices. The component $\hat{C}$ trivially exists if $\text{pipe-deg}(C) \in \{1, 2\}$, and if $\text{pipe-deg}(C) \geq 3$, it exists, since $\bar{E}_C$ forms a cycle containing all pipe-vertices. Modify C by deleting all vertices of $\bar{C} \setminus \hat{C}$, and update $\bar{C}$ (by deleting the same vertices from $\bar{C}$, as well); refer to Figure 3(b)–(c). Consequently, we may assume that $\bar{C}$ is biconnected and contains all pipe-vertices. Compute the SPQR tree T_C for $\bar{C}$. Set a node μ_r in T_C whose skeleton contains a pipe-vertex as the root of T_C . Traverse T_C using DFS. If a node μ is found such that $\text{skeleton}(\mu)$ contains no pipe-vertex, let $\{p, q\}$ be the 2-cut of $\bar{C}$ shared by $\text{skeleton}(\mu)$ and $\text{skeleton}(\text{parent}(\mu))$. Replace all irrelevant split components defined by $\{p, q\}$ by a single edge pq in C . If p or q now have degree 2, suppress p or q , respectively. Update $\bar{C}$ accordingly, and update T_C to reflect the changes in $\bar{C}$ by changing pq from virtual to real and possibly suppressing p and/or q in $\text{skeleton}(\text{parent}(\mu))$, which also deletes node μ and its descendants, since their skeletons contain edges in the deleted irrelevant split components; refer to Figure 3(c)–(d). Continue the DFS ignoring deleted nodes.

(2) While there is an R node μ in T_C , of a component C in G_u , that does not satisfy (P4), do the following: Let Y be the set of edges in $\text{skeleton}(\mu)$ adjacent to $\text{core}(\mu)$ (i.e., edges between a vertex in $\text{core}(\mu)$ and a pipe-vertex). Since μ is an R node, it represents a 3-connected planar graph that has a unique combinatorial embedding (up to reflection). If we contract $\text{core}(\mu)$ to a single vertex, the rotation of that vertex defines a cyclic order on Y . Let $e_i = p_i q_i$, where $p_i \in \text{core}(\mu)$ and q_i is a pipe-vertex, be the i th edge in this order. Recall that e_i is an edge in $\text{skeleton}(\mu)$ and therefore represents a subgraph $\bar{C}_{e_i}$ of $\bar{C}$. If e_i is real, $\bar{C}_{e_i}$ is a single edge. Otherwise, $\bar{C}_{e_i}$ contains all split components defined by the 2-cut $\{p_i, q_i\}$ that do not contain $\text{core}(\mu)$ as a subgraph. Do the following changes in C , which will incur changes in $\bar{C}$, as well: Replace $\text{core}(\mu)$ by the wheel graph $W_{|Y|+1}$ by deleting all interior edges and vertices of $\text{core}(\mu)$ and inserting a new vertex adjacent to all vertices of the outer cycle. If $\bar{C}_{e_i}$ is not a single edge, then replace p_i by two adjacent vertices, p'_i and p''_i where p'_i is in $W_{|Y|+1}$ and p''_i is in $\bar{C}_{e_i}$ (that is, p'_i is adjacent to p''_i and the three neighbors of p_i in the wheel, and p''_i is adjacent to p'_i and the remaining neighbors of p_i). Update T_C to reflect the changes in $\bar{C}$ by updating $\text{skeleton}(\mu)$ and adding an S node between μ and an adjacent node μ_i whose skeleton contained e_i for each virtual edge $e_i \in Y$ where p_i has not been not suppressed; refer to Figure 3(d)–(e).

LEMMA 2.3. *Given an instance $\varphi : G \rightarrow H$ in normal form and a cluster $u \in V(H)$, the instance $\varphi' = \text{simplify}(u)$ and φ are equivalent and u satisfies (P1)–(P4) in φ' . The operation runs in $O(|E(G_u)|)$ time. By successively applying *simplify* to all clusters in $V(H)$, we obtain an equivalent instance in simplified form in $O(m)$ time.*

PROOF. First, we prove that u satisfies (P1)–(P4) in φ' . Since φ is in normal form, u satisfies (P1)–(P2) in φ . By construction, u still satisfies (P1)–(P2) in φ' . For (P3), note that after step 1, $\bar{C}$ is biconnected and every node μ of T_C contains a pipe-vertex in its skeleton. Step 2 does not change this property. This implies that $\text{core}(\mu)$ contains only real edges for every node μ . Suppose for contradiction that there is a 2-cut $\{p, q\}$ such that neither p nor q is a pipe-vertex. Then $\{p, q\}$ must be in $\text{core}(\mu)$ where μ is an S node, or else either p or q would have been deleted for being in

$\bar{C} \setminus \hat{C}$. Then one split component of $\{p, q\}$ is a path of length two or more. But G'_u has no degree-2 vertex by (P2), a contradiction. Hence, u satisfies (P3) in φ' . By definition, after step 2 u satisfies (P4) in φ' .

We now show that the operation takes $O(|E(G_u)|)$ time. In step 1, planarity testing is done in linear time for each component C of G_u [17]. We obtain $\hat{C}$ by a DFS. We compute T_C in $O(|E(C)|)$ time [10]. Replacing irrelevant split components by one edge can be done in $O(|E(C)|)$ overall time. In step 2, we can obtain a list of R nodes in $O(|E(C)|)$ time. The changes in step 2 are local, both in C and T_C , and do not influence whether other R nodes satisfy (P4). Step 2 takes $O(|E(G_u)|)$ time overall by processing each R node sequentially. All the changes are local to u and, by successively applying simplify, we obtain a simplified form in $O(|E(G)|)$ time.

Finally, we show that φ and φ' are equivalent. Notice that there is a bijection between the terminals of φ and φ' . We show that $\Pi(\varphi) = \Pi(\varphi')$, i.e., given $\pi \in \Pi(\varphi)$, then $\pi \in \Pi(\varphi')$ and vice versa. Notice that for every $\pi \in \Pi(\pi)$ and $u \in V(H)$, two components C_1 and C_2 of G_u cross if and only if the corresponding components C'_1 and C'_2 of G'_u also cross. Then, it suffices to show that the SPQR trees of $\bar{C}$ and $\bar{C}'$ for corresponding components C of G_u and C' of G'_u represent the same constraints in the cyclic order of terminals around D_u . Step 1 deletes components of pipe-degree 0, which do not pose any restriction on the cyclic order of terminals. The subgraphs represented by irrelevant subtrees in the SPRQ tree of $\bar{C}$ can be flipped independently and, since they do not contain pipe-vertices, their embedding does not interfere with the order of edges adjacent to pipe-vertices. Hence, step 1 does not alter any constraint on the cyclic order of terminals. By construction, step 2 does not change the circular order of edges in $\text{core}(\mu)$. Replacing $\text{core}(\mu)$ by a wheel $W_{|Y|+1}$ does not change any of the constraints on the cyclic order of terminals. $\square$

LEMMA 2.4. *After $\text{simplify}(u)$, every component C of G_u contains $O(t_C)$ edges, where t_C is the number of terminals in C .*

PROOF. By (P4), every wheel subgraph in C is a maximal biconnected component. Let us contract every wheel in C into a single vertex and remove any loops created by the contraction. Let $\hat{C}$ denote the resulting component. We have $|E(C)| \leq 5|E(\hat{C})|$, since the number of edges decreases by at most $4|E(\hat{C})|$. Indeed, a wheel W_{k+1} has $2k$ edges, which are contracted, and its k external vertices are incident to k edges that are not contracted by (P4). We charge each of these k edges in $E(\hat{C})$ for two edges of W_k . Then every edge in $E(\hat{C})$ receives at most 2 units of charge from each of its endpoints, hence at most 4 units of charge overall. As all maximal biconnected components of C have been contracted to single vertices of degree at least three, $\hat{C}$ is a tree without degree-2 vertices whose leaves are precisely the terminals of C by (P1)–(P3). Since the number of edges in a tree is at most twice the number its leaves, by the above inequality, we have $|E(C)| \leq 5|E(\hat{C})| \leq 5 \cdot 2 \cdot t_C$, as claimed. $\square$

3 OPERATIONS

In this section, we present our two main operations, `clusterExpansion` and `pipeExpansion`, that we use successively in our recognition algorithm. Given an instance φ and a cluster u in simplified form, operation `clusterExpansion`(u) either finds a configuration that cannot be embedded locally in the neighborhood of u and reports that φ is not a weak embedding, or replaces cluster u with a group of clusters and pipes (in most cases reducing the number of edges in pipes). It first modifies the embedded graph H and then handles each component of G_u independently.

Operation `pipeExpansion`(uv) first merges two adjacent clusters, u and v (and the pipe uv) into a single cluster $\langle uv \rangle$ and invokes `clusterExpansion`($\langle uv \rangle$). We continue with the specifics.

Fig. 4. Changes in a cluster caused by clusterExpansion. Left: input. Note that the embedding of G_u is not known and the grey disks (rectangles) represent points (arcs). Middle: the instance after step 4. Right: output. Red rectangles indicate triples of edges defining a thick edge.

3.1 Cluster Expansion

For a cluster $u \in V(H)$ in an instance $\varphi : G \rightarrow H$, let the **expansion disk** Δ_u be a topological closed disk containing a single cluster $u \in V(H)$ and intersecting only pipes incident to u .

clusterExpansion(u). Input: an instance $\varphi : G \rightarrow H$ in simplified form and a cluster $u \in V(H)$. We either report that φ is not a weak embedding or return an instance $\varphi' : G' \rightarrow H'$. The instance φ' is computed incrementally: initially φ' is a copy of φ . Steps 0–3 will insert new clusters and pipes into H' that are within Δ_u without describing their embedding, step 4 will determine the rotation system for the new clusters and check whether the rotation system induces any crossing between new pipes within Δ_u , and step 5 brings φ' to its simplified form.

Step 0. For each pipe $uv \in E(H)$ incident to u , subdivide uv by inserting a cluster u_v in H' at the intersection of uv and $\partial\Delta_u$. If $\deg(u) \geq 3$, then add a cycle C_u of pipes through all clusters in $\partial\Delta_u$ (hence, the clusters u_v appear along C_u in the order given by the rotation of u); and if $\deg(u) = 2$, then add a pipe between the two clusters in $\partial\Delta_u$. Delete u (and all incident pipes). As a result, the interior of Δ_u contains no clusters.

Step 1: Components of pipe-degree 1. For each component C of G_u such that $\text{pipe-deg}(C) = 1$, let uv be the pipe to which the pipe-edges incident to C are mapped to. Move C to the new cluster u_v , i.e., set $\varphi'(C) = u_v$. (For example, see the component in u_x in Figure 4.)

Step 2: Components of pipe-degree 2. For each pair of clusters $\{v, w\}$ adjacent to u , denote by B_{vw} the set of components of G_u of degree 2 adjacent to pipe-edges in $\varphi^{-1}[uv]$ and $\varphi^{-1}[uw]$. For all nonempty sets B_{vw} do the following:

(a) Insert the pipe $u_v u_w$ into H' if it is not already present.

(b) For every component $C \in B_{vw}$, do the following:

(b1) Compute $\bar{C}$ (by (P3), $\bar{C}$ is biconnected). Compute the SPQR tree T_C of $\bar{C}$. Set a node μ_r as the root of T_C so $\text{skeleton}(\mu_r)$ contains both pipe-vertices, which we denote by v' and w' (i.e., consistently with Section 2). Note that μ_r cannot be of type P, otherwise C would not be connected.

(b2) If μ_r is of type S, then $E(\text{skeleton}(\mu_r)) \setminus \bar{E}_C$ forms a path between the pipe vertices v' and w' , which we denote by P , where the first and last edges may be virtual. Notice that path P contains at most 3 edges, otherwise (P2) or (P3) would not be satisfied. If $P = (v', w')$ has length 1, then subdivide P into 3 edges $P = (v', p_1, p_2, w')$. If $P = (v', p, w')$ has length 2, then $\{v', p\}$ is a 2-cut in $\bar{C}$ that defines two split components, C_v and C_w , containing v' and w' , respectively. Note that it is not possible that both $v'p$ and pw' are real edges, because there are no degree-2 vertices in G_u .

by (P2). Split p into two vertices, p_1 and p_2 , connected by an edge so p_1 (resp., p_2) is adjacent to every vertex in C_v (resp., C_w) that was adjacent to p . Finally, if $P = (v', p_1, p_2, w')$ has length 3, we do not modify P . Then, in all three cases, the edge p_1p_2 defines an edge cut in C that splits C into two components each incident to a single pipe, one to uv and the other to uw . We define φ' so it maps each of the two components into u_v or u_w accordingly. (See the components incident to pipe $u_{v_1}u_{v_3}$ in Figure 4.)

(b3) If μ_r is of type R, by (P4), $\text{core}(\mu_r)$ is a wheel subgraph W_k . Let k_v and k_w , where $k_v + k_w = k - 1$, be the number of edges between W_k and v' , and between W_k and w' , respectively. Replace W_k by two wheel graphs W_{k_v+4} and W_{k_w+4} connected by three edges so the circular order of the edges around v' and w' is maintained (recall that an R node has a unique embedding). The triple of edges between W_{k_v+4} and W_{k_w+4} is called a **thick edge**. The thick edge defines a 3-edge-cut that splits C into two components, each with a wheel graph. We define φ' so each of the two components is mapped to its respective vertex u_v or u_w . (See the components incident to pipe $u_{v_1}u_w$ in Figure 4.)

Step 3: Components of pipe-degree 3 or higher. For all the remaining components C (i.e., $\text{pipe-deg}(C) \geq 3$) of G_u , do the following: Assume C is incident to pipe-edges mapped to the pipes $uv_1, uv_2, \dots, uv_d$.

(a) Compute $\bar{C}$ and its SPQR tree T_C and let v'_i be the pipe-vertex corresponding to terminals adjacent to edges in uv_i for $i = 1, 2, \dots, d$. Set the node μ_r as the root of T_C such that $\text{skeleton}(\mu_r)$ contains the cycle $\bar{E}_C$. The type of μ_r is R, otherwise C would be disconnected.

(b) **Changes in H' .** By (P4), we have that $\text{core}(\mu_r)$ is a wheel graph W_k , where $k - 1 \geq d$. For $j = 1, 2, \dots, k - 1$, let p_j be the j th external vertex of W_k and p_C be its central vertex. Create a copy of W_k using clusters and pipes: Create a cluster u_{p_j} that represents each vertex p_j , a cluster u_C that represents vertex p_C (see Figure 4 (middle)). Insert the copy of W_k in H' . For every $1 \leq i \leq d$ and $1 \leq j \leq k - 1$, insert a pipe $u_{p_j}u_{v_i}$ in H' if an edge $p_jv'_i$ is present in $E(\text{skeleton}(\mu_r))$.

(c) **Changes in G' .** Delete all edges and the central vertex of W_k from G' , which splits C into $k - 1$ components. Set $\varphi'(p_j) = u_{p_j}$ for $j \in \{1, \dots, k - 1\}$. By (P4), every cluster u_{p_j} is adjacent to a single cluster, say u_{v_i} , outside of W_k . We modify φ' so it maps the vertices of the component of C containing p_j to u_{v_i} with the exception of p_j , which is mapped to u_{p_j} . (Note that the cluster u_C and all incident pipes are empty.)

Step 4: Local Planarity Test. Let H_u be the subgraph induced by the newly created clusters and pipes. Let $\widetilde{H}_u$ denote the graph obtained as the union of H_u and a star whose center is a new vertex (not in $V(H_u)$) and whose leaves are the clusters in $\partial\Delta_u$. Use a planarity testing algorithm to test whether $\widetilde{H}_u$ is planar. If $\widetilde{H}_u$ is not planar, report that φ is not a weak embedding and halt. Otherwise, find an embedding of $\widetilde{H}_u$ in which the center of the star is in the outer face. This defines a rotation system for H_u . The rotation system of H' outside of Δ_u is inherited from H .

Step 5: Normalize. Finally, apply normalize to each new cluster in H' . (This step subdivides edges so (P1) is satisfied as shown in Figure 4 (right).)

LEMMA 3.1. *Given an instance $\varphi : G \rightarrow H$ in simplified form containing a cluster u , $\text{clusterExpansion}(u)$ either reports that φ is not a weak embedding or produces an instance $\varphi' : G' \rightarrow H'$ in simplified form that is equivalent to φ in $O(|E(G_u)| + \deg(u))$ time.*

PROOF. All clusters of H other than u are clusters in H' as well. By assumption, $\varphi : G \rightarrow H$ is simplified form, hence these clusters satisfy (P1)–(P4) in φ' . By construction, the new clusters satisfy (P1)–(P4) in φ' , consequently $\varphi' : G' \rightarrow H'$ is in simplified form.

Step 5 of the operation receives as input an instance $\varphi^* : G^* \rightarrow H^*$ and returns an equivalent instance φ' by Lemma 2.1. It remains to show that φ^* and φ are equivalent, and to analyze the running time.

First assume that φ is a weak embedding, and so there is an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$. We need to show that there exists an embedding $\psi_{\varphi^*} : G^* \rightarrow \mathcal{H}^*$, and hence φ^* is a weak embedding. This can be done by performing steps 0–3 on the graphs G and H , and the embedding ψ_φ , which will produce G^* and H^* , and an embedding $\psi_{\varphi^*} : G^* \rightarrow \mathcal{H}^*$. By construction, every 2-cut in H_u^* consists of a pair of clusters in $\partial\Delta_u$. Consequently, H_u^* has a unique embedding with the given outer cycle. In particular, $\widetilde{H}_u^*$ is planar. The rotation of the new clusters of H^* is uniquely defined and therefore must be consistent with any embedding of G , including ψ_φ .

Next, assume that φ^* is a weak embedding. Given an embedding $\psi_{\varphi^*} : G^* \rightarrow \mathcal{H}^*$, we construct an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$ as follows: Let H_u^* be the subgraph of H^* induced by the clusters created by `clusterExpansion`(u). Note that H^* is a connected plane graph: the clusters created in step 0 are connected by a path (if $\deg(u) \leq 2$) or a cycle (if $\deg(u) \geq 3$); and any clusters created in step 3 (when $\deg(u) \geq 3$) are attached to this cycle. Since H_u^* is a connected plane graph, we may assume that there is a topological disk containing only the pipes and clusters of H_u^* ; let D_u denote such a topological disk.

Let G_u^* be the subgraph of G^* mapped to H_u^* . We show that steps 0–3 of the operation can be reversed without introducing crossings. For each component C of G_u with $\text{pipe-deg}(C) \geq 3$, embed a wheel W_k in the disk D_{u_C} around the cluster u_C and connect its external vertices to the vertices p_i , $i = 1, \dots, k-1$. Since the pipes incident to u_C are empty, and each p_i is a unique vertex in its cluster, this can be done without crossings. Now, every component C of G_u corresponds to a component C^* of G_u^* , and by (P1) every terminal vertex in C corresponds to terminal in C^* .

If we delete a component C^* from G^* , there will be a face F of D_u (a component of $D_u \setminus \psi_{\varphi^*}(G^*)$) that contains all terminals of C on its boundary. Denote by π_C the ccw cyclic order in which these terminals appear in the facial walk of F . If C admits an embedding in which the terminals appear in the outer face in the same order as π_C , we can embed C in F on the given terminals. We show that C admits such an embedding by proving that the SPQR trees of $\overline{C}$ and $\overline{C}^*$ impose the same constraints on the cyclic order of terminals. Subdividing edges do not change these constraints. Step 1 does not change C . If step 2(b2) adds an edge in the skeleton of an S node, the possible combinatorial embeddings remain the same. Step 2(b3) maintains the rotation of v' and w' . If $\text{pipe-deg}(C) \geq 3$, C and C^* are identical apart from subdivided edges. Because all steps maintain the same constraints on the rotation system of the terminals, we can construct $\psi_\varphi : G \rightarrow \mathcal{H}$ by incrementally replacing the embedding of C^* by an embedding of C in D_u for every component C of G_u . By Lemma 2.2, this is possible without introducing crossings, since no two components of G_u cross and every component is planar, because φ is in simplified form.

Finally, we show that `clusterExpansion`(u) runs in $O(|E(G_u)| + \deg(u))$ time. Step 0 takes $O(\deg(u))$ time. Steps 1–3 are local operations that take $O(|E(C)|)$ time for each component C of G_u . Step 4 takes $O(|E(G_u)| + \deg(u))$ time, since each component C inserts at most $O(|E(C)|)$ pipes in H' , where $|E(G_u)| = \sum_C |E(C)|$; the graph $\widetilde{H}_u$ has $\deg(u)$ more edges than H_u , and planarity testing takes linear time in the number of edges [17]. Step 5 takes $O(|E(G_u)|)$ time by Lemma 2.1 and (P1). $\square$

3.2 Pipe Expansion

A cluster $u \in V(H)$ is called a **base** of an incident pipe $uv \in E(H)$ if every component of G_u is incident to: (i) at least one pipe-edge in uv ; and (ii) at most one pipe-edge or one thick pipe-edge (i.e., a triple of pipe-edges), other than those in pipe uv . We call a pipe uv **safe** if both of its endpoints are bases of uv ; otherwise, it is **unsafe**. For a simplified instance $\varphi : G \rightarrow H$ and a safe pipe $uv \in E(H)$, operation `pipeExpansion`(uv) consists of the following steps:

Fig. 5. Pipe Expansion. A safe pipe uv (top left). The cluster $\langle uv \rangle$ obtained after contraction of uv (top right). The result of contracting the components in G_u and G_v (bottom left). The subsequent contraction of all components incident to $\partial_u D_{\langle uv \rangle}$ and $\partial_v D_{\langle uv \rangle}$ (all loops are already deleted), resp., and a Jordan curve that crosses every edge of the resulting bipartite plane multigraph for each component (bottom center). All Jordan curves can be combined into one (bottom right).

pipeExpansion (uv). Input: an instance $\varphi : G \rightarrow H$ in simplified form and a safe pipe $uv \in E(H)$. The operation either reports that φ is not a weak embedding or returns an instance $\varphi' : G' \rightarrow H'$. First, produce an instance $\varphi^* : G \rightarrow H^*$ by contracting the pipe uv into the new cluster $\langle uv \rangle \in V(H^*)$ while mapping the vertices of G_u and G_v to $\langle uv \rangle$ (see Figure 5 (top)); and then apply *simplify* and *clusterExpansion* to $\langle uv \rangle$.

We use the following folklore result in the proof of correctness of operation *pipeExpansion*(uv). This result is obtained by a Euler tour algorithm on the dual graph of a plane bipartite multigraph.

THEOREM 3.2 (BELYI [4]). *For every embedded connected bipartite multigraph G^* , there exists a Jordan curve that crosses every edge of G^* precisely once. Such a curve can be computed in $O(|E(G^*)|)$ time.*

LEMMA 3.3. *Given an instance $\varphi : G \rightarrow H$ and a safe pipe $uv \in E(H)$, *pipeExpansion*(uv) either reports that φ is not a weak embedding or produces an equivalent instance $\varphi' : G' \rightarrow H'$.*

PROOF. Let $\varphi : G \rightarrow H$ be an instance in simplified form, and let uv be a safe pipe. Recall that *pipeExpansion*(uv) starts with producing an instance $\varphi^* : G \rightarrow H^*$ by contracting the pipe uv into the new cluster $\langle uv \rangle \in V(H^*)$ while mapping the vertices of G_u and G_v to $\langle uv \rangle$. It is enough to prove that φ and φ^* are equivalent; the rest of the proof follows from Lemmas 2.1, 2.3, and 3.1.

One direction of the equivalence proof is trivial: Given an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$, we can obtain an embedding $\psi_{\varphi^*} : G \rightarrow \mathcal{H}^*$ by defining $D_{\langle uv \rangle}$ as a topological disk containing only D_u , D_v , and R_{uv} .

For the other direction, assume that we are given an embedding $\psi_{\varphi^*} : G \rightarrow \mathcal{H}^*$. We need to show that there exists an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$. We shall apply Theorem 3.2 after contracting certain subgraphs of G_u and G_v (as described below). If a component contains cycles, then the contraction of its embedding creates a bouquet of loops. We study the cycles formed by G_u , G_v , and thick edges incident to G_u or G_v to ensure that no other component is embedded in the interior of such cycles.

Components of $G_{\langle uv \rangle}$ of pipe-degree 0. Note that the terminals corresponding to the pipes incident to u and v lie in two disjoint arcs of $\partial D_{\langle uv \rangle}$, which we denote by $\partial_u D_{\langle uv \rangle}$ and $\partial_v D_{\langle uv \rangle}$, respectively. The components of graph $G_{\langle uv \rangle}$ with positive pipe-degree are incident to terminals

Fig. 6. Changing the embedding so: (a) no cycle encloses vertices that are not the center of a wheel; and (b) no induced cycle of a thick edge encloses a vertex.

in $\partial_u D_{\langle uv \rangle}$ or $\partial_v D_{\langle uv \rangle}$ (possibly both). The components of pipe-degree 0 can be relocated to any face of the embedding of all other components. Without loss of generality, we may assume that all components of pipe-degree 0 lie in a common face incident to both G_u and G_v in ψ_{φ^*} .

Cycles induced by G_u and G_v , and by thick edges. Notice that (P3) and (P4) imply that every maximal biconnected component in G_u and G_v is a wheel. Since each wheel is 3-connected, the circular order of their external vertices is determined by the embedding ψ_{φ^*} . We may assume that in the embedding no cycle of a wheel subgraph encloses any vertex other than the center of the wheel. Indeed, suppose a 3-cycle (p_1, p_2, p_c) of a wheel encloses some vertex, where p_c is the center of the wheel, and p_1 and p_2 are two consecutive external vertices. We can modify the embedding of the edge $p_1 p_2$ in ψ_{φ^*} so it closely follows the path (p_1, p_c, p_2) and the 3-cycle does not contain any vertex; see Figure 6(a).

Consider a thick edge θ in G between a wheel in G_u (or G_v) and a wheel in G_w for some adjacent cluster $w \notin \{u, v\}$. Recall that a thick edge gives rise to three paths, say $P_1 = (p_1, t_1, t_2, p_2)$, $P_2 = (p_3, t_3, t_4, p_4)$, and $P_3 = (p_5, t_5, t_6, p_6)$, where (p_1, p_3, p_5) and (p_2, p_4, p_6) are consecutive external vertices of the two wheels, resp., and p_i is the unique vertex in a cluster adjacent to terminal t_i for $i \in \{1, \dots, 6\}$; cf. (P1). If we suppress the terminals, then the two wheels incident to the thick edge would be in the same maximal 3-connected component of G and, therefore, their relative embedding is fixed within $D_u \cup R_{uw} \cup D_w$. We may assume that no vertex is enclosed by a cycle induced by the vertices of the thick edge (i.e., by any pair of paths from P_1 , P_2 , and P_3) in $D_u \cup R_{uw} \cup D_w$. Indeed, we can modify the embedding of the path P_1 and P_3 in ψ_{φ^*} so they closely follow the path $p_1 p_3 \cup P_2 \cup p_4 p_2$ and $p_5 p_3 \cup P_2 \cup p_4 p_6$, respectively. By (P4), such modification of the embedding is always possible without introducing crossings; see Figure 6(b). We conclude that a cycle induced by the thick edge θ does not enclose any vertices of G .

Separating G_u and G_v . We next show that there exists a closed Jordan curve that separates the embeddings of G_u and G_v , and crosses every edge between G_u and G_v precisely once. To use Theorem 3.2, we reduce $G_{\langle uv \rangle}$ and its embedding to a bipartite multigraph in two steps.

(1) Contract each component C of G_u (resp., G_v) to a single vertex w_C (see Figure 5 (bottom-left)). This results in a bouquet of loops at w_C that corresponds to faces bounded by C . As argued above, a cycle through the external vertices of a wheel in G_u or G_v encloses only its center. Hence, none of the loops at w_C encloses any other vertex of G , and they can be discarded.

(2) As uv is safe, every component C of G_u and G_v is incident to either at most one terminal in $\partial D_{\langle uv \rangle}$ or precisely three terminals corresponding to a thick edge. Contract the arc $\partial_u D_{\langle uv \rangle}$ (resp., $\partial_v D_{\langle uv \rangle}$) and for every component C of G_u (resp., G_v) all edges between w_C and terminals into a new vertex $u_\infty \in \partial D_{\langle uv \rangle}$ (resp., $v_\infty \in \partial D_{\langle uv \rangle}$) (see Figure 5 (bottom-center)) and insert an edge $u_\infty v_\infty$ in the exterior of $D_{\langle uv \rangle}$. This results in a bouquet of loops at u_∞ (resp., v_∞), corresponding to thick edges. As argued above, these loops do not enclose any vertices and can be eliminated. The subgraph of all components of pipe-degree 1 or higher has been transformed into a connected bipartite plane multigraph, and each component of pipe-degree 0 is also transformed into a connected bipartite plane multigraph.

We apply Theorem 3.2 for each of these plane multigraphs independently. By our assumption, the components C , $\text{pipe-deg}(C) = 0$, lie in a common face. Consequently, we can combine their Jordan curves with the Jordan curve of all remaining components into a single Jordan curve that crosses every edge between G_u and G_v precisely once (see Figure 5(bottom-right)). After reversing the contractions and loop deletions described above, we obtain a Jordan curve γ that separates G_u and G_v , and crosses every edge between G_u and G_v precisely once.

Since γ crosses the edge $u_\infty v_\infty$ precisely once, and $u_\infty v_\infty$ is homotopic to either subarcs of $\partial D_{\langle uv \rangle}$ between u_∞ and v_∞ , we can assume that γ crosses each of these arcs precisely once. Consequently, the Jordan arc $\gamma' = \gamma \cap D_{\langle uv \rangle}$ partitions the disk $D_{\langle uv \rangle}$ into two disks, D_u and D_v , adjacent to $\partial_u D_{\langle uv \rangle}$ and $\partial_v D_{\langle uv \rangle}$, respectively, such that every edge between G_u and G_v crosses γ' precisely once and all other edges of $G_{\langle uv \rangle}$ lie entirely in either D_u or D_v . This yields the required embedding $\psi_\varphi : G \rightarrow \mathcal{H}$. $\square$

4 ALGORITHM AND RUNTIME ANALYSIS

In this section, we present our algorithm for recognizing weak embeddings. We describe the algorithm and prove that it recognizes weak embeddings in Section 4.1. A naïve implementation would take $O(m^2)$ time, as explained below; we describe how to implement it in $O(m \log m)$ time using additional data structures in Section 4.3.

4.1 Main Algorithm

We are given a piecewise linear simplicial map $\varphi : G \rightarrow H$, where G is a graph and H is an embedded graph in an orientable surface. We introduce some terminology for an instance $\varphi : H \rightarrow G$.

- A component C of G_u , in a cluster $u \in V(H)$, is **stable** if $\text{pipe-deg}(C) = 2$ and, in each of the two pipes, C is incident to exactly one edge or exactly one thick edge. Otherwise, C is **unstable**.
- Recall that a pipe $uv \in E(H)$ is *safe* if both u and v are bases of uv ; otherwise, it is *unsafe*.
- A cluster $u \in V(H)$ (resp., a pipe $uv \in E(H)$) is **empty** if $\varphi^{-1}[u] = \emptyset$ (resp., $\varphi^{-1}[uv] = \emptyset$). Otherwise, u is **nonempty**.
- A safe pipe $uv \in E(H)$ is **useless** if uv is empty or if u and v are each incident to exactly two nonempty pipes and every component of G_u and G_v is stable. The safe pipe uv is **useful** otherwise (i.e., if uv is nonempty; and u or v is incident to at least 3 nonempty pipes, or G_u or G_v has an unstable component).
- A **generalized cycle** is a cycle in which each node is a component of G_u for some cluster u , and each edge is either an edge or a thick edge between wheels.

Notice that if a pipe uv is useless, then $\text{pipeExpansion}(uv)$ returns a combinatorially equivalent instance (i.e., there would be no progress). As we show below (Lemma 4.6), our algorithm reduces G to a collection of generalized cycles. (In particular, if we contract every wheel to a single vertex, then a generalized cycle reduces to a cycle, possibly with multiple edges coming from thick edges.)

Data Structures. The graph G is stored using adjacency lists. We store the combinatorial embedding of H using the rotation system of H (ccw order of pipes incident to each $u \in V(H)$). The mapping φ is encoded by its restriction to $V(G)$; that is, by the images $\varphi(p)$ for all $p \in V(G)$. For each cluster $u \in V(H)$, we store the set of components of G_u , each stored as the ID of a representative edge and the pipe-degree of the component. Every pipe $uv \in E(H)$ has two Boolean variables to indicate whether the respective endpoints are its bases; and an additional Boolean variable indicates whether uv is useful. These data structures are maintained dynamically as the algorithm modifies G , H , and φ .

Algorithm(φ). Input: an instance $\varphi : G \rightarrow H$ in simplified form.

Phase 1. Apply `clusterExpansion` to each $u \in V(H)$. Denote the resulting instance by $\varphi' : G' \rightarrow H'$. Build the data structures described above for $\varphi' : G' \rightarrow H'$.

Phase 2. While there is a useful pipe in H' , let $uv \in E(H')$ be an arbitrary useful pipe and apply `pipeExpansion`(uv).

Phase 3. If any component of G' that contains a wheel is nonplanar, then report that φ is not a weak embedding and halt. Otherwise, in each cluster contract every wheel component to a single vertex and turn every thick edge into a single edge by removing multiple edges. Denote the resulting instance by $\varphi'' : G'' \rightarrow H''$. If any component C of G'' is a cycle with k vertices but $\varphi''(C)$ is not a cycle with k clusters in H'' , then report that φ is not a weak embedding, else report that φ is a weak embedding. This completes the algorithm.

4.2 Analysis of Algorithm

We show that the algorithm recognizes whether the input $\varphi : G \rightarrow H$ is a weak embedding. The running time analysis follows in Section 4.3.

Termination. Since both Phase 1 and Phase 3 consist of for-loops only, it is enough to show that the while loop in Phase 2 terminates. We define a nonnegative potential function $\Phi_1(\varphi)$ for an instance $\varphi : G \rightarrow H$. For a pipe $uv \in E(H)$, let $\sigma(uv)$ be the number of pipe-edges in $\varphi^{-1}[uv]$ minus twice the number of thick edges (so each thick edge is counted as one). Let $s(u)$ be the number of stable components of G_u . We define the following quantities for an instance $\varphi : G \rightarrow H$:

$$\begin{aligned} N_\sigma(\varphi) &= \sum_{uv \in E(H)} \sigma(uv), \\ N_s(\varphi) &= \sum_{u \in V(H)} s(u). \end{aligned} \tag{1}$$

Let $Q(\varphi)$ be the number of nonempty pipes of an instance φ . We can now define the potential function

$$\Phi_1(\varphi) = 4(N_\sigma(\varphi) - Q(\varphi)) + N_s(\varphi). \tag{2}$$

Note that $\Phi_1(\varphi)$ is a nonnegative integer, since $N_\sigma(\varphi) \geq Q(\varphi)$.

The following lemma describes the effect of one iteration of the while loop in Phase 2 on Φ_1 and $N_\sigma - Q$.

LEMMA 4.1. *Assume we invoke `pipeExpansion`(uv) for a useful pipe uv in an instance φ , obtaining instance φ' . Then,*

- $\Phi_1(\varphi) > \Phi_1(\varphi')$, and
- $N_\sigma(\varphi) - Q(\varphi) \geq N_\sigma(\varphi') - Q(\varphi')$.

PROOF. Let $\langle uv \rangle$ be the cluster obtained after the first step of `pipeExpansion`(uv) (before applying simplify), and let C be a component of $G_{\langle uv \rangle}$. Let $\sigma_C(uv)$ be the number of pipe-edges of C in

$\varphi^{-1}[uv]$ minus twice the number of thick edges of C in $\varphi^{-1}[uv]$ before $\text{pipeExpansion}(uv)$. Let s_C be the number of stable components of G_u and G_v contained in C before $\text{pipeExpansion}(uv)$. Then

$$\sigma(uv) = \sum_C \sigma_C(uv) \quad \text{and} \quad s(u) + s(v) = \sum_C s_C,$$

where the summation is over all components C of $G_{\langle uv \rangle}$. We distinguish cases based on the pipe-degree of C and analyze their contribution to $N_\sigma - Q$ and N_s before and after $\text{pipeExpansion}(uv)$. The case when $\text{pipe-deg}(C) = 0$ is trivial, since C disappears and therefore its contribution to $N_\sigma - Q$ and N_s decreases.

Assume $\text{pipe-deg}(C) = 1$. Then $\sigma_C(uv) \geq 1$ by the definition of safe pipes. Step 1 of clusterExpansion (called inside $\text{pipeExpansion}(uv)$) creates component C' of pipe-degree 1 in some cluster on the boundary of $\Delta_{\langle uv \rangle}$ and no new pipe-edge in $\Delta_{\langle uv \rangle}$. Consequently, pipeExpansion decreases the contribution of C to N_σ and N_s to 0. The contribution of C to $N_\sigma - Q$ does not decrease (i.e., remains constant 0) only if $\sigma_C(uv) = 1$ and C is the only component in $G_{\langle uv \rangle}$. In this case, however, C contains precisely one component in each of G_u and G_v , of pipe-degree 1 and 2. The component of pipe-degree 2 is stable, since uv is safe and $\sigma_C(uv) = 1$, and so the number of stable components decreases by one, hence pipeExpansion decreases the contribution of C to Φ_1 .

Assume $\text{pipe-deg}(C) = 2$ (see Figure 2). Then step 2 of clusterExpansion (called inside $\text{pipeExpansion}(uv)$) creates two stable components connected by either a single edge or a thick edge in a pipe in $\Delta_{\langle uv \rangle}$. Consequently, $\text{pipeExpansion}(uv)$ changes the contribution of C to N_σ from at least 1 to precisely 1. The number of stable components increases if $s_C \in \{0, 1\}$, in which case at least one component of G_u and G_v in C is unstable. Since uv is useful, such a component is incident to at least two single or thick edges in uv , and so $\sigma_C(uv) \geq 2$. In this case, the contribution of C to N_σ decreases by at least one, while the number of stable components increases by at most two, hence Φ_1 strictly decreases. The contribution of C to Φ_1 is unchanged only if C contains precisely one component in each of G_u and G_v , where $\sigma_C(uv) = 1$ and $s_C = 2$ (that is, both are stable components).

Assume $\text{pipe-deg}(C) \geq 3$. Let W_k , $k \geq 4$, be the wheel created by $\text{clusterExpansion}(\langle uv \rangle)$. Since uv is a safe pipe, every component of G_u or G_v contained in C is incident to at most one single or thick pipe-edge outside of $\varphi^{-1}[uv]$. Hence, C contains at least $k - 2$ edges between these components of G_u and G_v (which are pipe-edges in uv), and its contribution to N_σ is at least $k - 2$. Step 3 of $\text{clusterExpansion}(uv)$ (called inside $\text{pipeExpansion}(uv)$) creates $k - 1$ components in distinct clusters of $\Delta_{\langle uv \rangle}$. Hence, the contribution of C to N_s increases by at most $k - 1$. Step 3 also creates $k - 1$ pipe-edges, each of which is the only pipe-edges in its pipe, hence they contribute zero to $N_\sigma - Q$. The difference between the new and the old value of $N_\sigma - Q$ due to C is at most $-(k - 2 - 1) = -k + 3$ (where the term -1 accounts for the case that C is the only component of $G_{\langle uv \rangle}$). Hence, we obtain the following upper bound on the difference between the new and the old contribution of C to Φ_1 : $-4(k - 3) + (k - 1) = 11 - 3k \leq 11 - 3 \cdot 4 = -1 < 0$, where the $(k - 1)$ -term corresponds to the increase in N_s . Consequently, the contribution of C to both, Φ_1 and $N_\sigma - Q$, decreases.

Summing over the contributions of all components of $G_{\langle uv \rangle}$, we have shown that neither Φ_1 nor $N_\sigma - Q$ increases. It remains to show that the Φ_1 decreases. By the previous case analysis, the contribution of a component C to neither Φ_1 nor $N_\sigma - Q$ increases. The contribution of C to Φ_1 is unchanged when C contains exactly one component in each G_u and G_v , and both are stable, and hence the contribution of C to N_s remains 2. For contradiction, assume that Φ_1 and $N_\sigma - Q$ are both unchanged. Then all components of $G_{\langle uv \rangle}$ are stable, and each contains precisely one (stable) component in G_u and G_v , respectively. Then N_s and N_σ remain unchanged. Therefore, Q must

also remain unchanged, which implies that all components of $G_{\langle uv \rangle}$ are adjacent to the same pair of pipes. Then, uv is useless, a contradiction. $\square$

COROLLARY 4.2. *The algorithm executes at most $15m$ iterations of Phase 2 and, therefore, terminates; and it creates at most $12m$ stable components.*

PROOF. First, we show that at the end of Phase 1, there are at most $3m$ pipe-edges. We charge the creation of $\text{pipe-deg}(C)$ pipe-edges in new pipes in step 3 of cluster expansion to $\text{pipe-deg}(C)$ original pipe-edges incident to C . Similarly we charge the creation of a (thick) pipe-edge created in step 2 to 1 or 3 (in the case we created a thick edge) original pipe-edges incident to C . In the latter case, C must be incident to at least 4 pipe-edges. Clearly, each original edge receives the charge of at most 2, one from each of its incident components. Hence, in total, we introduced at most $2m$ new pipe-edges.

By definition $N_\sigma - Q$ is a nonnegative integer, and its initial value does not exceed the number of pipe-edges. Therefore, at the end of Phase 1, we have $0 \leq N_\sigma - Q \leq 3m$. Since each pipe-edge can be incident to at most two stable components and each stable component must be incident to two pipe-edges, we have $0 \leq N_s \leq 3m$. Hence, $0 \leq \Phi_1 \leq 15m$.

By Lemma 4.1, Φ_1 strictly decreases in each iteration of the while loop in Phase 2. Therefore, Phase 2 has at most $15m$ iterations. Since both Phase 1 and Phase 3 consist of for-loops only, the algorithm terminates.

By Lemma 4.1, Φ_1 strictly decreases and $N_\sigma - Q$ does not increase in each iteration of the while loop in Phase 2. Consequently, we can charge the creation of a stable component to the decrease of one-fourth of a unit of $N_\sigma - Q$. Hence, Phase 2 creates at most $12m$ stable components. $\square$

Correctness. We show next that Phase 2 reduces G' to a collection of generalized cycles (see Lemma 4.6 below). First, we need a few observations.

LEMMA 4.3. *Every cluster u' created by an operation $\text{clusterExpansion}(u)$ satisfies the following:*

- (B1) u' is either empty or the base for some nonempty pipe.
- (B2) If u' lies in the interior of the disk Δ_u , then u' is either empty or the base of a unique nonempty pipe $u'v'$, where v' is on the boundary of Δ_u .
- (B3) If u' lies on the boundary of the disk Δ_u , then u' is either empty, or the base of at least one and at most two pipes, exactly one of which is outside of Δ_u . (In particular, the pipe outside of Δ_u is unique; this property will be crucial.)
- (B4) If u' is nonempty, then u' is incident to at most three empty pipes.

PROOF. Operation $\text{clusterExpansion}(u)$ creates a cluster u_v on the boundary of the disk Δ_u for every pipe uv incident to u . By construction, u_v is either empty (if the pipe uv was empty) or a base for the pipe vu_v , which lies outside of Δ_u . When u_v is nonempty, the only incident pipes that could be empty are created in step 0 (either two pipes of the cycle C_u or a single pipe if $\deg(u) = 2$). For each component C of G_u with $\text{pipe-deg}(C) \geq 3$, $\text{clusterExpansion}(u)$ creates a wheel where the center cluster is empty and every external cluster is incident to exactly one nonempty pipe (to a cluster on the boundary of Δ_u); consequently, it is a base for that pipe. The external clusters of the wheel are incident to exactly three empty pipes by (P4). By construction, clusters in $\partial\Delta_u$ are incident to at most two empty pipes. $\square$

COROLLARY 4.4. *In every step of Phase 2, every cluster $u \in V(H')$ is either empty or the base for at least one nonempty pipe.*

PROOF. Phase 1 of the algorithm performs $\text{clusterExpansion}(u)$ for every cluster of the input independently. At the end of Phase 1 (i.e., beginning of Phase 2), each cluster in H' has been created by a clusterExpansion operation, and (B1) follows from Lemma 4.3.

Subsequent steps of Phase 2 successively apply pipeExpansion operations, thereby creating new clusters. By Lemma 4.3, property (B1) is established for all new clusters, and it continues to hold for existing clusters. $\square$

LEMMA 4.5. *In every step of Phase 2, if H' contains a nonempty unsafe pipe, then it also contains a useful pipe.*

PROOF. Let u_0u_1 be a nonempty unsafe pipe in H' . Without loss of generality, assume that u_1 is not a base for u_0u_1 . We iteratively define a simple path $(u_0, u_1, \dots, u_\ell)$ for some $\ell \in \mathbb{N}$ as follows: Assume that $i \geq 1$, vertices $u_0, \dots, u_i$ have been defined, and $u_{i-1}u_i$ is a nonempty unsafe pipe in H' for which u_i is not a base. By (B1), u_i is a base for some nonempty pipe u_iw , where $w \neq u_{i-1}$. Put $u_{i+1} = w$. Since H is simple, this iterative process terminates when either u_iu_{i+1} is a nonempty safe pipe or $u_{i+1} = u_k$ for some $0 \leq k < i - 1$.

Case 1: The iterative process finds a nonempty safe pipe u_iu_{i+1} . We claim that u_iu_{i+1} is useful. Suppose, to the contrary, that u_iu_{i+1} is useless. Then u_i is incident to two nonempty pipes (which are necessarily $u_{i-1}u_i$ and u_iu_{i+1}), and every component in G_{u_i} has pipe-degree 2. Consequently, u_i is a base for both $u_{i-1}u_i$ and u_iu_{i+1} . This contradicts our assumption that u_i is not a base for $u_{i-1}u_i$ and proves the claim.

Case 2: The iterative process finds a cycle $U = (u_k, u_{k+1}, \dots, u_\ell)$ where $u_{\ell+1} = u_k$ and for every $i = k, \dots, \ell$, u_i is a base for u_iu_{i+1} but not a base for u_iu_{i-1} . We show that this case does not occur. All clusters in the cycle have been created by clusterExpansion operations (in Phase 1 or 2). We claim that not all clusters in U are created by the same clusterExpansion operation. For otherwise, by (B2) and (B3) u_i and u_{i+1} , for some i , are on the boundary of an expansion disk Δ_u , and u_{i+1} is not the base of u_iu_{i+1} by the construction of U . Furthermore, u_iu_{i+1} is nonempty by the definition of a base. By (B3), u_{i+1} is the base of a single pipe, which is outside of Δ_u , as the other base would have to be u_iu_{i+1} . Due to the previous claim and since U is a cycle, there are two consecutive clusters, u_j and u_{j+1} , in U such that u_j was created by an earlier invocation of clusterExpansion than u_{j+1} . By (B3), u_{j+1} is a base for u_ju_{j+1} . Indeed, $u_{i+1}u_i$ is the unique pipe at u_{i+1} outside of Δ_u for which u_{i+1} is also a base. This contradicts our assumption and proves that Case 2 does not occur. $\square$

LEMMA 4.6. *The following hold for the instance $\varphi' : G' \rightarrow H'$ at the end of Phase 2:*

- (1) *every pipe in $E(H')$ is empty or useless,*
- (2) *every component of G' is a generalized cycle, and*
- (3) *any two components of G' are mapped to the same or two vertex-disjoint cycles in H' .*

PROOF. 1. When the while loop of Phase 2 terminates, there are no useful pipes in H' . Consequently, every safe pipe is useless. By Lemma 4.5, every unsafe pipe is empty at that time. Overall, every pipe is empty or useless, as claimed.

2. Since every pipe $uv \in E(H')$ is empty or useless, every component in every cluster is stable. It follows that every component of G' must be a generalized cycle.

3. Since every pipe $uv \in E(H')$ is empty or useless, every cluster is incident to at most two nonempty pipes. Consequently, any two generalized cycles of G' are mapped to either the same cycle or two disjoint cycles in H' . $\square$

LEMMA 4.7. *The algorithm reports whether φ is a weak embedding.*

PROOF. By Lemmas 3.1 and 3.3, every operation either reports that the instance is not a weak embedding and halts or produces an instance equivalent to the input φ . Consequently, if any operation finds a negative instance, then φ is not a weak embedding. Otherwise the while loop in

Fig. 7. A nonplanar generalized cycle that does not embed in an annulus. It embeds in the Möbius band or the projective plane.

Phase 2 terminates and yields an instance $\varphi' : G' \rightarrow H'$ equivalent to φ . By Lemma 4.6, every component of G' is a generalized cycle, any two of which are mapped to the same or vertex-disjoint cycles in H' .

The strip system of the subgraph $\varphi'(C)$ of H' , where C is a generalized cycle, is homeomorphic to the annulus, since M is orientable. If C is nonplanar, then φ' is clearly not a weak embedding (see Figure 7). Furthermore, a generalized cycle C that winds around $\varphi'(C)$ several times cannot be embedded in that annulus (Figure 1(d)). However, one or more generalized cycles that each wind around $\varphi'(C)$ once can be embedded in nested annuli in the strip system $\mathcal{H}'$ (Figure 1(c)). This completes the proof of correctness of Algorithm(φ). $\square$

4.3 Efficient Implementation

Recall that $\text{pipeExpansion}(uv)$ first contracts the pipe uv into the new cluster $\langle uv \rangle \in V(H^*)$. Each stable component C of $G_{\langle uv \rangle}$ is composed of two stable components, in G_u and G_v , respectively. Then $\text{clusterExpansion}(uv)$ performed at the end of $\text{pipeExpansion}(uv)$ splits C into two stable components in two new clusters. That is, two adjacent stable components in G_u and G_v are replaced by two adjacent stable components in two new clusters. Consequently, we cannot afford to spend $O(|E_{\langle uv \rangle}|)$ time for $\text{pipeExpansion}(uv)$. We introduce auxiliary data structures to handle stable components efficiently: we use *set operations* to maintain a largest set of stable components in $O(1)$ time. A dynamic variant of the heavy path decomposition yields an $O(m \log m)$ bound on the total time spent on stable components in the main loop of the algorithm.

Data structures for stable components. For each pair (u, uv) of a cluster $u \in V(H)$ and an incident pipe $uv \in E(H)$, we store a set $L(u, uv)$ of all stable components of G_u adjacent to a (thick) edge in $\varphi^{-1}[uv]$. For every cluster $u \in V(H)$, let $w^*(u)$ be a neighbor of u maximizing the size of $L(u, uw^*(u))$; we maintain a pointer from u to the set $L(u, uw^*(u))$. The total number of sets $L(u, uv)$ and the sum of their sizes are $O(m)$. We can initialize them in $O(m)$ time. For each cluster $u \in V(H)$, we store the components of G_u in two sets: a set of stable and unstable components, respectively, each stored as the ID of a representative edge of the component. Hence, for each pipe, we can determine in $O(1)$ time whether uv is useful or useless. Every stable component in G_u has a pointer to the two sets $L(u, \cdot)$ in which it appears.

In each iteration of Phase 2, we implement $\text{pipeExpansion}(uv)$ for a useful pipe uv as follows: Let $G_{\langle uv \rangle}^\circ$ be the graph formed by the unstable components of $G_{\langle uv \rangle}$ and stable components C where $\sigma_C(uv) > 1$ (such as paths that “zig-zag” between u and v). Compute $G_{\langle uv \rangle}^\circ$ using DFS starting from each unstable component of G_u and G_v . If a stable component of G_u and G_v is absorbed by a new unstable component of $G_{\langle uv \rangle}$, delete it from the set $L(u, \cdot)$ or $L(v, \cdot)$ in which it appears.

(a) We perform simplify and clusterExpansion only on $G_{\langle uv \rangle}^\circ$.

(b) We handle the remaining stable components of $G_{\langle uv \rangle}$ as follows (refer to Figure 8): Since uv is safe, every stable component of G_u and G_v appears in $L(u, uv)$ and $L(v, uv)$, respectively. For all w ,

Fig. 8. Implementation of $\text{pipeExpansion}(uv)$ for stable components.

$w \notin \{v, w^*\}$, where $w^* = w^*(u)$, move all components from $L(u, uw)$ to the cluster $\langle uv \rangle_w$. Similarly, for all $x, x \notin \{u, x^*\}$, where $x^* = w^*(v)$, move all components from $L(v, vx)$ to the cluster $\langle uv \rangle_x$. Let $L'(u, uv)$ and $L'(v, uv)$ be the sets obtained from $L(u, uv)$ and $L(v, uv)$ after this step. Notice that the remaining stable components of G_u and G_v in $L'(u, uv)$ and $L'(v, uv)$ appear in $L(u, uw^*)$ and $L(v, vx^*)$, respectively. Move them to the clusters $\langle uv \rangle_{w^*}$ and $\langle uv \rangle_{x^*}$ by renaming all data structures of u and v , respectively, which can be done in $O(1)$ time. Update H so it reflects the changes in G by creating appropriate pipes between the clusters in $\partial\Delta_{\langle uv \rangle}$, checking for crossings by local planarity testing inside $\Delta_{\langle uv \rangle}$ (similar to step 4 of clusterExpansion).

(c) We add the new stable components (created by clusterExpansion from unstable components) that fit the definition of $L(\langle uv \rangle_{w^*}, \langle uv \rangle_{w^*} \langle uv \rangle_{x^*})$ to the current set $L'(u, uv)$ to obtain the set $L(\langle uv \rangle_{w^*}, \langle uv \rangle_{w^*} \langle uv \rangle_{x^*})$. Analogously, add new stable components to the current sets $L'(v, uv)$, $L(u, uw)$ where $w \neq v$, and $L(v, vx)$ where $x \neq u$ to obtain the new sets $L(\langle uv \rangle_{x^*}, \langle uv \rangle_{x^*} \langle uv \rangle_{w^*})$, $L(\langle uv \rangle_w, \langle uv \rangle_w w)$, and $L(\langle uv \rangle_x, \langle uv \rangle_x x)$, respectively. Compute any other sets of stable components from scratch.

We are now ready to show that the running time of the algorithm improves to $O(m \log m)$.

LEMMA 4.8. *Our implementation of the algorithm runs in $O(m \log m)$ time.*

PROOF. Phases 1 and 3 take $O(m)$ time by Lemmas 4.1, 3.1, and 4.6. The while loop in Phase 2 terminates after $O(m)$ iterations by Corollary 4.2. Using just Lemmas 2.1 and 3.1, each iteration of Phase 2 would take $O(m)$ time, leading to an overall running time of $O(m^2)$. We define a new potential function for an instance $\varphi : G \rightarrow H$ to show that each iteration of Phase 2 takes $O(\log m)$ amortized time.

For every $u \in V(H)$, let $\mathcal{L}(u)$ be the number of stable components in G_u . Let s be the number of stable components created from the beginning of Phase 2 up to the current iteration. We define a new potential function as¹

$$\Phi_2(\varphi) = \Phi_1(\varphi) + (12m - s) \log(48m) + \sum_{u \in V(H)} \mathcal{L}(u) \log \mathcal{L}(u).$$

By Corollary 4.2, the second term is nonnegative. Note that $\Phi_2(\varphi) = O(m \log m)$, since $\Phi_1(\varphi) = O(m)$ and $\sum_{u \in V(H)} \mathcal{L}(u) = O(m)$. We show that Φ_2 strictly decreases in Phase 2. As argued above (cf. Lemma 4.1), Φ_1 strictly decreases. The second term of Φ_2 can only decrease, since s increments when stable components are created (but never decrements). The term $\sum_{u \in V(H)} \mathcal{L}(u) \log \mathcal{L}(u)$ increases when new stable components are created. However, this increase is offset by the decrease in the second term of Φ_2 . It suffices to consider the case that $\mathcal{L}(u)$ increments

¹ All logarithms are of base 2.

from k to $k + 1$. Then

$$\begin{aligned}
 k \log k + \log(48m) &\geq k \log k + \log(4s) \\
 &\geq \left(k \log(k + 1) - k \log \frac{k + 1}{k} \right) + (2 + \log(k + 1)) \\
 &\geq (k + 1) \log(k + 1) + 2 - \log \left(1 + \frac{1}{k} \right)^k \\
 &\geq (k + 1) \log(k + 1) + 2 - \log e \\
 &> (k + 1) \log(k + 1),
 \end{aligned}$$

that is, the decrease of $\log(48m)$ in the second term offsets the increase of $(k + 1) \log(k + 1) - k \log k$ in the third term.

We next show that the time spent on each iteration of the while loop in Phase 2 is bounded above by a constant times the decrease of the potential Φ_2 . This will complete the proof, since the potential Φ_2 is nonincreasing throughout the execution of the algorithm, as we have just shown.

For a useful pipe $uv \in E(H)$, let $\Lambda(uv) = \Lambda$ be the collection of sets consisting of all $L(u, uw)$ where $w \notin \{v, w^*(u)\}$ and all $L(v, vx)$ where $x \notin \{u, w^*(v)\}$. Our implementation of $\text{pipeExpansion}(uv)$ spends $O(|E(G_{\langle uv \rangle}^\circ)| + \deg(\langle uv \rangle))$ time to process the components of $G_{\langle uv \rangle}^\circ$, by Lemma 3.1, and $O(1 + \sum_{L \in \Lambda} |L|)$ time to process the remaining stable components. However, by (B4), $\deg(\langle uv \rangle) = O(|E(G_{\langle uv \rangle}^\circ)| + 1 + \sum_{L \in \Lambda} |L|)$. Hence, the running time of $\text{pipeExpansion}(uv)$ is

$$O\left(|E(G_{\langle uv \rangle}^\circ)| + 1 + \sum_{L \in \Lambda} |L|\right). \quad (3)$$

First, let C be a component of $G_{\langle uv \rangle}^\circ$. We have seen (in the proof of Lemma 4.1) that $\text{pipeExpansion}(uv)$ decreases N_σ by at least $\sigma_C(uv) - 1$ due to edges in C . By the definition of safe pipes, C contains a nonempty set A of components of G_u and a nonempty set B of components of G_v . By Lemma 2.4, each component $C' \in A \cup B$ contains $O(t_{C'})$ edges, where $t_{C'}$ is the number of terminals of C' . Notice that $\sigma_C(uv) \geq \sum_{C' \in A} (t_{C'} - 3)$, because at most three out of $t_{C'}$ terminals are not adjacent to a pipe-edge in $\varphi^{-1}[uv]$ by the definition of safe pipes (equality occurs when a thick edge is incident to C' in a pipe other than uv). As such, for a component C in $G_{\langle uv \rangle}^\circ$, the contribution of C to $N_\sigma - Q$ decreases by $\Omega(|E(C)|)$. Therefore, the first term of (3) is charged to the decrease in $N_\sigma - Q$. By Lemma 4.1, summation over pipe expansions in Phase 2 yields $O(m)$. This takes care of steps (a) and (c) of the efficient implementation (Section 4.3). It remains to bound the time complexity of step (b), which deals exclusively with stable components.

Second, we show that the time $\text{pipeExpansion}(uv)$ spends on stable components is absorbed by the decrease in the last term of Φ_2 . When we move the components of $L(u, uw)$, $w \notin \{v, w^*\}$ to the cluster $\langle uv \rangle_w$, we spend linear time on all but a maximal set, which can be moved in $O(1)$ time using a set operation. In what follows, we show that a constant times the corresponding decrease in the term $\sum_{u \in V(H)} \mathcal{L}(u) \log \mathcal{L}(u)$ subsumes this work.

We adapt the analysis from the classic heavy path decomposition. Suppose we partition a set of size $k = \mathcal{L}(u) = \mathcal{L}(v)$ into ℓ subsets of sizes $k_1 \geq \dots \geq k_\ell$. Note that $k_j \leq k/2$ for $j \geq 2$. Then

$$k \log k = \sum_{i=1}^{\ell} k_i \log k \geq k_1 \log k_1 + \sum_{j=2}^{\ell} k_j \log(2k_j) = \sum_{i=1}^{\ell} k_i \log k_i + \sum_{j=2}^{\ell} k_j.$$

Hence, the decrease in $\sum_{u \in V(H)} \mathcal{L}(u) \log \mathcal{L}(u)$, which is equal to $k \log k - \sum_{i=1}^{\ell} k_i \log k_i$, is bounded from below by $k - k_1$. Therefore, if we spend $O(1)$ time on a maximal subset (of size k_1), we can

afford to spend linear time on all other subsets. Thus, the decrease in $\sum_{u \in V(H)} \mathcal{L}(u) \log \mathcal{L}(u)$ subsumes the actual work and this concludes the proof. $\square$

This completes the proof of Theorem 1.1(i). Part (ii) of Theorem 1.1 is shown in Section 5.

5 CONSTRUCTING AN EMBEDDING

Our recognition algorithm in Section 4 decides in $O(m \log m)$ time whether a given instance φ is a weak embedding. However, if φ turns out to be a weak embedding, it does not provide an embedding ψ_φ , since at the end of the algorithm, we have an equivalent “reduced” instance φ' at hand. In this section, we show how to compute the combinatorial representation of an embedding ψ_φ for the input φ .

Assume that $\varphi : G \rightarrow H$ is a weak embedding. By Lemmas 4.7 and 4.8, we can obtain a combinatorial representation of an embedding $\pi_{\varphi'} \in \Pi(\varphi')$ of the instance $\varphi' : G' \rightarrow H'$ produced by the algorithm at the end of Phase 2 in $O(n \log n)$ time.

We sequentially reverse the steps of the algorithm and maintain combinatorial embeddings for all intermediate instances until we obtain a combinatorial representation $\pi \in \Pi(\varphi)$. By Lemma 2.2, we can then obtain an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$ in $O(m)$ time. Reversing a `clusterExpansion(u)` operation is trivial: the total orders of pipes within Δ_u can be ignored, and the total order for every pipe uv is the same as the order for $u_v v$. This can be done in $O(\deg(u))$ time.

Let $\varphi^{(1)} : G^{(1)} \rightarrow H^{(1)}$ be the input instance of `pipeExpansion(uv)`, $\varphi^{(2)} : G^{(2)} \rightarrow H^{(2)}$ be the instance obtained by contracting uv and $\varphi^{(3)} : G^{(3)} \rightarrow H^{(3)}$ be the instance after `clusterExpansion( $\langle uv \rangle$ )`. By the previous argument, the total orders of pipe-edges $\pi^{(2)}(\langle uv \rangle w)$ of all pipes $\langle uv \rangle w$ can be obtained from a combinatorial representation $\pi^{(3)}$ in $O(\deg(\langle uv \rangle))$ time. These orders also correspond to $\pi^{(1)}(uw)$ and $\pi^{(1)}(vx)$ for pipes uw and vx in $\varphi^{(1)}$ where $w \neq v$ and $x \neq u$.

To obtain an order $\pi^{(1)}(uv)$, we embed $G_{\langle uv \rangle}^{(1)}$ into $D_{\langle uv \rangle}$ using Lemma 2.2 in $O(|E(G_{\langle uv \rangle}^{(1)})|)$ time and find the Jordan curve defined in the proof of Lemma 3.3. The order $\pi^{(1)}(uv)$ of the pipe-edges in $\varphi^{-1}[uv]$ is given by the order in which the Jordan curve intersects these edges. This takes $O(|E(G_{\langle uv \rangle}^{(1)})|)$ time by Theorem 3.2, though we first need to obtain an embedding of $G^{(1)}$, where triangles of wheels in $G_{\langle uv \rangle}^{(1)}$ and 4-cycles induced by thick edges in and incident to $G_{\langle uv \rangle}^{(1)}$ are empty. This can be done by changing the rotation at the vertices of the wheels and 4-cycles corresponding to thick edges one-by-one in $O(|E(G_{\langle uv \rangle}^{(1)})|)$ time, since uv is safe in the resulting instance.

However, this would lead to a $O(m^2)$ worst-case time complexity because of stable components. We show how to reduce the running time to $O(m \log m)$. Let us call a pipe-edge (thick edge) **stable** if it connects two stable components in two adjacent clusters. In each total order $\pi^*(uv)$ for a pipe uv in an instance φ^* , we arrange maximal blocks of consecutive stable edges into a **bundle** that takes a single position in the order, and we store the order among the stable edges in the bundle in a separate linked list. We can substitute each bundle with one **representative** stable component. Then, using $\pi^{(2)}(\langle uv \rangle w)$ for all pipes $\langle uv \rangle w$ incident to $\langle uv \rangle$, we can obtain a list of at most $\deg(\langle uv \rangle) + c$ representative stable components, where c is the number of components of $G_{\langle uv \rangle}^\circ$. We can proceed by embedding all components in $G_{\langle uv \rangle}^\circ$ and the representatives of the remaining stable components in $D_{\langle uv \rangle}$. Obtaining a Jordan curve that encloses all vertices in $(\varphi^{(1)})^{-1}[u]$ now takes $O(E(G_{\langle uv \rangle}^\circ) + \deg(\langle uv \rangle))$ time. The order in which the Jordan curve crosses the edges in $G_{\langle uv \rangle}$ defines $\pi^{(1)}(uv)$, where the size of $\pi^{(1)}(uv)$ is $O(E(G_{\langle uv \rangle}^\circ) + \deg(\langle uv \rangle))$ and each pipe-edge obtained from a representative stable component represents a bundle of stable edges. We can merge consecutive bundles of stable edges in $\pi^{(1)}(uv)$, as needed, in $O(E(G_{\langle uv \rangle}^\circ) + \deg(\langle uv \rangle))$ time. By (B4), this

running time is bounded above by the running time of our implementation of `pipeExpansion`(uv), as argued in the proof of Lemma 4.8. Therefore, we can reverse every operation and obtain an embedding $\psi_\varphi : G \rightarrow \mathcal{H}$ in $O(m \log m)$ time. This completes the proof of Theorem 1.1(ii).

6 ALGORITHM FOR NONORIENTABLE SURFACES

We show that our algorithm can be adapted to recognize weak embeddings $\varphi : G \rightarrow H$ when H is embedded in a *nonorientable* manifold M . We discuss the adaptation to nonorientable surfaces in a separate section to reduce notational clutter in Sections 2–5.

First, we adapt the definition of the strip system. The embedding of a graph H into a (orientable or nonorientable) surface M is given by a rotation system that specifies, for each vertex of H , the ccw cyclic order of incident edges and a signature $\lambda : E(H) \rightarrow \{-1, 1\}$. To define the strip system $\mathcal{H}$, we proceed exactly as in the case that M is orientable (Section 2), except that for every edge $e = uv$, if $\lambda(e) = -1$, we identify $A_{u,v}$ with ∂R_{uv} via an orientation-reversing homeomorphism and $A_{v,u}$ with ∂R_{uv} via an orientation-preserving homeomorphism, or vice versa. If we represent M as a sphere with a finite number of holes that are turned into **cross-caps**, the signature of an edge is interpreted as the parity of the number of times an edge passes through a cross-cap.

Second, we adapt the operation of cluster expansion as follows: We put $\lambda(u_v v) := \lambda(uv)$ for all neighbors v of u , and $\lambda(e) := 1$ for all newly created edges e .

Third, we adapt the operation `pipeExpansion`(uv) as follows: If $\lambda(uv) = -1$, before creating the cluster $\langle uv \rangle$, we flip the value of λ from -1 to 1 , and vice versa, for every edge of H adjacent to u . This corresponds to pushing the edge uv off all the cross-caps that it passes through. Then the values of λ on the edges incident to $\langle uv \rangle$ in H^* are naturally inherited from the values of λ on the edges adjacent to u and v in H . The value of λ for all other edges in H^* remain the same as in H .

The first two phases of the algorithm remain the same except that they use the adapted operations of cluster and pipe expansion. Phase 3 is modified as follows: If a generalized cycle C in G' satisfies $\prod_{e \in E(\varphi'(C))} \lambda(e) = 1$ (that is, the strip system of $\varphi'(C)$ is homeomorphic to an annulus), we proceed as in the orientable case. The modifications affect only the generalized cycles C such that $\prod_{e \in E(\varphi'(C))} \lambda(e) = -1$; or in other words, generalized cycles C , for which the strip system of $\varphi'(C)$ is homeomorphic to the Möbius band. For such a generalized cycle C , we report that the instance is negative if C winds more than two times around $\varphi'(C)$; or if all pipe edges of C are thick edges, the underlying graph of C is planar (resp., not planar), and C winds exactly once (resp., twice) around $\varphi'(C)$. Furthermore, we report that the instance is negative if there exist two distinct generalized cycles C_1 and C_2 in G' winding once around $\varphi'(C)$ such that $\varphi'(C) = \varphi'(C_1) = \varphi'(C_2)$ in H' . Else, we can report that φ is a weak embedding at the end.

The correctness of the algorithm in the Möbius band case is implied by a stronger statement in Reference [15, Section 8], but we can easily verify it by the following contradictions: Suppose two or more cycles wind exactly once around $\varphi'(C)$. In any embedding φ , the total orders in π (defined in Section 2.3) reverse in each traversal of the cycle, which is a contradiction. Similarly, if C winds $k > 2$ many times around $\varphi'(C)$, we divide it into k paths $P_1, \dots, P_k$ sharing end vertices each of which is mapped injectively into $\varphi'(C)$. We assume that in π the paths $P_1, \dots, P_k$ appear in the given order up to the choice of orientation. The order of P_i 's in π reverses with respect to a fixed orientation if we traverse $\varphi'(C)$. Hence, P_1 must precede and follow P_k along C and therefore $k \leq 2$, a contradiction.

7 CONCLUSIONS

We have shown (Theorem 1.1) that it takes $O(m \log m)$ time to decide whether a piecewise linear simplicial map $\varphi : G \rightarrow H$ from an abstract graph G with m edges to an embedded graph H is a weak embedding (i.e., whether, for every $\varepsilon > 0$, there exists an embedding $\psi_\varepsilon : G \rightarrow \mathcal{H}$ of G into

a neighborhood $\mathcal{H}$ of H such that $\|\varphi - \psi_\varepsilon\| < \varepsilon$). The only previously known algorithm for this problem takes $O(n^{2\omega}) \leq O(n^{4.75})$ time, where ω is the matrix multiplication constant [15], and until recently no polynomial-time algorithm was available even in the special case that H is embedded in the plane. Only the trivial lower bound of $\Omega(m)$ is known for the time complexity of recognizing weak embeddings in our setting. Closing the gap between $\Omega(m)$ and $O(m \log m)$ remains open.

If $\varphi : G \rightarrow H$ is a continuous map, but not necessarily simplicial (i.e., an image of an edge may pass through vertices of H), then the running time increases to $O(mn \log(mn))$, where $n = |V(G)|$ (cf. Corollary 1.2). In the special case that G is a cycle and H is a planar straight-line graph, an $O(m \log m)$ -time algorithm was recently given in Reference [1]. It remains an open problem whether a similar improvement is possible for arbitrary G and H .

Atomic Embeddings. An interesting generalization of our problem is deciding whether a given 2-dimensional simplicial complex embeds into some 3-dimensional manifold, also known as the *thickenability of 2-dimensional simplicial complexes*. Indeed, Reference [22, Lemma] implies² that a polynomial-time algorithm for the thickenability of 2-dimensional simplicial complexes would directly translate into a polynomial-time algorithm for our problem. Studying this more general problem is one of the next natural steps in our investigation: Currently, we do not know whether this problem is tractable. Furthermore, due to the same argument, a polynomial-time algorithm for deciding whether a 2-dimensional simplicial complex embeds in $\mathbb{R}^3$ would already imply a polynomial-time algorithm for our problem if we restrict ourselves to orientable surfaces. However, this problem is NP-hard [9], so the existence of such an algorithm is highly unlikely.

The thickenability problem, originally due to Neuwirth [22], can be seen as the following variant of the problem of deciding whether a simplicial map $\varphi : G \rightarrow H$ is a weak embedding,³ which we call the **atomic embeddability problem**: We are given an abstract graph H and a strip system $\mathcal{H}$ as a 2-dimensional surface without boundary, partitioned into regions representing “clusters” and “pipes.” Every vertex $u \in V(H)$ is represented by a surface S_u with boundary (instead of a disk D_u), obtained from a 2-sphere by removing disjoint open disks (holes), where the number of holes equals $\deg(u)$ and each hole corresponds to an edge in H incident to u . Every edge $uv \in E(H)$ is associated with a cylinder C_{uv} (rather than a rectangle R_{uv}) whose boundary components are homeomorphically identified with the boundaries of the corresponding holes on S_u and S_v , respectively. The technique that we developed in Sections 2–4 is not directly applicable to the setting of atomic embeddings. It remains an open problem to recognize atomic embeddability efficiently.

ACKNOWLEDGMENTS

A preliminary version of this article appeared in the *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms* (SODA), 2018, pp. 274–292, <https://doi.org/10.1137/1.9781611975031.20>.

REFERENCES

- [1] Hugo A. Akitaya, Greg Aloupis, Jeff Erickson, and Csaba D. Tóth. 2017. Recognizing weakly simple polygons. *Discrete Comput. Geom.* 58 (2017), 785–821. Issue 4. DOI : <https://doi.org/10.4230/LIPIcs.SoCG.2016.8>
- [2] Patrizio Angelini and Giordano Da Lozzo. 2019. Clustered planarity with pipes. *Algorithmica* 81, 6 (2019), 2484–2526. DOI : <https://doi.org/10.1007/s00453-018-00541-w>
- [3] Patrizio Angelini, Giordano Da Lozzo, Giuseppe Di Battista, and Fabrizio Frati. 2017. Strip planarity testing for embedded planar graphs. *Algorithmica* 77, 4 (2017), 1022–1059. DOI : <https://doi.org/10.1007/s00453-016-0128-9>

²The lemma in Reference [22] is stated only for connected graphs and only in the case that the target surface is a sphere. However, it is easy to see that an analogous statement holds for disconnected graphs and for orientable surfaces of arbitrary genus.

³In fact, for Neuwirth’s algorithm it is enough to assume that G is a finite union of pairwise disjoint 3-cycles.

- [4] S. B. Belyi. 1983. Self-nonintersecting and non intersecting chains. *Math. Notes Acad. Sci. USSR* 34, 4 (1983), 802–804. DOI : <https://doi.org/10.1007/BF01157400> Translated from *Matematicheskie Zametki*, Vol. 34, No. 4, pp. 625–628, 1983.
- [5] Henning Bruhn and Reinhard Diestel. 2009. MacLane’s theorem for arbitrary surfaces. *J. Combin. Theory Ser. B* 99 (2009), 275–286.
- [6] Hsien-Chih Chang, Jeff Erickson, and Chao Xu. 2015. Detecting weakly simple polygons. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms (SODA’15)*. 1655–1670. DOI : <https://doi.org/10.1137/1.9781611973730.110> Full paper available at [arXiv:1407.3340](https://arxiv.org/abs/1407.3340).
- [7] Pier Francesco Cortese, Giuseppe Di Battista, Maurizio Patrignani, and Maurizio Pizzonia. 2009. On embedding a cycle in a plane graph. *Discrete Math.* 309, 7 (2009), 1856–1869. DOI : <https://doi.org/10.1016/j.disc.2007.12.090>
- [8] Pier Francesco Cortese and Maurizio Patrignani. 2018. Clustered planarity = flat clustered planarity. In *Proceedings of the 26th Symposium on Graph Drawing and Network Visualization (GD’18) (LNCS)*, Therese C. Biedl and Andreas Kerren (Eds.), Vol. 11282. Springer, Cham, 23–38. DOI : https://doi.org/10.1007/978-3-030-04414-5_2
- [9] Arnaud de Mesmay, Yoav Rieck, Eric Sedgwick, and Martin Tancer. 2018. Embeddability in $\mathbb{R}^3$ is NP-hard. In *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms (SODA’18)*. 1316–1329. Retrieved from: <http://dl.acm.org/citation.cfm?id=3174304.3175356>.
- [10] Giuseppe Di Battista and Roberto Tamassia. 1996. On-line planarity testing. *SIAM J. Comput.* 25 (1996), 956–997.
- [11] Reinhard Diestel. 2017. *Graph Theory* (5th ed.). Graduate Texts in Mathematics, Vol. 173. Springer, Heidelberg.
- [12] Qing-Wen Feng, Robert F. Cohen, and Peter Eades. 1995. How to draw a planar clustered graph. In *Proceedings of the 1st Conference on Computing and Combinatorics (COCOON’95)*. LNCS, Vol. 959. Springer, Berlin, 21–30. DOI : <https://doi.org/10.1007/BFb0030816>
- [13] Qing-Wen Feng, Robert F. Cohen, and Peter Eades. 1995. Planarity for clustered graphs. In *Proceedings of the 3rd European Symposium on Algorithms (ESA’95) (LNCS)*, Vol. 979. Springer, Berlin, 213–226.
- [14] Radoslav Fulek. 2017. Embedding graphs into embedded graphs. In *Proceedings of the 28th International Symposium on Algorithms and Computation (ISAAC’17) (Leibniz International Proceedings in Informatics (LIPIcs))*, Yoshio Okamoto and Takeshi Tokuyama (Eds.), Vol. 92. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 34:1–34:12. DOI : <https://doi.org/10.4230/LIPIcs.ISAAC.2017.34>
- [15] Radoslav Fulek and Jan Kynčl. 2018. Hanani-Tutte for approximating maps of graphs. In *Proceedings of the 34th Symposium on Computational Geometry (LiPICS)*, Vol. 99. Schloss Dagstuhl. Retrieved from: [arXiv:1705.05243](https://arxiv.org/abs/1705.05243).
- [16] Haim Hanani. 1934. Über wesentlich unplättbare Kurven im drei-dimensionalen Raume. *Fund. Math.* 23 (1934), 135–142.
- [17] John Hopcroft and Robert Tarjan. 1974. Efficient planarity testing. *J. ACM* 21 (1974), 549–568.
- [18] Michael Jünger, Sebastian Leipert, and Petra Mutzel. 1998. Level planarity testing in linear time. In *Proceedings of the 6th Symposium on Graph Drawing (GD’98)*. LNCS, Vol. 1547. Springer, Berlin, 224–237.
- [19] Bojan Mohar. 1999. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM J. Discrete Math.* 12, 1 (1999), 6–26.
- [20] Dušan Repovš and Arkadij B. Skopenkov. 1998. A deleted product criterion for approximability of maps by embeddings. *Topology Appl.* 87, 1 (1998), 1–19. DOI : [https://doi.org/10.1016/S0166-8641\(97\)00121-1](https://doi.org/10.1016/S0166-8641(97)00121-1)
- [21] Karol Sieklucki. 1969. Realization of mappings. *Fund. Math.* 65 (1969), 325–343.
- [22] Arkadij B. Skopenkov. 1995. A generalization of Neuwirth’s theorem on thickening 2-dimensional polyhedra. *Math. Notes* 58, 5 (1995), 1244–1247.
- [23] Mikhail Skopenkov. 2003. On approximability by embeddings of cycles in the plane. *Topology Appl.* 134, 1 (2003), 1–22. DOI : [https://doi.org/10.1016/S0166-8641\(03\)00069-5](https://doi.org/10.1016/S0166-8641(03)00069-5)
- [24] Carsten Thomassen. 1989. The graph genus problem is NP-complete. *J. Algor.* 10, 4 (1989), 568–576.
- [25] William T. Tutte. 1970. Toward a theory of crossing numbers. *J. Combin. Theory* 8 (1970), 45–53.
- [26] Virginia Vassilevska Williams. 2012. Multiplying matrices faster than Coppersmith-Winograd [extended abstract]. In *Proceedings of the 44th Symposium on Theory of Computing (STOC’12)*. ACM Press, New York, NY, 887–898. DOI : <https://doi.org/10.1145/2213977.2214056>

Received July 2018; revised June 2019; accepted June 2019