

CANONICAL DECOMPOSITIONS AND ALGORITHMIC RECOGNITION OF SPATIAL GRAPHS

STEFAN FRIEDL¹ , LARS MUNSER¹ , JOSÉ PEDRO QUINTANILHA²  AND
YURI SANTOS REGO³ 

¹*Fakultät für Mathematik, Universität Regensburg, Regensburg, Germany*
(stefan.friedl@ur.de; lars.munser@ur.de)

²*Fakultät für Mathematik, Universität Bielefeld, Bielefeld, Germany*
(jquintan@math.uni-bielefeld.de)

³*Fakultät für Mathematik (IAG), Otto-von-Guericke-Universität Magdeburg,
Magdeburg, Germany* (yuri.santos@ovgu.de)

(Received 2 May 2023)

Abstract We prove that there exists an algorithm for determining whether two piecewise-linear spatial graphs are isomorphic. In its most general form, our theorem applies to spatial graphs furnished with vertex colourings, edge colourings and/or edge orientations.

We first show that spatial graphs admit canonical decompositions into blocks, that is, spatial graphs that are non-split and have no cut vertices, in a suitable topological sense. Then, we apply a result of Haken and Matveev in order to algorithmically distinguish these blocks.

Keywords: spatial graphs; 3-manifolds with boundary pattern; Haken manifolds;
piecewise-linear topology

2020 Mathematics subject classification: Primary 57M15; 57-08; 57Q35
Secondary 57K10; 57Q40

1. Introduction

1.1. Our main result

This article is concerned with spatial graphs, which are finite graphs embedded in oriented 3-spheres – not merely as subspaces, but with an explicit decomposition into vertices and edges. They thus generalize knots and links. We give a precise definition using piecewise-linear (PL) topology (Definition 2.1), which makes it easy to formalize the notion of an isomorphism of spatial graphs: an orientation-preserving PL homeomorphism of the ambient 3-spheres mapping vertices to vertices and edges to edges bijectively (Definition 2.2). The reader might be amused to try and decide which of the pairs of spatial graphs in Figure 1.1 are isomorphic.






Figure 1.1. Two pairs of (decorated) spatial graphs with the same combinatorial structure, but *a priori* different topology.

Our main result is the following:

Theorem 1.1 (Algorithmic recognition of spatial graphs). *There exists an algorithm for determining whether two spatial graphs are isomorphic.*

Theorem 1.1 is restated and proved in a more general context as Theorem 7.13, where we allow spatial graphs to come equipped with colourings of vertices/edges and/or edge orientations (as in the right hand side of Figure 1.1). These decorations must of course be respected by isomorphisms.

We reduce the task of testing whether two spatial graphs are isomorphic, to the application of algorithms in Matveev’s text on computational 3-manifold topology [14]. An algorithm for such a test would then take as input a pair of oriented 3-spheres given as finite simplicial complexes, with the vertices and edges of the spatial graphs specified as subcomplexes, and also possibly the data of decorations.

Much of the article is devoted to setting up a rigorous and self-contained theory of spatial graphs, with the algorithmic aspects introduced only towards the end. We emphasize that our proof of Theorem 1.1, alongside Matveev’s text, actually supplies an explicit algorithm. Even though we make no claims about its computational efficiency, we believe the existence of such an algorithm to be of independent theoretical interest. It is also conceivable that our result might have applications to other decision problems, namely when they involve comparing objects that can be encoded as spatial graphs.

1.2. Main idea of the proof

Our proof uses as a main ingredient a ‘Recognition Theorem’ of Matveev [14, Theorem 6.1.6] that extends work of Haken [7], Johannson [8] and others, concerning PL 3-manifolds equipped with a 1-dimensional subcomplex of their boundary called a ‘boundary pattern’ (Definition 6.4). The Recognition Theorem (reproduced below as Theorem 7.4) states that it is possible to algorithmically detect whether two such 3-manifolds with boundary pattern are PL-homeomorphic (via a homeomorphism respecting the boundary patterns), provided that they are ‘Haken’ (Definition 7.2).

We construct a PL 3-manifold with boundary pattern out of a spatial graph, its ‘marked exterior’ (Definition 6.5), such that two spatial graphs are isomorphic precisely if their marked exteriors are PL-homeomorphic. The marked exterior is built in a two-step process by first removing a suitably chosen open neighbourhood of the vertices, and then one of what is left of the edges. The boundary is then marked with a pattern that allows

for easily reconstructing the spatial graph; Figures 6.1 and 6.3 illustrate the general idea of the construction. The main feature of this boundary pattern is that it encodes the curves playing the role of meridians, bypassing the usual difficulties in recognizing knots and links from their exteriors (see Proposition 6.3 and the ensuing discussion).

We then wish to apply the Recognition Theorem to the marked exteriors to deduce Theorem 1.1, but face the difficulty that it applies only to 3-manifolds with boundary pattern that are Haken. This property encompasses, besides an easy to guarantee technical condition, a requirement about triviality of embedded 2-spheres (irreducibility) and one about triviality of properly embedded discs (boundary-irreducibility). As we shall see, the marked exterior of a spatial graph may very well fail to be irreducible and boundary-irreducible, and this issue will take considerable effort to resolve.

1.3. Decomposition results

To overcome the aforementioned difficulties, we establish a decomposition theory of spatial graphs (§ 3 and 4). Consider the irreducibility requirement. The starting point is the observation (Proposition 6.9) that irreducibility of the marked exterior of a spatial graph Γ is equivalent to it not being the ‘disjoint union’ $\Gamma_1 \sqcup \Gamma_2$ of non-empty spatial sub-graphs, where this disjoint union is the operation of placing Γ_1, Γ_2 ‘next to one another’ in the same ambient 3-sphere (Definition 3.2). A non-empty spatial graph that is not a non-trivial disjoint union is called a ‘piece’ (Definition 3.9). In Proposition 3.4, we show that the decomposition of a spatial graph into pieces is canonical in a suitable sense. This reduces the task of determining whether two spatial graphs are isomorphic, to testing whether the pieces in their decompositions are pairwise isomorphic.

The next step is to find a decomposition of non-split graphs into spatial graphs whose marked exteriors are moreover boundary-irreducible. Here the role of the disjoint union is played by the operation of ‘vertex sum’ (Definition 4.1). Roughly, the vertex sum of two spatial graphs, each with a distinguished vertex, is obtained by ‘gluing them’ along those vertices. For non-split spatial graphs Γ , there is a close correspondence between Γ having boundary-irreducible marked exterior, and Γ being indecomposable as a non-trivial vertex sum (Propositions 6.11 and 6.14). We show that non-split spatial graphs admit a canonical decomposition as an iterated vertex sum (Propositions 4.21 and 4.22) of non-split spatial graphs without cut vertices (called ‘blocks’, see Definition 4.18). This reduces the comparison of the isomorphism type of two non-split spatial graphs, to comparing the blocks in their decomposition. Except for one easy special case, these blocks have marked exteriors amenable to the algorithm in the Recognition Theorem.

Iterated vertex sums can be performed along different vertices, so the canonical decomposition must come bundled with the combinatorial data of which vertices from different blocks are glued to which. To package this information, we introduce ‘trees of spatial graphs’ (Definition 4.11), and in case the spatial graphs being glued are blocks, we call it a ‘tree of blocks’ (Definition 4.18). Our main results on decompositions as iterated vertex sums are summarized in the following theorem (see Propositions 4.21 and 4.22 for precise statements):

Theorem 1.2 (Canonical decomposition as a tree of blocks). *Non-split spatial graphs other than a one-point graph admit a unique decomposition as a tree of blocks.*

Though most of our results on iterated vertex sums are perhaps unsurprising, their proofs are often rather involved. In fact, proving that the vertex sum operation is well-defined is one of the most technically demanding points of our program, with most of the work contained in the proof of Proposition 4.2.

We point out that our theory of decompositions has analogues in the setting of abstract graphs [9, Exercise 8.3.3]. In the topological setting, Suzuki [16] has established a unique factorization result with respect to a ‘composition’ operation similar in spirit to our vertex sum, but only for connected 1-subcomplexes of the 3-sphere and up to a ‘neighbourhood congruence’ relation. Our Theorem 1.2 differs from Suzuki’s in the following: our spatial graphs come with vertex/edge decompositions and possibly decorations, we broaden the connectedness assumption to being non-split, and we have no identification ‘up to neighbourhood congruence’, instead keeping track of the vertices along which to glue.

We also deduce a triviality result for spatial forests, that is, spatial graphs Γ whose underlying abstract graph $\langle \Gamma \rangle$ is a forest:

Theorem 1.3 (Spatial forests as their underlying graphs). *If Γ_1, Γ_2 are spatial forests, then every isomorphism $\langle \Gamma_1 \rangle \rightarrow \langle \Gamma_2 \rangle$ of their underlying abstract graphs is induced by an isomorphism of spatial graphs $\Gamma_1 \rightarrow \Gamma_2$. Moreover, the spatial graphs whose isomorphism type is determined by the underlying graph are precisely the spatial forests.*

This result combines Propositions 6.16 and 6.17 in the text. It reduces the isomorphism test for spatial forests to an isomorphism test for abstract forests, bypassing the machinery of Haken–Matveev.

1.4. The piecewise-linear setting

The PL category is the natural home for results in computational topology, such as our main theorem and the machinery in Matveev’s text. It is a standard framework in the field [3, 10, 16, 17], although a theory of smooth, rather than PL, spatial graphs has also been introduced by the first author and Herrmann [4, 5].

We refer to the textbook of Rourke and Sanderson [15] for the standard notions in PL topology. We often give precise references for the results we import, but knowledge of basic concepts such as that of a polyhedron, or a PL manifold (possibly oriented, or with boundary) is assumed. In particular, PL spaces (also called polyhedra) are subspaces of some \mathbb{R}^n whose points admit a star neighbourhood. The ambient space \mathbb{R}^n is equipped with the metric induced from the ℓ^∞ -norm, so by ‘balls’ and ‘spheres’ we mean polyhedra that are PL-homeomorphic to cubes $[-1, 1]^n$ and their boundaries, respectively.

We also make heavy usage of regular neighbourhoods. If $X \subseteq P$ are polyhedra, with X compact, then one may think of a regular neighbourhood of X in P as a ‘small, well-behaved neighbourhood’ of X that deformation-retracts onto X [15, Chapter 3]. If P_0 is a closed sub-polyhedron of P , there is also the notion of a regular neighbourhood N of X in the pair (P, P_0) [15, p. 52]. In this case we use a lighter notation than Rourke–Sanderson, who would instead have written that *the pair* $(N, N \cap P_0)$ is a regular neighbourhood of *the pair* $(X, X \cap P_0)$ in (P, P_0) .

The reader might be more familiar with the combinatorial definition of a PL space as (the topological realization of) an abstract simplicial complex. The two theories are

equivalent: every polyhedron is a union of geometric simplices intersecting along faces [15, Theorem 2.11], and such a decomposition can be abstracted to a combinatorial setup [15, Exercise 2.27(1)]. Moreover, PL maps between polyhedra can be expressed as simplicial maps between subdivisions of the abstract simplicial complexes that they realize [15, Theorem 2.14] (and although the notion of a simplicial subdivision is not combinatorial, the Alexander–Newman Theorem [13, Theorem 4.5] allows one to phrase purely combinatorially the property of two abstract simplicial complexes having PL-homeomorphic realizations).

1.5. Outline of the article

After laying out basic terminology (§2), we introduce the disjoint union of spatial graphs (§3) and prove that the decomposition as a disjoint union of pieces is unique. This program is mirrored in §4, where we define the vertex sum, explain how to specify iterated vertex sums as trees of spatial graphs, and show uniqueness of the decomposition a tree of blocks, thus completing the proof of Theorem 1.2. Section 5 summarizes how to extend the theory developed thus far to spatial graphs decorated with vertex/edge colourings and/or edge orientations.

Section 6 introduces the marked exterior of a (decorated) spatial graph. We define it, explain how it encodes the spatial graph used to construct it, and translate indecomposability properties of spatial graphs into features of their marked exteriors. This is also where we briefly discuss spatial forests and prove Theorem 1.3. Finally, in §7 we import, results from computational 3-manifold topology in order to show that the canonical decompositions of Theorem 1.2 can be computed algorithmically, and apply the Recognition Theorem to marked exteriors of decorated blocks. Altogether, this culminates in the proof of our main result, Theorem 1.1.

2. Basic terminology

We remind the reader that all spaces to be considered are polyhedra: subspaces of \mathbb{R}^n having local cone neighbourhoods at every point, and PL maps are defined as preserving this local cone structure [15, Chapter 1]. Standard models of balls and spheres are defined using the ℓ^∞ -norm, so they are effectively cubes and their boundaries. Orientations of PL manifolds are PL isotopy classes of embeddings of balls [15, pp. 43–46].

Definition 2.1. A *spatial graph* Γ is a triple (\mathcal{S}, V, E) , where:

- \mathcal{S} is an oriented PL 3-sphere, called the **ambient sphere** of Γ . We say that Γ is a spatial graph in \mathcal{S} .
- V is a finite subset of \mathcal{S} , whose elements are called **vertices** of Γ , and
- E is a finite set of subpolyhedra of \mathcal{S} , called **edges** of Γ , such that:
 - each edge is PL-homeomorphic to an interval or to a PL circle,
 - each edge that is PL-homeomorphic to an arc intersects V precisely at its endpoints,
 - each edge that is PL-homeomorphic to a circle contains precisely one element of V (such edges are called **loops**),

– for every two distinct edges e, e' , we have $e \cap e' \subseteq V$.

The **support** of Γ is the union:

$$|\Gamma| := V \cup \bigcup_{e \in E} e \subseteq \mathcal{S}.$$

The **underlying graph** $\langle \Gamma \rangle$ of Γ is the (undirected) abstract graph with vertex set V , edge set E , and where each edge is incident to the one or two elements of V that it contains. We say that an edge of Γ is **incident** to a vertex if this is true in $\langle \Gamma \rangle$. The **degree** of a vertex v is its degree in $\langle \Gamma \rangle$, that is, the number of edges incident to v , with loops counting twice. A vertex of degree 0 is called an **isolated vertex**, and a vertex of degree 1 is called a **leaf**.

A **sub-graph** of $\Gamma = (\mathcal{S}, V, E)$ is a spatial graph $\Gamma' = (\mathcal{S}, V', E')$, where $V' \subseteq V$ and $E' \subseteq E$.

Observe that the two subsets $|\Gamma|$ and V of \mathcal{S} determine E , since there is a canonical bijection between E and $\pi_0(|\Gamma| \setminus V)$.

Definition 2.2. Let $\Gamma_1 = (\mathcal{S}_1, V_1, E_1)$ and $\Gamma_2 = (\mathcal{S}_2, V_2, E_2)$ be spatial graphs. An **isomorphism** $\Phi: \Gamma_1 \rightarrow \Gamma_2$ is a PL homeomorphism of triples $\Phi: (\mathcal{S}_1, |\Gamma_1|, V_1) \rightarrow (\mathcal{S}_2, |\Gamma_2|, V_2)$ respecting the orientation of the ambient spheres. We write $\Gamma_1 \cong \Gamma_2$ if Γ_1, Γ_2 are isomorphic.

By the characterization of E_1 in terms of $|\Gamma_1| \setminus V_1$, and similarly for E_2 , such a Φ also induces a bijection $E_1 \rightarrow E_2$, and we get an isomorphism of abstract graphs $\langle \Phi \rangle: \langle \Gamma_1 \rangle \rightarrow \langle \Gamma_2 \rangle$. We loosen notation by writing ‘ $\Gamma_1 = \Gamma_2$ ’ whenever $\langle \Gamma_1 \rangle = \langle \Gamma_2 \rangle$ and there is an isomorphism $\Phi: \Gamma_1 \rightarrow \Gamma_2$ such that $\langle \Phi \rangle$ is the identity morphism.

Up to isomorphism, there is a unique spatial graph with no vertices (and hence no edges), called the **empty spatial graph**, and denoted by $\mathbf{0}$. Similarly, since the group of PL self-homeomorphisms of a 3-sphere acts transitively on its points [15, Lemma 3.33], there is a unique spatial graph (up to isomorphism) with one vertex and no edges, called the **one-point** spatial graph, denoted by $\mathbf{1}$.

3. The disjoint union of spatial graphs

We will define and establish basic properties of two operations on spatial graphs, which have analogues in the setting of abstract graphs. For the first one, the disjoint union, most arguments are fairly straightforward. However, the overall strategy will serve as a guide for treating the more involved vertex sum operation (§4).

3.1. Assembling spatial graphs through disjoint unions

To define the disjoint union of spatial graphs, we need:

Theorem 3.1 (Disc Theorem [15, Theorem 3.34]). *Every two orientation-preserving PL embeddings of an n -ball into the interior of a connected, oriented n -manifold M are PL-ambient-isotopic relative ∂M .¹*

Definition 3.2. *An enclosing ball for a spatial graph Γ in \mathcal{S} , is a PL-embedded 3-ball $B \subset \mathcal{S}$ such that $|\Gamma| \subset \text{int}(B)$.*

For each $i \in \{1, 2\}$, let $\Gamma_i = (\mathcal{S}_i, V_i, E_i)$ be a spatial graph and B_i an enclosing ball for Γ_i . Suppose $f: \partial B_1 \rightarrow \partial B_2$ is an orientation-reversing PL homeomorphism. Then the spatial graph:

$$\Gamma_1 \sqcup_f \Gamma_2 := (B_1 \cup_f B_2, V_1 \sqcup V_2, E_1 \sqcup E_2),$$

where $B_1 \cup_f B_2$ denotes the 3-sphere obtained by attaching B_1 to B_2 using f , is said to be a disjoint union of Γ_1 and Γ_2 .

The underlying graph $\langle \Gamma_1 \sqcup_f \Gamma_2 \rangle$ is the disjoint union $\langle \Gamma_1 \rangle \sqcup \langle \Gamma_2 \rangle$, as usually defined for abstract graphs.

Lemma 3.3. (Uniqueness of enclosing balls). *Let Γ be a spatial graph in \mathcal{S} , and let B, B' be enclosing balls for Γ . Then every orientation-preserving PL homeomorphism $\Phi_\partial: \partial B \rightarrow \partial B'$ extends to an orientation-preserving PL homeomorphism $\Phi_B: B \rightarrow B'$ that restricts to the identity on $|\Gamma|$.*

Proof. Fix a regular neighbourhood N_Γ of $|\Gamma|$ in \mathcal{S} disjoint from $\partial B \cup \partial B'$. The subspace $\overline{N_\Gamma} := \mathcal{S} \setminus \text{int}(N_\Gamma)$ is also a PL 3-manifold [15, Corollary 3.14]. Moreover, $\overline{B} := \mathcal{S} \setminus \text{int}(B) \subseteq \text{int}(\overline{N_\Gamma})$ is a 3-ball (similarly for $\overline{B'} := \mathcal{S} \setminus \text{int}(B')$) since closures of complements of PL-embedded n -balls in n -spheres are n -balls [15, Corollary 3.13].

Since a PL homeomorphism between the boundaries of two balls extends to a PL homeomorphism of their interiors [15, Lemma 1.10], we may extend Φ_∂ to an orientation-preserving PL homeomorphism $\Phi_{\overline{B}}: \overline{B} \rightarrow \overline{B'}$. Apply Theorem 3.1 to produce a PL ambient isotopy of $\overline{N_\Gamma}$ from the inclusion $\overline{B} \hookrightarrow \overline{N_\Gamma}$ to the composition $\overline{B} \xrightarrow{\Phi_{\overline{B}}} \overline{B'} \hookrightarrow \overline{N_\Gamma}$. As this ambient isotopy keeps ∂N_Γ fixed, the homeomorphism $\Phi_{\overline{N_\Gamma}}: \overline{N_\Gamma} \rightarrow \overline{N_\Gamma}$ extends to \mathcal{S} by setting it to the identity on N_Γ . This extension $\Phi_{\mathcal{S}}: \mathcal{S} \rightarrow \mathcal{S}$, when restricted to B , is a PL homeomorphism $\Phi_B: B \rightarrow B'$ satisfying the conclusion of the lemma. \square

Proposition 3.4. (Disjoint union is well-defined). *For each $i \in \{1, 2\}$, let Γ_i be a spatial graph in \mathcal{S}_i with two enclosing balls B_i, B'_i , and suppose $f: \partial B_1 \rightarrow \partial B_2, f': \partial B'_1 \rightarrow \partial B'_2$ are orientation-reversing PL homeomorphisms. Then $\Gamma_1 \sqcup_f \Gamma_2 = \Gamma_1 \sqcup_{f'} \Gamma_2$.*

Recall from § 2 that the equality in this proposition means that there is an isomorphism $\Phi: \Gamma_1 \sqcup_f \Gamma_2 \rightarrow \Gamma_1 \sqcup_{f'} \Gamma_2$ such that $\langle \Phi \rangle$ is the identity on $\langle \Gamma_1 \rangle \sqcup \langle \Gamma_2 \rangle$.

Proof. By Lemma 3.3 there is an orientation-preserving PL homeomorphism $\Phi_1: B_1 \rightarrow B'_1$ with $\Phi_1|_{|\Gamma_1|} = \text{id}$. Similarly, let $\Phi_2: B_2 \rightarrow B'_2$ be an orientation-preserving

¹ The reference does not state that the ambient isotopy fixes ∂M , but this follows from the proof. Later, a stronger version of the Disc Theorem (Theorem 4.4) will include the boundary condition.

PL homeomorphism fixing $|\Gamma_2|$ and whose restriction to ∂B_2 is $f' \circ \Phi_1|_{\partial B_1} \circ f^{-1}$. The maps Φ_i assemble to a PL homeomorphism $\Phi: B_1 \sqcup_f B_2 \rightarrow B'_1 \sqcup_{f'} B'_2$ giving the required isomorphism between $\Gamma_1 \sqcup_f \Gamma_2$ and $\Gamma_1 \sqcup_{f'} \Gamma_2$. □

From now on we suppress f from the notation $\Gamma_1 \sqcup_f \Gamma_2$ when no confusion arises.

Lemma 3.5. (Disjoint union summands as sub-graphs). *Let $\Gamma = \Gamma_1 \sqcup \Gamma_2$ be a disjoint union of spatial graphs, and denote by Γ'_1 the sub-graph of Γ obtained by discarding all vertices and edges of Γ_2 . Then $\Gamma'_1 = \Gamma_1$.*

Proof. Let $B_i \subset \mathcal{S}_i$ be the enclosing balls from which the disjoint union was formed, with attaching map $f: \partial B_1 \rightarrow \partial B_2$. To construct a PL homeomorphism $\Phi: \mathcal{S}_1 \rightarrow B_1 \cup_f B_2$ with $\Phi|_{|\Gamma_1|} = \text{id}$, take Φ as the identity on B_1 and on $\overline{B_1} := \mathcal{S}_1 \setminus \text{int}(B_1)$ choose any extension $\overline{B_1} \rightarrow B_2$ of f [15, Lemma 1.10]. □

Lemma 3.6. (Disjoint union of isomorphisms). *Let $\Phi_1: \Gamma_1 \rightarrow \Gamma'_1$ and $\Phi_2: \Gamma_2 \rightarrow \Gamma'_2$ be isomorphisms of spatial graphs. Then there exists an isomorphism:*

$$\Phi_1 \sqcup \Phi_2: \Gamma_1 \sqcup \Gamma_2 \rightarrow \Gamma'_1 \sqcup \Gamma'_2,$$

such that for each $i \in \{1, 2\}$ the underlying isomorphism of abstract graphs $\langle \Phi_1 \sqcup \Phi_2 \rangle$ restricts to $\langle \Phi_i \rangle$ on $\langle \Gamma_i \rangle$.

Proof. Form the disjoint union $\Gamma_1 \sqcup_f \Gamma_2$ by using a suitable PL homeomorphism $f: \partial B_1 \rightarrow \partial B_2$ between the boundaries of enclosing balls for Γ_1, Γ_2 . Writing $B'_i := \Phi_i(B_i)$ and defining $f': \partial B'_1 \rightarrow \partial B'_2$ as $f' := \Phi_2|_{\partial B_2} \circ f \circ \Phi_1^{-1}|_{\partial B_1}$, we can form the disjoint union $\Gamma'_1 \sqcup_{f'} \Gamma'_2$. The restrictions $\Phi_i|_{B_i}$ then assemble to the desired isomorphism $\Phi_1 \sqcup \Phi_2$. □

Note that Lemma 3.6 strongly depends on the ambient 3-spheres carrying an orientation, which is preserved by isomorphisms. Without this requirement, a spatial graph Γ comprised of one vertex and one loop in the shape of a trefoil would be isomorphic to its mirror-image $\tilde{\Gamma}$, while $\Gamma \sqcup \Gamma \not\cong \Gamma \sqcup \tilde{\Gamma}$.

Proposition 3.7. (Properties of the disjoint union). *Let $\Gamma_1, \Gamma_2, \Gamma_3$ be spatial graphs. Then the following conditions hold:*

- **0 is the identity element:** $\Gamma_1 \sqcup \mathbf{0} = \Gamma_1$,
- **commutativity:** $\Gamma_1 \sqcup \Gamma_2 = \Gamma_2 \sqcup \Gamma_1$,
- **associativity:** $(\Gamma_1 \sqcup \Gamma_2) \sqcup \Gamma_3 = \Gamma_1 \sqcup (\Gamma_2 \sqcup \Gamma_3)$.

Proof. Lemma 3.5 implies the first claim, and commutativity follows by using the same enclosing balls for both disjoint unions, and mutually inverse attaching maps.

Associativity is illustrated in Figure 3.1. Let \mathcal{S}_i be the ambient sphere of Γ_i and B_1, B_3 enclosing balls for Γ_1, Γ_3 , respectively, and let B_{21}, B_{23} be enclosing balls for Γ_2 such that $\text{int}(B_{21}) \cup \text{int}(B_{23}) = \mathcal{S}_2$ (that is, $\mathcal{S}_2 \setminus \text{int}(B_{21}) \cap \mathcal{S}_2 \setminus \text{int}(B_{23}) = \emptyset$). Choosing attaching maps $f_1: \partial B_1 \rightarrow \partial B_{21}, f_3: \partial B_{23} \rightarrow \partial B_3$, it follows that $B_1 \cup_{f_1} (B_{21} \cap B_{23})$ is an enclosing ball for $\Gamma_1 \sqcup_{f_1} \Gamma_2$, and $(B_{21} \cap B_{23}) \cup_{f_3} B_3$ is an enclosing ball for $\Gamma_2 \sqcup_{f_3} \Gamma_3$. The spatial graphs $(\Gamma_1 \sqcup_{f_1} \Gamma_2) \sqcup_{f_3} \Gamma_3$ and $\Gamma_1 \sqcup_{f_1} (\Gamma_2 \sqcup_{f_3} \Gamma_3)$ are thus the same. □




Figure 3.1. The proof of associativity of the disjoint union.

We can thus unambiguously write down iterated disjoint unions. More precisely, if $\{\Gamma_i\}_{i \in I}$ is a collection of spatial graphs with I finite, then $\bigsqcup_{i \in I} \Gamma_i$ is well-defined up to isomorphism inducing the identity on $\bigsqcup_{i \in I} \langle \Gamma_i \rangle$.

3.2. Decomposing spatial graphs as disjoint unions

We want to show that spatial graphs canonically decompose into iterated disjoint unions. We (often implicitly) use the fact that every PL-embedded 2-sphere in a PL 3-sphere decomposes it into two 3-balls. (Recall that this holds in the PL category but not in the topological setting, by work of Alexander [1, 2].)

Lemma 3.8. (If it looks like a disjoint union, it is a disjoint union). *Let Γ be a spatial graph in \mathcal{S} and $S \subset \mathcal{S} \setminus |\Gamma|$ a PL-embedded 2-sphere. Denote the closures of the two components of $\mathcal{S} \setminus S$ by B_1 and B_2 . For each $i \in \{1, 2\}$, let Γ_i be the sub-graph of Γ comprised of the vertices and edges contained in B_i . Then $\Gamma = \Gamma_1 \sqcup \Gamma_2$.*

Proof. Use Lemma 3.5 to regard each Γ_i as a sub-graph of $\Gamma_1 \sqcup \Gamma_2$, and take B_i as an enclosing ball for Γ_i . If $f: S \rightarrow S$ is the identity map, then $\Gamma = \Gamma_1 \sqcup_f \Gamma_2$. \square

Definition 3.9. *Let Γ be a spatial graph in \mathcal{S} .*

- *If $S \subset \mathcal{S}$ is a 2-sphere as in Lemma 3.8, we say that ‘ S decomposes Γ as $\Gamma_1 \sqcup \Gamma_2$ ’.*
- *Γ is **split** if it is the disjoint union of two non-empty spatial graphs; otherwise it is **non-split**.*
- *If S is a 2-sphere in \mathcal{S} decomposing Γ as $\Gamma_1 \sqcup \Gamma_2$ with Γ_1, Γ_2 non-empty, then S is a **splitting sphere** for Γ .*



Figure 3.2. A non-split spatial graph with disconnected support.

- A spatial graph is a **piece** if it is non-empty and non-split. We also say that a spatial graph Λ is ‘a piece of Γ ’ if Λ is a piece and $\Gamma = \Gamma' \sqcup \Lambda$ for some Γ' .

We use the word ‘piece’ rather than ‘component’ to avoid suggesting that $|\Lambda|$ (or equivalently the abstract graph $\langle \Lambda \rangle$) is connected; see Figure 3.2.

By induction, every spatial graph can be expressed as a disjoint union of finitely many pieces. In the sequel we establish uniqueness of such a decomposition.

Lemma 3.10. (Spheres sort pieces). *Let Λ be a piece in \mathcal{S} , and let $S \subset \mathcal{S} \setminus |\Lambda|$ be a PL-embedded 2-sphere. Denote the closures of the two components of $\mathcal{S} \setminus S$ by B_1, B_2 . Then, $|\Lambda|$ is contained in exactly one of the B_i .*

Proof. Since $\Lambda \neq \mathbf{0}$, certainly $|\Lambda|$ cannot be contained in both B_i . Denote by Λ_i the sub-graph of Λ whose vertices and edges are contained in B_i . By Lemma 3.8, we see Λ decomposes as $\Lambda_1 \sqcup \Lambda_2$. Since Λ is non-split, one of the summands, say Λ_1 , is empty. By the first part of Proposition 3.7, it follows that $\Lambda_2 = \Lambda$. □

Proposition 3.11. (Uniqueness of decomposition into pieces). *Let $(\Lambda_i)_{i \in I_1}$ and $(\Lambda_i)_{i \in I_2}$ be collections of pieces with I_1, I_2 finite. If $\Phi: \bigsqcup_{i \in I_1} \Lambda_i \rightarrow \bigsqcup_{i \in I_2} \Lambda_i$ is an isomorphism, there is a bijection $f: I_1 \rightarrow I_2$ such that for each $i \in I_1$, the PL homeomorphism Φ is an isomorphism of the sub-graphs $\Phi: \Lambda_i \rightarrow \Lambda_{f(i)}$.*

Proof. Write $\Gamma_1 := \bigsqcup_{i \in I_1} \Lambda_i$ and $\Gamma_2 := \bigsqcup_{i \in I_2} \Lambda_i$. We induct on $\#I_1$.

If $I_1 = \emptyset$ then $\Gamma_1 = \mathbf{0} = \Gamma_2$, whence $I_2 = \emptyset$ and there is nothing left to show. If $I_1 = \{i_1\}$, then $\Gamma_1 = \Lambda_{i_1}$ is a piece. Hence Γ_2 is also a piece and therefore $I_2 = \{i_2\}$. We thus set $f(i_1) := i_2$.

If I_1 has more than one element, choose any partition into non-empty subsets $I_1 = I_1^+ \sqcup I_1^-$. Let S_1 be a 2-sphere decomposing Γ_1 as $(\bigsqcup_{i \in I_1^+} \Lambda_i) \sqcup (\bigsqcup_{i \in I_1^-} \Lambda_i)$, and write $\Gamma_1^+ := \bigsqcup_{i \in I_1^+} \Lambda_i$ and $\Gamma_1^- := \bigsqcup_{i \in I_1^-} \Lambda_i$. Now $S_2 := \Phi(S_1)$ is a 2-sphere in the ambient sphere of Γ_2 disjoint from $|\Gamma_2|$ whose sides correspond to Γ_1^+ and Γ_1^- . By Lemma 3.10, each $|\Lambda_i|$ is contained in either the ‘+’-side or the ‘-’-side of S_2 . Partition I_2 accordingly as $I_2 = I_2^+ \sqcup I_2^-$, and write $\Gamma_2^\pm := \bigsqcup_{i \in I_2^\pm} \Lambda_i$. Since Φ maps the support $|\Gamma_1^\pm|$ into $|\Gamma_2^\pm|$, we conclude that Φ doubles as a pair of isomorphisms of sub-graphs $\Phi^\pm: \Gamma_1^\pm \rightarrow \Gamma_2^\pm$. Both I_1^\pm have fewer elements than I_1 , so by induction we obtain bijections $f^\pm: I_1^\pm \rightarrow I_2^\pm$, which assemble to the required $f: I_1 \rightarrow I_2$. □

By Lemma 3.6, the converse of Proposition 3.11 also holds. Hence, if one decomposes Γ_1, Γ_2 as disjoint unions of pieces, then testing if $\Gamma_1 \cong \Gamma_2$ boils down to testing whether

the pieces are pairwise isomorphic. To do so, we first need to further decompose pieces, which is the content of the next section.

4. The vertex sum of spatial graphs

The next operation combines pairs of spatial graphs with distinguished vertices. Many definitions and results have analogues in §3.

4.1. Defining the vertex sum

A **pointed spatial graph** is a pair (Γ, v) consisting of a spatial graph Γ and a vertex v of Γ . The underlying graph of a pointed spatial graph is pointed with the same distinguished vertex. An isomorphism of pointed spatial graphs is an isomorphism of the spatial graphs preserving distinguished vertices.

Definition 4.1. An **enclosing ball** for a pointed spatial graph (Γ, v) in \mathcal{S} is a PL-embedded 3-ball $B \subset \mathcal{S}$ such that $|\Gamma| \subset B$ and $|\Gamma| \cap \partial B = \{v\}$.

For each $i \in \{1, 2\}$, let $\Gamma_i = (\mathcal{S}_i, V_i, E_i)$ be a non-empty spatial graph, let $v_i \in V_i$, and let B_i be an enclosing ball for (Γ_i, v_i) . Moreover, let $f: \partial B_1 \rightarrow \partial B_2$ be an orientation-reversing PL homeomorphism mapping v_1 to v_2 . We consider the spatial graph:

$$\Gamma_1 \bullet_{v_1 \sim v_2} \Gamma_2 := (B_1 \cup_f B_2, (V_1 \sqcup V_2)/v_1 \sim v_2, E_1 \sqcup E_2),$$

where $B_1 \cup_f B_2$ denotes the 3-sphere obtained by attaching B_1 to B_2 using f , and define the pointed spatial graph $(\Gamma_1 \bullet_{v_1 \sim v_2} \Gamma_2, v_1 = v_2)$ to be a **vertex sum** of (Γ_1, v_1) and (Γ_2, v_2) .

We use the same notation for the analogous operation on pointed abstract graphs, so $\langle \Gamma_1 \bullet_{v_1 \sim v_2} \Gamma_2 \rangle = \langle \Gamma_1 \rangle_{v_1 \sim v_2} \langle \Gamma_2 \rangle$.

The following result is key in showing that the vertex sum is well-defined.

Proposition 4.2. (Uniqueness of enclosing balls for pointed spatial graphs).

Let (Γ, v) be a pointed spatial graph in \mathcal{S} , let B, B' be enclosing balls for (Γ, v) . Then every orientation-preserving PL homeomorphism $\Phi_\partial: (\partial B, v) \rightarrow (\partial B', v)$ extends to an orientation-preserving PL homeomorphism $\Phi_B: B \rightarrow B'$ restricting to the identity on $|\Gamma|$.

Proving this proposition requires substantially more work than its non-pointed counterpart, Lemma 3.3, due to the particular behaviour demanded of Φ near v . One of the ingredients is a generalization of the Disc Theorem.

Definition 4.3. ([15, pp. 50, 51]). An **unknotted ball pair** (B, B_0) is a pair of polyhedra PL-homeomorphic to a standard ball pair $([-1, 1]^n, [-1, 1]^m \times \{0\}^{n-m})$ (for some $n \geq m \geq 0$). A **PL manifold pair** (M, M_0) is a pair of polyhedra that are manifolds, such that $\partial M \cap M_0 = \partial M_0$ ('properness'), and such that each point of M_0 has a neighbourhood in (M, M_0) PL-homeomorphic to an unknotted ball pair ('local flatness').²

² The definition given by Rourke-Sanderson on p. 50 requires only that M, M_0 both be manifolds, but the remark on p. 51 adds the local flatness and properness conditions.

Theorem 4.4 (Disc Theorem for pairs [15, Theorem 4.20]). *Let (M, M_0) be a pair of connected, oriented PL manifolds, let (B, B_0) be an unknotted ball pair with the same dimensions as (M, M_0) , and let $\iota_1, \iota_2: (B, B_0) \rightarrow (\text{int}(M), \text{int}(M_0))$ be PL embeddings that preserve the orientation on both components. Then there is a PL ambient isotopy of (M, M_0) relative ∂M that carries ι_1 to ι_2 .*

Corollary 4.5. (Disc Theorem at the boundary). *Let M be a connected, oriented PL n -manifold, let $N \subseteq \partial M$ be a connected PL-embedded $(n - 1)$ -manifold that is closed in ∂M , let B be a PL n -ball, and $D \subset \partial B$ a PL $(n - 1)$ -ball. For every two orientation-preserving PL embeddings $\iota_1, \iota_2: (B, D) \rightarrow (\text{int}(M) \cup \text{int}(N), \text{int}(N))$, there is a PL ambient isotopy of (M, N) relative $\partial M \setminus \text{int}(N)$ carrying ι_1 to ι_2 .*

Proof. Consider the double $\mathcal{D}_N(M)$ of M along N , which is a union of two copies of M glued along the identity map on N , one of the copies with its orientation reversed. Using the fact that N is closed in ∂M one sees that $(\mathcal{D}_N(M), N)$ is a PL manifold pair, and its boundary is $(\mathcal{D}_{\partial N}(\partial M \setminus \text{int}(N)), \partial N)$. Doubling also B along D yields an unknotted ball pair $(\mathcal{D}_D(B), D)$.

Now, the maps ι_1, ι_2 extend to orientation-preserving PL embeddings:

$$\mathcal{D}(\iota_1), \mathcal{D}(\iota_2): (\mathcal{D}_D(B), D) \rightarrow (\text{int}(\mathcal{D}_N(M)), \text{int}(N)).$$

Theorem 4.4 yields a PL ambient isotopy of $(\mathcal{D}_N(M), N)$ relative $\mathcal{D}_{\partial N}(\partial M \setminus \text{int}(N))$ that carries $\mathcal{D}(\iota_1)$ to $\mathcal{D}(\iota_2)$. A connectivity argument shows that it restricts to an isotopy from ι_1 to ι_2 relative $\partial M \setminus \text{int}(N)$. □

We also need Lemma 4.6 below, but first recall some terminology. Given a polyhedron $P \subseteq \mathbb{R}^n$ and $v \in \mathbb{R}^n$, we denote by vP the polyhedron comprised of all points of the form $tp + (1 - t)v$, with $p \in P$ and $t \in [0, 1]$. If each point of vP admits a unique such expression, we say vP is a **cone** with base P and vertex v . If vP, wQ are cones with bases P, Q and vertices v, w , respectively, the **cone** of a PL map $f: P \rightarrow Q$ (with respect to v, w) is the PL map $vP \rightarrow wQ$ given by $tp + (1 - t)v \mapsto tf(p) + (1 - t)w$ [15, Exercise 1.6(3)].

Lemma 4.6. (Interpolating annulus). *Let A_0 be a PL annulus in \mathbb{R}^n , and let $v \in \mathbb{R}^n$ be such that vA_0 is a cone with base A_0 and vertex v . Denote the boundary circles of A_0 by γ_0, δ_0 , and let $\gamma \subset v\gamma_0$ and $\delta \subset v\delta_0$ be PL circles such that $v\gamma, v\delta$ are cones with bases γ, δ respectively, and vertex v . Then there is a PL annulus $A \subset vA_0$ with $\partial A = \gamma \cup \delta$, such that vA is a cone with base A and vertex v .*

This statement is illustrated in Figure 4.1.

Proof. We may assume that $A_0 = C \times [0, 1] \subset \mathbb{R}^n$ for some PL circle $C \subset \mathbb{R}^{n-1}$, with $\gamma_0 = C \times \{0\}$ and $\delta_0 = C \times \{1\}$, because the cone $v(C \times [0, 1]) \rightarrow vA_0$ preserves cones at v for every PL homeomorphism $C \times [0, 1] \rightarrow A_0$.

Choose a finite set of points in $\gamma \subset v(C \times \{0\})$ subdividing γ into straight line segments [15, Theorem 2.2]. Pushing these points radially into $\gamma_0 = C \times \{0\}$ and projecting onto C yields a finite set of points in C (note that since $v\gamma$ is a cone, no two points of γ get pushed to the same point of γ_0). Doing the same with δ yields a second finite subset




Figure 4.1. A PL annulus A ‘interpolating’ between γ and δ .

of C . Finally, choose a third finite subset of C subdividing C itself into straight line segments. We denote by p_1, \dots, p_k the points in the union of these three subsets, ordered cyclically around C (with indices $1, \dots, k$ modulo k). Now push $(p_1, 0), \dots, (p_k, 0) \in \gamma_0$ radially into γ to obtain points $p_1^\gamma, \dots, p_k^\gamma$. Similarly, pushing $(p_1, 1), \dots, (p_k, 1)$ radially yields $p_1^\delta, \dots, p_k^\delta$.

Since the p_i subdivide C into straight line segments, we see that for each $i \in \mathbb{Z}/k$, the points $(p_i, 0), (p_{i+1}, 0), (p_i, 1), (p_{i+1}, 1)$ are the vertices of a rectangle R_i contained in A_0 . In particular, the cone $vR_i \subset vA_0$ is convex.

For each $i \in \mathbb{Z}/k$, denote by T_i^γ the triangle spanned by the points $p_i^\gamma, p_{i+1}^\gamma, p_i^\delta$, and by T_i^δ the one spanned by $p_i^\delta, p_{i+1}^\delta, p_{i+1}^\gamma$. By the previous observation, these triangles are contained in vR_i . The union $A := \bigcup_{i \in \mathbb{Z}/k} (T_i^\gamma \cup T_i^\delta)$ is then a PL annulus embedded in vA_0 , with $\partial A = \gamma \cup \delta$. It is also clear that each point of A lies in a unique ray from v through a point in A . The cone condition on vA follows. \square

Proof of Proposition 4.2. Write $\overline{B} := \mathcal{S} \setminus \text{int}(B)$ and $\overline{B}' := \mathcal{S} \setminus \text{int}(B')$, and choose any extension of Φ_∂ to a PL homeomorphism $\Phi_{\overline{B}}: \overline{B} \rightarrow \overline{B}'$. We will find an extension $\Phi_{\mathcal{S}}: \mathcal{S} \rightarrow \mathcal{S}$ of $\Phi_{\overline{B}}$ that fixes $|\Gamma|$, and whose restriction Φ_B to B therefore satisfies the conclusion of the lemma. The intricate construction of $\Phi_{\mathcal{S}}$, for which we need some notation, is illustrated in Figure 4.2.

First, choose a star neighbourhood N_0 for v in the pair $(\mathcal{S}, \overline{B}' \cup |\Gamma|)$. More explicitly, N_0 is a 3-ball such that $(\overline{B}' \cup |\Gamma|) \cap N_0$ is a cone with base its intersection with ∂N_0 , and vertex v [15, p. 50]. In particular, $D_0 := \overline{B}' \cap \partial N_0$ is a disc and $\overline{B}' \cap N_0$ is the cone vD_0 with base D_0 and vertex v .

We then pick a smaller star neighbourhood $N_v \subset \text{int}(N_0)$ of v in $(\mathcal{S}, \overline{B}' \cup |\Gamma|)$, such that N_v is also a star neighbourhood of v in $(\mathcal{S}, \overline{B} \cup |\Gamma|)$, and $\overline{B} \cap N_v$ is mapped conically by $\Phi_{\overline{B}}$ into $\text{int}(N_0)$. Denoting by D the disc $\overline{B} \cap \partial N_v$, so $\overline{B} \cap N_v$ is a cone vD with base D and vertex v , this means that $\Phi_{\overline{B}}(vD)$ is a cone vD' with base the disc $D' := \Phi_{\overline{B}}(D)$ and vertex v , and that $\Phi_{\overline{B}}|_{vD}: vD \rightarrow vD'$ is the cone of $\Phi_{\overline{B}}|_D: D \rightarrow D'$. The existence




Figure 4.2. The 3-balls \bar{B} and \bar{B}' .




Figure 4.3. The 3-ball \bar{B}' and its Ψ -image \tilde{B} .

of such N_v follows from the definitions of PL map and polyhedron, say, by taking N_v to be a sufficiently small ϵ -neighbourhood of v . We denote by \bar{N}_v the 3-ball $\mathcal{S} \setminus \text{int}(N_v)$.

To apply the disc theorem at the boundary, we need to move \bar{B}' to the more convenient configuration given by the following claim and illustrated in Figure 4.3.

Claim. There exists an orientation-preserving PL homeomorphism $\Psi : \mathcal{S} \rightarrow \mathcal{S}$ fixing $|\Gamma|$ and such that:

- Ψ maps the pair $(\Phi_{\bar{B}}(\bar{B} \cap \bar{N}_v), D')$ into the pair $(\bar{N}_v, \partial N_v)$, and
- writing $\tilde{D} := \Psi(D')$, the map Ψ is given on vD' as the cone $vD' \rightarrow v\tilde{D}$ of the PL homeomorphism $D' \rightarrow \tilde{D}$.

Assuming the claim for the moment, let us see how to use the resulting Ψ to construct the desired extension $\Phi_{\mathcal{S}}$ of $\Phi_{\bar{B}}$.

Let \tilde{B} be the 3-ball $\Psi(\bar{B}')$ and choose a regular neighbourhood N_{Γ} of $|\Gamma| \cap \bar{N}_v$ in \bar{N}_v , small enough to be disjoint from \bar{B} and \tilde{B} . Denote by M the closure of $\bar{N}_v \setminus N_{\Gamma}$ in \bar{N}_v , and consider the closed codimension-0 submanifold $N := \partial N_v \cap M$ of ∂M . By construction of Ψ , its restriction to $\Phi_{\bar{B}}(\bar{B} \cap \bar{N}_v)$ is a PL homeomorphism of pairs $(\Phi_{\bar{B}}(\bar{B} \cap \bar{N}_v), D') \rightarrow (\tilde{B} \cap \bar{N}_v, \tilde{D})$. We may thus apply Corollary 4.5 to the inclusion $(\bar{B} \cap \bar{N}_v, D) \hookrightarrow (M, N)$




Figure 4.4. Applying the Disc Theorem at the boundary to ambiently isotope $(\overline{B} \cap \overline{N}_v, D)$ onto $(\tilde{B} \cap \overline{N}_v, \tilde{D})$ within (M, N) .

and the composition:

$$(\overline{B} \cap \overline{N}_v, D) \xrightarrow{\Phi_{\overline{B}}} (\Phi_{\overline{B}}(\overline{B} \cap \overline{N}_v), D') \xrightarrow{\Psi} (\tilde{B} \cap \overline{N}_v, \tilde{D}) \hookrightarrow (M, N).$$

This is illustrated in Figure 4.4.

The final PL homeomorphism $\tilde{\Phi}_M : M \rightarrow M$ of the resulting PL isotopy of M extends the composition $\Psi|_{\Phi_{\overline{B}}(\overline{B} \cap \overline{N}_v)} \circ \Phi_{\overline{B}}|_{\overline{B} \cap \overline{N}_v}$ and fixes $\partial M \setminus \text{int}(N) = \partial M \cap N_\Gamma$. We may thus extend $\tilde{\Phi}_M$ to a PL homeomorphism $\tilde{\Phi}_{N_v} : \overline{N}_v \rightarrow \overline{N}_v$ by setting it to be the identity on N_Γ . In particular, $\tilde{\Phi}_{N_v}$ fixes $|\Gamma| \cap \overline{N}_v$. Finally, extend $\tilde{\Phi}_{N_v}$ to a PL homeomorphism $\tilde{\Phi}_S : S \rightarrow S$ by defining it on $N_v = v(\partial N_v)$ as the cone of the already prescribed PL homeomorphism $\partial N_v \rightarrow \partial N_v$.

The restriction $\tilde{\Phi}_S|_{\overline{B}}$ is now the composition $\Psi|_{\overline{B}} \circ \Phi_{\overline{B}}$: indeed, we have seen that the two maps agree on $\overline{B} \cap \overline{N}_v$, and on $\overline{B} \cap N_v = vD$ both are defined as the cone of $D \xrightarrow{\Phi_{\overline{B}}} D' \xrightarrow{\Psi} \tilde{D}$. Moreover, $\tilde{\Phi}_S$ fixes $|\Gamma|$. Hence, the map $\Phi_S := \Psi^{-1} \circ \tilde{\Phi}_S$ extends $\Phi_{\overline{B}}$ and fixes $|\Gamma|$, as desired. \square

Proof of the Claim. The argument is illustrated in Figure 4.5. Choose a collar for ∂D_0 in $B' \cap \partial N_0$, that is, a PL embedding $c : \partial D_0 \times [0, 1] \rightarrow B' \cap \partial N_0$ such that $c(-, 0)$ is the identity on ∂D_0 , and $c(\partial D_0 \times [0, 1])$ is an open neighbourhood of ∂D_0 in $B' \cap \partial N_0$. We may also assume that the image A_0 of c is disjoint from $|\Gamma|$; see [15, p. 24] for more on collars.

Let D_0^+ be the ‘enlarged disc’ $D_0 \cup A_0$, and consider the β -ball vD_0^+ , which is a cone with base D_0^+ and vertex v . We shall define Ψ as the identity on $S \setminus \text{int}(vD_0^+)$ and then suitably extend the identity on $\partial(vD_0^+)$ to all of vD_0^+ .

Denote by \tilde{D}^+ the disc $vD_0^+ \cap \partial N_v$ and consider the PL circles $\partial D'$ and $\partial \tilde{D}^+$, each lying in the cone of a distinct component of ∂A_0 . Each of these circles is the base of a cone with vertex v , so we can use Lemma 4.6 to find an annulus A with $\partial A = \partial D' \cup \partial \tilde{D}^+$, such that vA is a cone with base A and vertex v . Denote by D'^+ the disc $D' \cup A$ and note that by construction, $\partial D'^+ = \partial \tilde{D}^+$.




Figure 4.5. The construction of Ψ .

To define Ψ inside vD_0^+ , we first choose any extension of the identity map $\partial D'^+ \rightarrow \partial \tilde{D}^+$ to a PL homeomorphism $D'^+ \rightarrow \tilde{D}^+$. Since both vD'^+ and $v\tilde{D}^+$ are cones at v , we can define Ψ on vD'^+ as the cone of the above PL homeomorphism $D'^+ \rightarrow \tilde{D}^+$. This is consistent with the definition of Ψ as the identity on $\partial(vD_0^+)$.

It remains only to define Ψ on $vD_0^+ \setminus vD'^+$, whose closure in \mathcal{S} is a 3-ball C (because it is the complement in vD_0^+ of an open regular neighbourhood of a boundary point). Denoting by \tilde{C} the closure in \mathcal{S} of $vD_0^+ \setminus v\tilde{D}^+$, which is a 3-ball, this amounts to choosing a PL homeomorphism $C \rightarrow \tilde{C}$ extending the prescribed map $\partial C \rightarrow \partial \tilde{C}$. We choose any extension, and this completes the construction of Ψ . It is straightforward to verify that all required conditions on Ψ are satisfied. □

With the claim established, the proof of Proposition 4.2 is complete.

Proposition 4.7 (Vertex sum is well-defined). *Any two vertex sums of pointed spatial graphs $(\Gamma_1, v_1), (\Gamma_2, v_2)$ are isomorphic via an isomorphism that induces the identity on $(\langle \Gamma_1 \rangle_{v_1 \bullet v_2}, v_1 = v_2)$.*

Proof. The argument can be copied almost word-by-word from the proof of Proposition 3.4, with the role of Lemma 3.3 now played by Proposition 4.2. □

The ambiguity about enclosing balls and attaching maps in the notation $\langle \Gamma_1 \rangle_{v_1 \bullet v_2} \Gamma_2$ is thus immaterial.

We remark that, for abstract graphs, we can define the vertex sum along an ordered k -tuple of distinct vertices. For spatial graphs, however, we would need to define an enclosing ball as a 3-ball containing the support of the spatial graph, and whose boundary intersects it precisely at the k distinguished vertices. But such balls might be non-unique in the sense of Proposition 4.2; see Figure 4.6.

Lemmas 3.5 and 3.6 have analogues for vertex sums, with the same proofs:

Lemma 4.8 (Vertex summands as sub-graphs). *Let $\Gamma = \Gamma_1 \bullet_{v_1 v_2} \Gamma_2$ be a vertex sum of pointed spatial graphs, and denote by Γ'_1 the sub-graph of Γ obtained by discarding all vertices and edges that are not in Γ_1 . Then $\Gamma'_1 = \Gamma_1$.*




Figure 4.6. An example of non-uniqueness of enclosing balls for a spatial graph with two distinguished vertices.

We have slightly extended our ongoing abuse of notation when writing ‘ $\Gamma'_1 = \Gamma_1$ ’. Implicit in this statement is an equality between v_1 and the vertex of $\Gamma_1 \bullet_{v_1 \bullet v_2} \Gamma_2$ obtained from the identification $v_1 \sim v_2$.

Lemma 4.9 (Vertex sum of isomorphisms). *Let $\Phi_1: (\Gamma_1, v_1) \rightarrow (\Gamma'_1, v'_1)$ and $\Phi_2: (\Gamma_2, v_2) \rightarrow (\Gamma'_2, v'_2)$ be isomorphisms of pointed spatial graphs. Then there exists an isomorphism:*

$$\Phi_1 \bullet_{v_1 \bullet v_2} \Phi_2: (\Gamma_1 \bullet_{v_1 \bullet v_2} \Gamma_2, v_1 = v_2) \rightarrow (\Gamma'_1 \bullet_{v'_1 \bullet v'_2} \Gamma'_2, v'_1 = v'_2),$$

such that for each $i \in \{1, 2\}$ the underlying isomorphism $\langle \Phi_1 \bullet_{v_1 \bullet v_2} \Phi_2 \rangle$ restricts to $\langle \Phi_i \rangle$ on $\langle \Gamma_i \rangle$.

Proposition 4.10. (Properties of the vertex sum). *Let $(\Gamma_1, v_1), (\Gamma_2, v_2), (\Gamma_3, v_3)$ be pointed spatial graphs. Then the following conditions hold:*

- **1** is the identity element: $(\Gamma_1 \bullet_{v_1 \bullet v_1} \mathbf{1}, v_1) = (\Gamma_1, v_1)$,
- commutativity: $(\Gamma_1 \bullet_{v_1 \bullet v_2} \Gamma_2, v_1 = v_2) = (\Gamma_2 \bullet_{v_2 \bullet v_1} \Gamma_1, v_1 = v_2)$,
- associativity: for v_{21} and v_{23} (not necessarily distinct) vertices of Γ_2 , we have

$$(\Gamma_1 \bullet_{v_1 \bullet v_{21}} \Gamma_2) \bullet_{v_{23} \bullet v_3} \Gamma_3 = \Gamma_1 \bullet_{v_1 \bullet v_{21}} (\Gamma_2 \bullet_{v_{23} \bullet v_3} \Gamma_3).$$

Proof. Identical to the proof of Proposition 3.7 save for the following modifications:

- The first item relies on vertex summands being sub-graphs (Lemma 4.8), rather than disjoint union summands being sub-graphs (Lemma 3.5).
- To prove associativity in the case $v_{21} = v_{23}$, the requirement on enclosing balls is that $\text{int}(B_1) \cup \text{int}(B_2) = \mathcal{S} \setminus \{v_{21}\}$ (equivalently, $\overline{B_{21}} \cap \overline{B_{23}} = \{v_{21}\}$). □

4.2. Iterated vertex sums and trees of spatial graphs

To elucidate how iterated vertex sums are constructed without keeping track of the order in which they are performed, we need to encode the involved combinatorics in the notation.

First, consider the case where only one vertex of each summand is used. We denote by $(\star_{i \in I} (\Gamma_i, v_i), v)$ the vertex sum of a collection of pointed spatial graphs (Γ_i, v_i) indexed

by a finite set I (with $\mathbf{1}$ being the vertex sum over the empty set). If V_i is the vertex set of Γ_i , then the vertex set of $\star_{i \in I}(\Gamma_i, v_i)$ is $(\bigsqcup_{i \in I} V_i) / \sim$, where $v_i \sim v_{i'}$ for all $i, i' \in I$. The distinguished vertex v is the one obtained from identifying all the v_i . Here it is clear from commutativity and the ‘ $v_{21} = v_{23}$ ’ case of associativity in Proposition 4.10 that the omission of parentheses or an ordering of I is harmless. At the level of underlying graphs this operation identifies the vertices v_i of the $\langle \Gamma_i \rangle$:

$$\langle \star_{i \in I}(\Gamma_i, v_i) \rangle = \star_{i \in I}(\langle \Gamma_i \rangle, v_i).$$

To deal with vertex sums where the vertices being glued in each spatial graph can vary, we introduce the following notion.

Definition 4.11. A *tree of spatial graphs* is a tuple:

$$\mathcal{T} = (T, I, J, L, (\Gamma_i)_{i \in I}, (v(l))_{l \in L}),$$

where:

- T is an abstract finite tree with vertex set $I \sqcup J$ and edge set L .
- The partition of the vertex set of T into I and J is a bipartition of T , that is, each edge $l \in L$ has one endpoint in I and the other in J . We write $i(l), j(l)$, respectively, to denote the endpoints of l in I and J .
- Each vertex in J is incident to at least two edges of T .
- The Γ_i are spatial graphs indexed by I .
- For each $l \in L$, $v(l)$ is a vertex of $\Gamma_{i(l)}$.
- If two different edges $l, l' \in L$ satisfy $i(l) = i(l')$, then $v(l) \neq v(l')$.

One should think of \mathcal{T} as a blueprint for assembling a spatial graph $[\mathcal{T}]$, called its **realization**, out of the Γ_i through iterated vertex sums. Roughly, when two distinct edges $l, l' \in L$ satisfy $j(l) = j(l')$ (and hence $i(l) \neq i(l')$), we understand this as an instruction to glue Γ_i to $\Gamma_{i'}$ along $v_l, v_{l'}$. Before making this precise, we invite the reader to study the example in Figure 4.7.

We will use an inductive procedure to define $[\mathcal{T}]$, and along the way verify that its underlying graph $\langle [\mathcal{T}] \rangle$ is as expected:

- the vertex set of $\langle [\mathcal{T}] \rangle$ is $(\bigsqcup_{i \in I} V_i) / \sim$, where V_i is the vertex set of Γ_i and $v(l) \sim v(l')$ whenever $j(l) = j(l')$,
- the edge set of $\langle [\mathcal{T}] \rangle$ is $\bigsqcup_{i \in I} E_i$, where E_i is the edge set of Γ_i .

We construct $[\mathcal{T}]$ by induction on $\#J$. If $J = \emptyset$, then either T is the empty tree, in which case we set $[\mathcal{T}] := \mathbf{0}$, or T has a single vertex $i \in I$, in which case $[\mathcal{T}] := \Gamma_i$. Either way, $\langle [\mathcal{T}] \rangle$ is as claimed.

For the inductive step, introduce the following notation: for each edge $l \in L$, the sub-graph of T obtained by removing l has precisely two connected components, each containing one endpoint of l . Let T_l be the component containing $i(l)$. Moreover, denote by I_l, J_l, L_l , respectively, the subsets of I, J, L comprised of the vertices/edges in T_l . We then define the tree of spatial graphs $\mathcal{T}_l := (T_l, I_l, J_l, L_l, (\Gamma_i)_{i \in I_l}, (v(l'))_{l' \in L_l}$.

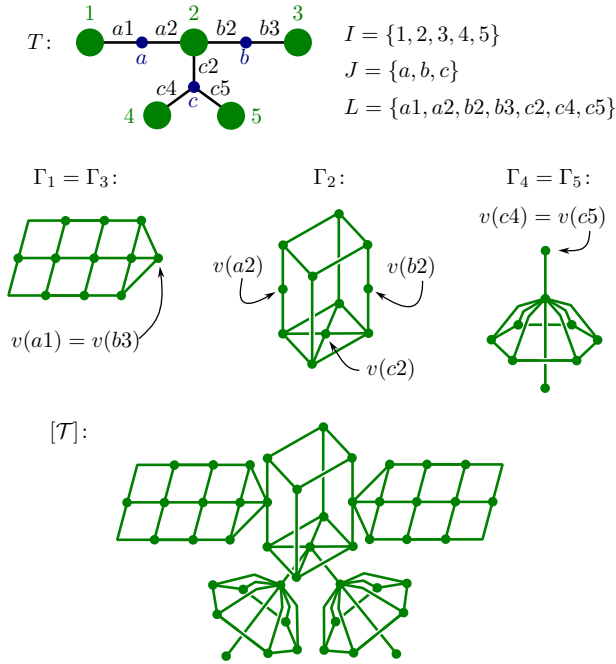


Figure 4.7 The realization of $\mathcal{T} = (T, I, J, L, (\Gamma_i)_{i \in I}, (v(l))_{l \in L})$.

Now, if J contains at least one vertex j_0 (whose choice will be shown to be immaterial), let $L_0 \subseteq L$ be the set of edges incident to j_0 . For each $l \in L_0$, the set J_l has strictly fewer elements than J . Hence we have by induction constructed realizations $[\mathcal{T}_l]$, whose underlying graphs $\langle [\mathcal{T}_l] \rangle$ are as described. In particular, $\langle [\mathcal{T}_l] \rangle$ has $v(l)$ as a vertex, hence so does $[\mathcal{T}_l]$. We then define

$$[\mathcal{T}] := \star_{l \in L_0} ([\mathcal{T}_l], v(l)).$$

By the earlier description of the underlying graph of an iterated vertex sum of pointed spatial graphs, we see,

$$\langle [\mathcal{T}] \rangle = \langle \star_{l \in L_0} ([\mathcal{T}_l], v(l)) \rangle = \star_{l \in L_0} (\langle [\mathcal{T}_l] \rangle, v(l)),$$

which matches the claimed description by some straightforward bookkeeping.

This finishes a construction of $[\mathcal{T}]$ such that $\langle [\mathcal{T}] \rangle$ is independent of choices. Next we show that $[\mathcal{T}]$ itself is also independent of the choice of j_0 .

Lemma 4.12 (Tree realizations are well defined). *Any two realizations of a tree of spatial graphs $\mathcal{T} = (T, I, J, L, (\Gamma_i)_{i \in I}, (v(l))_{l \in L})$ are isomorphic via an isomorphism inducing the identity on $\langle [\mathcal{T}] \rangle$.*




Figure 4.8. The tree T , the sub-trees T_{l_1}, T_{l_2} and their intersection \dot{T} . Large vertices represent elements of I , and small ones elements of J .

Proof. We again induct on $\#J$. When J has at most one element, no choices are made in defining $[\mathcal{T}]$, so there is nothing to show.

Suppose then that J contains two elements $j_1 \neq j_2$. For each $k \in \{1, 2\}$, denote by $[\mathcal{T}]_k$ the realization of \mathcal{T} constructed by splitting T at j_k . Moreover, let $L_k \subset L$ be the set of edges incident with j_k , and consider, for each $l \in L_k$, the tree of spatial graphs $\mathcal{T}_l := (T_l, I_l, J_l, L_l, (\Gamma_i)_{i \in I_l}, (v(l'))_{l' \in L_l})$ defined as before.

Now, there is exactly one edge $l_1 \in L_1$ such that T_{l_1} contains the vertex j_2 , and one edge $l_2 \in L_2$ such that T_{l_2} contains j_1 . Consider the tree $\dot{T} := T_{l_1} \cap T_{l_2}$. We have a tree of spatial graphs:

$$\dot{\mathcal{T}} := (\dot{T}, \dot{I}, \dot{J}, \dot{L}, (\Gamma_i)_{i \in \dot{I}}, (v(l'))_{l' \in \dot{L}}),$$

where $\dot{I} := I_{l_1} \cap I_{l_2}$, $\dot{J} := J_{l_1} \cap J_{l_2}$, and $\dot{L} := L_{l_1} \cap L_{l_2}$; see Figure 4.8.

By induction, the realizations $[\mathcal{T}_{l_k}]$ are well-defined. One readily checks that:

$$[\mathcal{T}_{l_1}] = [\dot{\mathcal{T}}]_{v(l_2) \bullet v_2} \star_{l \in L_2 \setminus \{l_2\}} ([\mathcal{T}_l], v(l)),$$

$$[\mathcal{T}_{l_2}] = [\dot{\mathcal{T}}]_{v(l_1) \bullet v_1} \star_{l \in L_1 \setminus \{l_1\}} ([\mathcal{T}_l], v(l)),$$

where v_k is the result of identifying the vertices $v(l)$ with $l \in L_k \setminus \{l_k\}$. We finish the proof by applying the ‘ $v_{21} \neq v_{23}$ ’ case of associativity from Proposition 4.10:

$$[\mathcal{T}]_1 = [\mathcal{T}_{l_1}]_{v(l_1) \bullet v_1} \left(\star_{l \in L_1 \setminus \{l_1\}} ([\mathcal{T}_l], v(l)) \right)$$

$$= \left(\left(\star_{l \in L_2 \setminus \{l_2\}} ([\mathcal{T}_l], v(l)) \right)_{v_2 \bullet v(l_2)} [\dot{\mathcal{T}}]_{v(l_1) \bullet v_1} \left(\star_{l \in L_1 \setminus \{l_1\}} ([\mathcal{T}_l], v(l)) \right) \right)$$

$$= \left(\star_{l \in L_2 \setminus \{l_2\}} ([\mathcal{T}_l], v(l)) \right)_{v_2 \bullet v(l_2)} \left([\dot{\mathcal{T}}]_{v(l_1) \bullet v_1} \left(\star_{l \in L_1 \setminus \{l_1\}} ([\mathcal{T}_l], v(l)) \right) \right)$$

$$= \left(\star_{l \in L_2 \setminus \{l_2\}} ([\mathcal{T}_l], v(l)) \right)_{v_2 \bullet_{v(l_2)} [\mathcal{T}_{l_2}]} = [\mathcal{T}]_2.$$

□

Since realizations of trees are constructed by iterated vertex sums, Lemmas 4.8 and 4.9 have the following generalizations.

Lemma 4.13 (Sub-graphs of a tree of spatial graphs). *For a tree of spatial graphs $\mathcal{T} = (T, I, J, L, (\Gamma_i)_{i \in I}, (v(l))_{l \in L})$ and for each $i \in I$, let Γ'_i be the sub-graph of $[\mathcal{T}]$ comprised of the vertices and edges of Γ_i . Then $\Gamma'_i = \Gamma_i$.*

Lemma 4.14 (Trees of isomorphisms). *For each $k \in \{1, 2\}$ fix a tree of spatial graphs $\mathcal{T}_k = (T_k, I_k, J_k, L_k, (\Gamma_i)_{i \in I_k}, (v(l))_{l \in L_k})$. Fix also the data of:*

- an isomorphism of trees $f: T_1 \rightarrow T_2$ such that $f(I_1) = I_2$ (hence $f(J_1) = J_2$), and
- for each $i \in I_1$ an isomorphism $\Phi_i: \Gamma_i \rightarrow \Gamma_{f(i)}$, such that the collection $(\Phi_i)_{i \in I_1}$ respects the assignments $l \mapsto v(l)$ on L_1 and L_2 , that is, for every $l \in L_1$, we have $\Phi_{i(l)}(v(l)) = v(f(l))$.

Then, there is an isomorphism $\Phi: [\mathcal{T}_1] \rightarrow [\mathcal{T}_2]$ such that for every $i \in I_1$, the underlying isomorphism $\langle \Phi \rangle$ restricts to $\langle \Phi_i \rangle$ on the sub-graph $\langle \Gamma_i \rangle$ of $\langle [\mathcal{T}_1] \rangle$.

4.3. Decomposing pieces as trees of blocks

From now on several statements include a non-split assumption on spatial graphs. Incidentally, we collect the following observation:

Lemma 4.15 (Vertex sum preserves being non-split). *Let $(\Gamma_1, v_1), (\Gamma_2, v_2)$ be pointed spatial graphs. Then $\Gamma_1 \bullet_{v_1 \bullet_{v_2}} \Gamma_2$ is non-split if and only if both Γ_1, Γ_2 are non-split.*

Proof. If a vertex summand, say Γ_1 , is split, let \mathcal{S}_1 be its ambient sphere and let $S \subset \mathcal{S}_1$ be a splitting sphere. Choose an enclosing ball B_1 for (Γ_1, v_1) containing S in its interior. Then if we use B_1 to form the vertex sum, S will be contained in the ambient sphere of $\Gamma_1 \bullet_{v_1 \bullet_{v_2}} \Gamma_2$, with both sides of S intersecting $|\Gamma_1 \bullet_{v_1 \bullet_{v_2}} \Gamma_2|$. Hence S is a splitting sphere for $\Gamma_1 \bullet_{v_1 \bullet_{v_2}} \Gamma_2$ by Lemma 3.8.

Conversely, suppose S is a splitting sphere for $\Gamma_1 \bullet_{v_1 \bullet_{v_2}} \Gamma_2$. The component of $\mathcal{S} \setminus S$ not containing the vertex $v_1 = v_2$ has non-empty intersection with the support of a summand, say Γ_1 . Then both components of $\mathcal{S} \setminus S$ intersect $|\Gamma_1|$ and so, regarding Γ_1 as a sub-graph of $\Gamma_1 \bullet_{v_1 \bullet_{v_2}} \Gamma_2$, we see S is a splitting sphere for Γ_1 . □

Corollary 4.16 (Trees of spatial graphs preserve being non-split). *Let,*

$$\mathcal{T} = (T, I, J, L, (\Gamma_i)_{i \in I}, (v(l))_{l \in L})$$

be a tree of spatial graphs. Then $[\mathcal{T}]$ is non-split if and only if each Γ_i is non-split.




Figure 4.9. A spatial graph Γ whose vertices are cut in $\langle \Gamma \rangle$ but not in Γ .

The following analogue of Lemma 3.8 has essentially the same proof:

Lemma 4.17 (If it looks like a vertex sum, it is a vertex sum). *Let Γ be a spatial graph in \mathcal{S} and $S \subset \mathcal{S}$ a PL-embedded 2-sphere intersecting $|\Gamma|$ precisely at one vertex v of Γ . Denote the closures of the two components of $\mathcal{S} \setminus S$ by B_1 and B_2 . For each $i \in \{1, 2\}$, let Γ_i be the sub-graph of Γ comprised of the vertices and edges contained in B_i . Then $\Gamma = \Gamma_1 \bullet_v \Gamma_2$.*

Definition 4.18. *Let Γ be a spatial graph in \mathcal{S} .*

- *If $S \subset \mathcal{S}$ is a 2-sphere as in Lemma 4.17, we say that ‘ S decomposes Γ as $\Gamma_1 \bullet_v \Gamma_2$ ’.*
- *If S is a 2-sphere decomposing Γ as $\Gamma_1 \bullet_v \Gamma_2$ such that for each $k \in \{1, 2\}$ the piece of Γ_k containing v is not isomorphic to $\mathbf{1}$, then v is called a **cut vertex** of Γ and S a **cut sphere** of Γ .*
- *Γ is called a **block** if Γ is a piece without cut vertices and $\Gamma \not\cong \mathbf{1}$.*
- *A tree of spatial graphs $\mathcal{T} = (T, I, J, L, (\Lambda_i)_{i \in I}, (v(l))_{l \in L})$ where each Λ_i is a block is called a **tree of blocks**, and we say ‘ \mathcal{T} is a tree of blocks for $[\mathcal{T}]$ ’.*

There is a standard notion of cut vertex in the abstract setting: a vertex v of a connected abstract graph G is **cut** if G is the union of sub-graphs G_1, G_2 intersecting precisely at v , with neither G_i consisting of a single vertex. It is however possible for a vertex of a spatial graph Γ to be cut in $\langle \Gamma \rangle$ but not in Γ ; see Figure 4.9.

Lemma 4.19 (Spheres sort blocks). *Let Λ be a block in \mathcal{S} , and let $S \subset \mathcal{S}$ be a PL-embedded 2-sphere that intersects $|\Lambda|$ either at a single vertex of Λ , or not at all. Denote the closures of the two components of $\mathcal{S} \setminus S$ by B_1, B_2 . Then $|\Lambda|$ is contained in exactly one of the B_i .*

Proof. Since Λ is a piece, the case where $S \cap |\Lambda| = \emptyset$ follows from Lemma 3.10. If $S \cap |\Lambda|$ consists of one vertex v of Λ , then since $\Lambda \not\cong \mathbf{1}$, certainly $|\Lambda|$ is not contained in both B_i . By Lemma 4.17, S decomposes Λ as $\Lambda_1 \bullet_v \Lambda_2$. As Λ has no cut vertices, one of the summands, say Γ_1 , is isomorphic to $\mathbf{1}$. This means $\Lambda = \Lambda_2$. □

Cut vertices are easily read off in the realization of a tree of blocks:

Proposition 4.20 (Cut vertices in a tree of blocks). *Fix a tree of blocks $\mathcal{T} = (T, I, J, L, (\Lambda_i)_{i \in I}, (v(l))_{l \in L})$. For each $j \in J$, denote by $v(j)$ the vertex of $[\mathcal{T}]$ that results from identifying all $v(l)$ with l incident to j . Then the correspondence $j \mapsto v(j)$ is a bijection between J and the set of cut vertices of $[\mathcal{T}]$.*

Proof. To see that each $v(j)$ is cut: by definition of a tree of spatial graphs, j has degree at least 2, so if $L_0 \subseteq L$ is the set of edges incident to j , one has a non-trivial partition $L_0 = L_1 \sqcup L_2$. For each $k \in \{1, 2\}$, choose $l_k \in L_k$. Then $\Lambda_{i(l_k)}$, being a block, has an edge, hence $[\mathcal{T}_{l_k}]$ also has an edge. Thus each vertex summand in:

$$[\mathcal{T}] = \left(\star_{l \in L_1} ([\mathcal{T}_l], v(l)) \right)_{v(l_1) \bullet v(l_2)} \left(\star_{l \in L_2} ([\mathcal{T}_l], v(l)) \right),$$

has an edge and so is not isomorphic to $\mathbf{1}$. Thus $v(j) = v(l_1) = v(l_2)$ is cut. Looking at the vertex set of $[\mathcal{T}]$, as given by the description of $\langle [\mathcal{T}] \rangle$, it is clear that the assignment $j \mapsto v(j)$ is injective.

Conversely, suppose v is a vertex of $[\mathcal{T}]$ that does not result from such an identification, and consider a PL-embedded 2-sphere S in the ambient sphere \mathcal{S} of $[\mathcal{T}]$ intersecting $[[\mathcal{T}]]$ precisely at v . Say S decomposes $[\mathcal{T}]$ as $\Gamma_1 \bullet_v \Gamma_2$; we argue that some Γ_i is $\mathbf{1}$. The on v implies that all edges of $[\mathcal{T}]$ incident to v come from the same block Λ_i . Using Lemma 4.13 to regard Λ_i as a sub-graph of $[\mathcal{T}]$, we see from Lemma 4.19 that all edges of $[\mathcal{T}]$ incident to v are in one of the Γ_i , say in Γ_1 . Hence v is an isolated vertex of Γ_2 . But since $[\mathcal{T}]$ is non-split by Corollary 4.16, Γ_2 is non-split by Lemma 4.15, which implies $\Gamma_2 \cong \mathbf{1}$. \square

Proposition 4.21 (Existence of trees of blocks). *Every non-split spatial graph $\Gamma \not\cong \mathbf{1}$ is the realization of some tree of blocks.*

Proof. We induct on the number of edges in Γ to produce a tree of blocks $\mathcal{T} = (T, I, J, L, (\Lambda_i)_{i \in I}, (v(l))_{l \in L})$ realizing Γ . If Γ has no edges, then $\Gamma = \mathbf{0}$ and we take T to be the empty tree.

Now suppose Γ has edges. If Γ is a block, take T to be a tree with a single vertex $i \in I$ and set $\Lambda_i := \Gamma$. Otherwise Γ can be expressed as $\Gamma_1 \bullet_v \Gamma_2$ with each Γ_k not isomorphic to $\mathbf{1}$, and also non-split by Lemma 4.15. Hence both Γ_k have at least one edge, and thus fewer edges than Γ , so the induction hypothesis applies.

Let $\mathcal{T}_k = (T_k, I_k, J_k, L_k, (\Lambda_i)_{i \in I_k}, (v(l))_{l \in L_k})$ be a tree of blocks for Γ_k . We construct T as in Figure 4.10 from modified versions T'_k of the T_k , according to the following two cases:

- If v is not a cut vertex of Γ_k , so by Proposition 4.20 there is no edge $l_k \in L_k$ with $v(l_k) = v$, construct T'_k from T_k by adding a new vertex j_k and a new edge l_k connecting j_k to the vertex $i_k \in I$ whose block Λ_{i_k} contains v . We also write

$$J'_k := J_k \sqcup \{j_k\}, \quad L'_k := L_k \sqcup \{l_k\}, \quad L_k^0 := \{l_k\},$$

and set $v(l_k) := v$. (In this case, T'_k with its vertex set partitioned as $I_k \sqcup J'_k$ is not admissible as the tree in a tree of spatial graphs, since j_k is a leaf.)

- If v is a cut vertex of Γ_k , then there is a corresponding $j_k \in J_k$, whose set of incident edges we denote by L_k^0 . In Γ_k , the vertices $v(l)$ with $l \in L_k^0$ are identified into the vertex v . Set $T'_k := T_k$, $J'_k := J_k$, $L'_k := L_k$.




Figure 4.10. Constructing T from T_1 and T_2 . Large (resp. small) vertices represent elements of I_1, I_2 (resp. of J_1, J_2). Elements of I_1, I_2 whose corresponding blocks contain v are indicated. Here, v is not a cut vertex in Γ_1 , but it is in Γ_2 , where it corresponds to j_2 .

We define

$$T := T_1' \bullet_{j_1} T_2, \quad I := I_1 \sqcup I_2, \quad J := (J_1' \sqcup J_2') / j_1 \sim j_2, \quad L := L_1' \sqcup L_2'$$

and this turns \mathcal{T} into a tree of spatial graphs whose realization is Γ :

$$\begin{aligned} [\mathcal{T}] &= \star_{\substack{l \in L \\ j(l)=j_1=j_2}} ([\mathcal{T}_l], v(l)) \\ &= \left(\star_{l \in L_1^0} ([\mathcal{T}_l], v(l)) \right) \bullet_v \left(\star_{l \in L_2^0} ([\mathcal{T}_l], v(l)) \right) \\ &= [\mathcal{T}_1] \bullet_v [\mathcal{T}_2] = \Gamma_1 \bullet_v \Gamma_2. \end{aligned}$$

□

Proposition 4.22 (Uniqueness of trees of blocks). *For each $k \in \{1, 2\}$, let $\mathcal{T}_k = (T_k, I_k, J_k, L_k, (\Lambda_i)_{i \in I_k}, (v(l))_{l \in L_k})$ be a tree of blocks, and let $\Phi: [\mathcal{T}_1] \rightarrow [\mathcal{T}_2]$ be an isomorphism. Then there is an isomorphism of trees $f: T_1 \rightarrow T_2$ satisfying $f(I_1) = I_2$, and such that:*

- for each $i \in I_1$, the map Φ is an isomorphism $\Lambda_i \rightarrow \Lambda_{f(i)}$, and
- for each $l \in L_1$, we have $\Phi(v(l)) = v(f(l))$.

The second item implies that f respects the bijection given by Proposition 4.20 between the J_k and the set of cut vertices of $[\mathcal{T}_k]$: for each $j \in J_1$ we have $\Phi(v(j)) = v(f(j))$.

Proof. We induct on $\#I_1$. If $I_1 = \emptyset$, then $[\mathcal{T}_1] = \mathbf{0} = [\mathcal{T}_2]$, so T_2 is the empty tree and there is nothing to show. If I_1 is comprised of a single element i_1 , and hence $J_1 = \emptyset$, then $[\mathcal{T}_1] = \Lambda_{i_1}$ is a block, so $[\mathcal{T}_2]$ is a block. In particular, $[\mathcal{T}_2]$ has no cut vertices and so by Proposition 4.20 we conclude $J_2 = \emptyset$. Hence I_2 contains exactly one element i_2 , with $[\mathcal{T}_2] = \Lambda_{i_2}$, and we set $f(i_1) := i_2$.

Assume now that I_1 contains at least two elements, so $J_1 \neq \emptyset$. Choose $j_1 \in J_1$, write $v_1 := v(j_1)$, and let S_1 be a cut sphere for $[\mathcal{T}_1]$ decomposing it as $[\mathcal{T}_1] = \Gamma_1^+ \bullet_{v_1} \Gamma_1^-$, so Γ_1^+, Γ_1^- are pieces non-isomorphic to $\mathbf{1}$. Similarly, $v_2 := \Phi(v_1)$ is a cut vertex for $[\mathcal{T}_2]$, so let $j_2 \in J_2$ be the corresponding element. The sphere $S_2 := \Phi(S_1)$ is now a cut sphere for $[\mathcal{T}_2]$ decomposing it as $[\mathcal{T}_2] = \Gamma_2^+ \bullet_{v_2} \Gamma_2^-$, with Φ giving isomorphisms of sub-graphs $\Phi^\epsilon: \Gamma_1^\epsilon \rightarrow \Gamma_2^\epsilon$, for each $\epsilon \in \{+, -\}$.

Let $k \in \{1, 2\}$. Our goal is to extract from \mathcal{T}_k a description of the Γ_k^+, Γ_k^- as realizations of trees of blocks, to which we then apply the induction hypothesis. The procedure is analogous for all four spatial graphs, so fix $\epsilon \in \{+, -\}$.

Denote by $L_k^0 \subseteq L_k$ the set of edges incident to j_k . By Lemma 4.19, for each $i \in I_k$, the block Λ_i is a sub-graph of exactly one among Γ_k^+, Γ_k^- . Consider the partition $L_k^0 = L_k^{0+} \sqcup L_k^{0-}$, where $l \in L_k^0$ is in $L_k^{0\epsilon}$ if $\Lambda_{i(l)}$ is a sub-graph of Γ_k^ϵ . Then $[\mathcal{T}_k]$ decomposes as:

$$[\mathcal{T}_k] = \left(\star_{l \in L_k^{0+}} ([(\mathcal{T}_k)_l], v(l)) \right) \bullet_{v_k} \left(\star_{l \in L_k^{0-}} ([(\mathcal{T}_k)_l], v(l)) \right).$$

We claim that this is the same as the decomposition given by S_k , that is:

$$\Gamma_k^\epsilon = \star_{l \in L_k^{0\epsilon}} ([(\mathcal{T}_k)_l], v(l)).$$

To see this, first notice that since the Λ_i are non-split, each $[(\mathcal{T}_k)_l]$ is also non-split by Corollary 4.16. Now, for every $l \in L_k^0$, it follows from Proposition 4.20 that $v(l)$ is not a cut vertex of $[(\mathcal{T}_k)_l]$. Therefore, S_k decomposes $[(\mathcal{T}_k)_l]$ as a trivial vertex sum $[(\mathcal{T}_k)_l] \bullet_{v(l)} \mathbf{1}$. In other words, $[(\mathcal{T}_k)_l]$ is entirely contained in one side of S_k , which must be the same as $|\Lambda_{i(l)}|$. Thus, each $\star_{l \in L_k^{0\epsilon}} ([(\mathcal{T}_k)_l], v(l))$ is a sub-graph of Γ_k^ϵ , whence the above description of the Γ_k^ϵ holds.

Next, we write down an explicit tree of blocks \mathcal{T}_k^ϵ for $\star_{l \in L_k^{0\epsilon}} ([(\mathcal{T}_k)_l], v(l))$. If $L_k^{0\epsilon}$ has only one element l_k^ϵ , put $\mathcal{T}_k^\epsilon := (\mathcal{T}_k)_{l_k^\epsilon}$. Otherwise, recover the notation introduced when defining the realization of a tree of spatial graphs:

$$(\mathcal{T}_k)_l = ((T_k)_l, (I_k)_l, (J_k)_l, (L_k)_l, (\Lambda_i)_{i \in (I_k)_l}, (v(l'))_{l' \in (L_k)_l}),$$

and set $\mathcal{T}_k^\epsilon := (T_k^\epsilon, I_k^\epsilon, J_k^\epsilon, L_k^\epsilon, (\Lambda_i)_{i \in I_k^\epsilon}, (v(l))_{l \in L_k^\epsilon})$ to be the tree of blocks comprised of the branches of \mathcal{T}_k at j_k that stem from edges in $L_k^{0\epsilon}$. Explicitly, T_k^ϵ is the sub-tree of T_k

with vertex and edge sets given by:

$$I_k^\epsilon := \bigsqcup_{l \in L_k^{0\epsilon}} (I_k)_l, \quad J_k^\epsilon := \{j_k\} \sqcup \bigsqcup_{l \in L_k^{0\epsilon}} (J_k)_l, \quad L_k^\epsilon := L_k^{0\epsilon} \sqcup \bigsqcup_{l \in L_k^{0\epsilon}} (J_k)_l.$$

Observe that in the first case $[\mathcal{T}_k^\epsilon]$ does not have v_k as a cut vertex, and in the second case it does, with j_k being the corresponding element of J_k^ϵ .

It is now clear that, in either case, $[\mathcal{T}_k^\epsilon] = \star_{l \in L_k^{0\epsilon}} ([(\mathcal{T}_k)_l], v(l)) = \Gamma_i^\epsilon$. By induction hypothesis, the isomorphisms $\Phi^\epsilon: [\mathcal{T}_1^\epsilon] \rightarrow [\mathcal{T}_2^\epsilon]$ yield tree isomorphisms $f^\epsilon: T_1^\epsilon \rightarrow T_2^\epsilon$, which we now assemble to the desired $f: T_1 \rightarrow T_2$. On each sub-tree T_1^ϵ of T_1 , we want to set $f = f^\epsilon$ but have to ensure that f^+ and f^- agree where they overlap, and we must also define f on vertices and edges of T_1 that are not in any of the T_1^ϵ .

Fix $\epsilon \in \{+, -\}$ for this paragraph. The isomorphism Φ^ϵ ensures that v_1 is a cut vertex of $[\mathcal{T}_1^\epsilon]$ if and only if v_2 is a cut vertex of $[\mathcal{T}_2^\epsilon]$. If this is the case, then for both $k \in \{1, 2\}$, the vertex j_k of T_k is in T_k^ϵ , along with the edges in $L_k^{0\epsilon}$. Moreover, in this situation we have in \mathcal{T}_2^ϵ that $v(j_2) = v_2 = \Phi^\epsilon(v_1) = \Phi^\epsilon(v(j_1)) = v(f^\epsilon(j_1))$, whence it follows by injectivity of $j \mapsto v(j)$ that $j_2 = f^\epsilon(j_1)$. On the other hand, if one (hence both) v_k is not cut in $[\mathcal{T}_k^\epsilon]$, then the corresponding j_k and the unique edge l_k^ϵ in $L_k^{0\epsilon}$ are not in T_k^ϵ . In this situation, $i(l_k^\epsilon)$ is the only element of I_k^ϵ whose corresponding block $\Lambda_{i(l_k^\epsilon)}$ contains v_k as a vertex.

At this point there are three cases to consider:

- If for some (hence both) $k \in \{1, 2\}$ the vertex v_k is cut in $[\mathcal{T}_k^+]$ and $[\mathcal{T}_k^-]$, then the sub-trees T_k^+, T_k^- jointly cover T_k , and they overlap precisely at j_k . As $f^+(j_1) = j_2 = f^-(j_1)$, we can glue together the f^ϵ into the desired $f: T_1 \rightarrow T_2$.
- Suppose, for both $k \in \{1, 2\}$, the vertex v_k is cut in $[\mathcal{T}_k^+]$ but not in $[\mathcal{T}_k^-]$ (the reverse situation being analogous). Then T_k^+ and T_k^- do not overlap, and jointly they cover all of T_k except for the edge l_k^- described above. In this case, extend the definition of f^+, f^- to all of T_1 by setting $f(l_1^-) := l_2^-$. This respects the endpoints of the edge: we have seen that $f^-(j_1) = j_2$, and the characterization of $i(l_k^\epsilon)$ given above, together with the fact that $\Phi^-(v_1) = v_2$, shows that $f^-(i(l_1^-)) = i(l_2^-)$.
- If for both $k \in \{1, 2\}$ the vertex v_k is not cut in $[\mathcal{T}_k^+]$ nor in $[\mathcal{T}_k^-]$, then T_k^+, T_k^- are disjoint and cover all of T_k except for the vertex j_k and its only two incident edges l_k^+, l_k^- . We extend f^+, f^- by putting $f(j_1) := j_2$ and, for each $\epsilon \in \{+, -\}$, setting $f(l_1^\epsilon) := l_2^\epsilon$. This respects the incidence of each l_1^ϵ at the endpoint j_1 , and for the other endpoint we argue as in the previous case.

Having defined the isomorphism $f: T_1 \rightarrow T_2$, almost all properties stated in the proposition are inherited from the f^ϵ . We are only left to check that, in the second and third cases above, the definition of f on the new edge(s) l_1^ϵ satisfies $\Phi(v(l_1^\epsilon)) = v(f(l_1^\epsilon))$. And indeed it does: $\Phi(v(l_1^\epsilon)) = \Phi(v_1) = v_2 = v(l_2^\epsilon) = v(f(l_1^\epsilon))$. □

Proposition 4.22 works in tandem with Proposition 3.11, which reduced recognition of spatial graphs to recognition of pieces. To test whether two pieces are isomorphic, first decompose them as trees of blocks – the algorithmic details are given below in Lemma 7.8. Then, Proposition 4.22 guarantees that the pieces are isomorphic if and only

if the blocks are pairwise isomorphic, via isomorphisms respecting the structure of the trees. Comparing the isomorphism type of blocks is within reach using the Recognition Theorem (Theorem 7.4).

5. Extension to decorated spatial graphs

The theory developed in the previous sections generalizes to spatial graphs equipped with additional structure. Three natural extensions are directed spatial graphs, and spatial graphs with colourings of edges and/or vertices. Here we formalize these concepts and comment on how the operations and decompositions are adapted to such settings. Proofs require no additional insight, so we omit them.

A **directed** spatial graph Γ together with a choice of orientation of each edge. If e is a non-loop edge of Γ and $h: [-1, 1] \rightarrow e$ is a PL homeomorphism orienting e , we say the vertex $h(-1)$ is the **source** of e , and $h(1)$ is its **target**. When e is a loop, the only vertex of Γ contained in e is simultaneously the source and the target. We denote the source and target of an edge e by $s(e)$ and $t(e)$, respectively. An isomorphism of directed spatial graphs is an isomorphism of the spatial graphs such that all induced PL homeomorphisms between the edges are orientation-preserving.

A **vertex colouring** of a spatial graph $\Gamma = (\mathcal{S}, V, E)$ is a function $f: V \rightarrow \mathbb{N}$ from the vertex set to the non-negative integers. For each vertex $v \in V$, we refer to $f(v)$ as the **colour** of v . Given spatial graphs $\Gamma_1 = (\mathcal{S}_1, V_1, E_1), \Gamma_2 = (\mathcal{S}_2, V_2, E_2)$ with vertex colourings f_1, f_2 , an isomorphism $\Phi: \Gamma_1 \rightarrow \Gamma_2$ is said to preserve vertex colourings if the bijection $\Phi|_{V_1}: V_1 \rightarrow V_2$ satisfies $f_1 = f_2 \circ \Phi|_{V_1}$. Analogously we define an **edge colouring** $g: E \rightarrow \mathbb{N}$, and what it means for an isomorphism of spatial graphs to preserve edge colourings.

One may consider spatial graphs with any (possibly empty) combination of these three types of structure, and we will broadly refer to such spatial graphs as **decorated**. By two spatial graphs carrying a decoration ‘of the same type’, we mean that the combination of additional structures is the same. An isomorphism of spatial graphs with decorations of the same type is an isomorphism of the corresponding undecorated spatial graphs that respects all extra structure. Sub-graphs and the underlying abstract graphs also inherit decorations of the same type. The induced decoration on $\langle \Gamma \rangle$ determines the decoration of Γ , except for one ambiguity: if Γ is directed, the orientation of a loop e cannot be inferred from $s(e)$ and $t(e)$.

Lemma 5.1 (Compatibility with decorations via underlying graphs). *Let Γ_1, Γ_2 be spatial graphs with decorations of the same type, and $\Phi: \Gamma_1^- \rightarrow \Gamma_2^-$ an isomorphism of the corresponding undecorated spatial graphs. Assume moreover that the Γ_i have no loops, or are not directed. Then Φ respects the decorations on the Γ_i if and only if $\langle \Phi \rangle: \langle \Gamma_1^- \rangle \rightarrow \langle \Gamma_2^- \rangle$ respects the decorations on the $\langle \Gamma_i \rangle$.*

We summarize how to adapt main definitions and statements regarding disjoint unions and vertex sums of spatial graphs with decorations are of the same type:

- The disjoint union of decorated spatial graphs is the disjoint union of the underlying spatial graphs, carrying a decoration of the same type.

- The vertex sum of pointed decorated spatial graphs is similarly defined provided that, in case a vertex colouring is part of the decoration, the basepoints are of the same colour.
- Isomorphisms between decorated spatial graphs can be assembled along disjoint unions and vertex sums in the sense of Lemmas 3.6 and 4.9.
- In the decorated set-up, there is a well-defined identity element $\mathbf{0}$ for the disjoint union, but if a vertex colouring is part of the decoration, there is one isomorphism type $\mathbf{1}_c$ of one-point spatial graph for each colour $c \in \mathbb{N}$.
- The properties listed in Propositions 3.7 and 4.10 hold for decorated spatial graphs. If a vertex colouring is part of the decoration, the occurrence of $\mathbf{1}$ in Proposition 4.10 should read $\mathbf{1}_c$, where c is the colour of v_1 .
- The definitions of splitting sphere, split spatial graph and piece (Definition 3.9) remain unchanged in the decorated setting. Every decorated spatial graph can be expressed as an iterated disjoint union of (decorated) pieces, in a way that is unique in the sense of Proposition 3.11.
- In a tree of (decorated) spatial graphs (Definition 4.11), all Γ_i should have a decoration of the same type. If a vertex colouring is part of the decoration we additionally require that, for each $j \in J$, the vertices $v(l)$ with l incident to j be all of the same colour. Realizations of trees of decorated spatial graphs are well-defined in the sense of Lemma 4.12, carrying a canonical decoration of the same type.
- In the vertex-coloured version of cut vertex, cut sphere and block (Definition 4.18), occurrences of the expression ‘not isomorphic to $\mathbf{1}$ ’ should read ‘not isomorphic to any $\mathbf{1}_c$ ’. The definition of tree of blocks does not change.
- Propositions 4.21 and 4.22 apply to decorated graphs: every non-split decorated spatial graph that is not isomorphic to a one-point graph is the realization of a tree of decorated blocks in a unique way.

Vertex colourings will be used in an essential way for establishing algorithmic recognition of pieces (Proposition 7.10), even when the pieces do not have a vertex colouring. Explicitly, we will use vertex colourings in the proof of Lemma 7.7 to encode the requirement that the isomorphisms between blocks in a tree of blocks respect the combinatorial structure (the second item of Proposition 4.22).

6. The marked exterior of a spatial graph

In this section, we construct the marked exterior of a (decorated) spatial graph $\Gamma = (\mathcal{S}, V, E)$, which will be a ‘manifold with boundary pattern’. We explain how it encodes the spatial graph used to construct it and translate indecomposability properties of spatial graphs into properties of their marked exteriors.

6.1. Construction and faithfulness

We first describe a simpler variant, the *oriented* marked exterior, which will be a pair $(X_\Gamma^\circ, P_\Gamma^\circ)$, with X_Γ° an oriented PL 3-manifold and P_Γ° an oriented one-dimensional submanifold of ∂X_Γ° ; see Figure 6.1 for an illustration.




Figure 6.1. An oriented marked exterior of a spatial graph. Arrows indicate orientations.

First, choose a regular neighbourhood N_V of V in the pair $(\mathcal{S}, |\Gamma|)$. Since V is discrete, N_V is a disjoint union of β -balls [15, Corollary 3.12], each containing exactly one vertex $v \in V$. Denote the corresponding ball by N_v and let X_V be the compact PL 3-manifold $\mathcal{S} \setminus \text{int}(N_V)$.

Next, choose a regular neighbourhood N_E of $|\Gamma| \cap X_V$ in X_V . As $|\Gamma| \cap X_V$ is the disjoint union of the properly embedded arcs $e \cap X_V$, with $e \in E$, the regular neighbourhood N_e of each $e \cap X_V$ is a β -ball and $(N_e, e \cap X_V)$ is an unknotted ball pair [15, Corollary 3.27].³

We now set $X_\Gamma^\circ := \mathcal{S} \setminus \text{int}(N_V \cup N_E)$, which is a compact PL 3-manifold with $\partial X_\Gamma^\circ \subseteq \partial X_V \cup \partial N_E$. For each vertex $v \in V$, we call $R_v := \partial N_v \cap \partial X_\Gamma^\circ$ the **vertex region** of v and for each edge $e \in E$ we call $R_e := \partial N_e \cap \partial X_\Gamma^\circ$ the **edge region** of e . Notice that R_e is always an annulus, and if v has degree d , then R_v is a surface of genus 0 with d boundary components.

To define P_Γ° , note that the vertex regions and edge regions of ∂X_Γ° intersect along circles, which we call **junctions**. We set $P_\Gamma^\circ := (\bigcup_{v \in V} R_v) \cap (\bigcup_{e \in E} R_e)$ as the union of these junctions. Then P_Γ° separates ∂X_Γ° into the vertex and edge regions. As X_Γ° inherits an orientation from \mathcal{S} , also ∂X_Γ° has a canonical orientation. Now we orient $P_\Gamma^\circ = \partial(\bigcup_{v \in V} R_v)$ as the boundary of $\bigcup_{v \in V} R_v$ (viewing P_Γ° as the boundary of $\bigcup_{e \in E} R_e$ instead would induce the opposite orientation). Hence, the orientation of P_Γ° determines which regions of ∂X_Γ° are vertex regions. As each edge is incident to at least one vertex, each vertex or edge region of ∂X_Γ° that has no boundary has to be a vertex region (corresponding to an isolated vertex).

Definition 6.1. An *oriented marked exterior* of a spatial graph Γ is a pair $(X_\Gamma^\circ, P_\Gamma^\circ)$, where X_Γ° and P_Γ° are oriented PL manifolds obtained by the above construction.

Proposition 6.2. (Oriented marked exteriors are well-defined). Let Γ be a spatial graph with two oriented marked exteriors (X_k°, P_k°) , $k \in \{1, 2\}$. Then there exists an orientation-preserving homeomorphism of pairs $\Phi: (X_1^\circ, P_1^\circ) \rightarrow (X_2^\circ, P_2^\circ)$ such that, for each vertex v of Γ , the map Φ sends the corresponding vertex region $R_{v,1}$ to $R_{v,2}$, and similarly for edge regions.

³ The reference states only that N_e is a β -ball, but the proof of Rourke–Sanderson’s Theorem 3.26 in our particular case reveals that the ball pair $(N_e, e \cap X_V)$ is unknotted.

Proof. Let $N_{V,k}$ and $N_{E,k}$ be the regular neighbourhoods from the construction of X_k° . Similarly, use the corresponding notation $N_{v,k}, N_{e,k}$ for the components corresponding to single vertices and edges, as well as $X_{V,k}$ for $\mathcal{S} \setminus \text{int}(N_{V,k})$.

By the Regular Neighbourhood Theorem for pairs [15, Theorem 4.11] there is a PL isotopy of \mathcal{S} carrying $N_{V,1}$ to $N_{V,2}$ and fixing $|\Gamma|$. This induces an orientation-preserving PL homeomorphism $\Phi_V: N_{V,1} \rightarrow N_{V,2}$. As $|\Gamma|$ is fixed during the isotopy, Φ_V maps $N_{v,1}$ to $N_{v,2}$ for each vertex v of Γ and fixes the edges as well.

Since PL homeomorphisms take regular neighbourhoods to regular neighbourhoods, $\Phi_V(N_{E,1})$ and $N_{E,2}$ are regular neighbourhoods of $|\Gamma| \cap X_{V,2}$ in $X_{V,2}$. Using the Regular Neighbourhood Theorem again, we find a PL isotopy of $X_{V,2}$ that carries $\Phi_V(N_{E,1})$ to $N_{E,2}$. This isotopy carries $\Phi_V(X_1^\circ)$ to X_2° . This shows that the two marked exteriors are in fact PL-homeomorphic via an orientation-preserving PL homeomorphism $\Phi: X_1^\circ \rightarrow X_2^\circ$. As these isotopies fix $|\Gamma|$, each vertex/edge region of the boundary gets mapped to the corresponding vertex/edge region, thus mapping P_1° to P_2° . As Φ is orientation-preserving and the orientation of the junctures is determined by the orientation of X_k , Φ also preserves the orientation of the P_k° . \square

We thus refer to *the* oriented marked exterior of a spatial graph.

Proposition 6.3 (Faithfulness of oriented marked exteriors). *Let Γ_1, Γ_2 be spatial graphs, (X_1°, P_1°) and (X_2°, P_2°) their oriented marked exteriors, and let $\Phi_X: (X_1^\circ, P_1^\circ) \rightarrow (X_2^\circ, P_2^\circ)$ be a PL homeomorphism preserving the orientation of both factors. Then Φ_X extends to an isomorphism $\Phi: \Gamma_1 \rightarrow \Gamma_2$.*

Before the proof, we remark that although one can encode links as spatial graphs (for example, with each link component consisting of one vertex and one loop), Proposition 6.3 does not contradict the well-known fact that there exist non-isotopic links with homeomorphic exteriors. Indeed, the difficulty in reconstructing a link from its exterior lies in determining the curve on each boundary component that should play the role of the meridian. If a link is realized as a spatial graph, however, each link component has at least one vertex, so in the marked exterior, the meridians are explicitly visible as the junctures. This is also the reason why Proposition 6.3 is quite distinct in spirit from Gordon-Luecke’s classical result that knots are determined by their exteriors [6].

Proof. Recall that the orientation of P_1° determines which regions of ∂X_1° are vertex regions. Thus, preserving the orientation of P_k° is equivalent to mapping vertex regions to vertex regions (and edge regions to edge regions).

Denote the regular neighbourhoods used in the construction of X_k by $N_{V,k}$ and $N_{E,k}$. The discs $N_{E,k} \cap \partial N_{V,k}$ are bounded by the junctures in P_k° , and each disc intersects $|\Gamma_k|$ in exactly one point in its interior. Since discs are PL-homeomorphic to cones over any of their interior points, we can use the cone construction to extend Φ_X to a PL homeomorphism $\Phi_{X+}: X_1^\circ \cup (N_{E,1} \cap \partial N_{V,1}) \rightarrow X_2^\circ \cup (N_{E,2} \cap \partial N_{V,2})$ mapping the intersection points of $|\Gamma_k|$ with $N_{E,k} \cap \partial N_{V,k}$ to each other.

Note that each $(N_{e,k}, e \cap X_{V,k})$ is an unknotted ball pair. As any PL homeomorphism of the boundary of an unknotted ball pair extends to the interior [15, Theorem 4.4], we can extend Φ_{X+} to a PL homeomorphism $\Phi_{X_V}: X_{V,1} \rightarrow X_{V,2}$ that maps $|\Gamma_1| \cap X_{V,1}$ to $|\Gamma_2| \cap X_{V,2}$.

Since each $N_{v,k}$ is PL-homeomorphic to a cone with base a 2-sphere containing $R_{v,k}$ and cone point v , such that $|\Gamma| \cap N_{v,k}$ corresponds to the cone over $|\Gamma| \cap \partial N_{v,k}$, we may cone the already defined map on each $\partial N_{v,1}$. This finishes the extension of Φ_X to the ambient sphere of Γ_1 . □

We have thus shown how to encode a spatial graph as its oriented marked exterior. Since our aim is to distinguish exteriors by applying the Recognition Theorem, which is insensitive to orientations, we need to refine our boundary pattern so that it encodes orientations, and also possibly decorations of the spatial graph.

Definition 6.4 ([14, Definition 3.3.9]). *A manifold with boundary pattern (M, P) is a PL 3-manifold M together with a 1-dimensional subpolyhedron $P \subset \partial M$ containing no isolated points. A **homeomorphism** of manifolds with boundary pattern is a PL homeomorphism of pairs $(M_1, P_1) \rightarrow (M_2, P_2)$.*

If we ignore orientations, then $(X_\Gamma^\circ, P_\Gamma^\circ)$ as defined above is an example of a manifold with boundary pattern. In the sequel we describe the required modification of P_Γ° . In order to avoid restricting spatial graph isomorphisms that can be detected by comparing marked exteriors, we need to ensure that all changes are independent of artificial choices.

First, we shall encode which regions of ∂X_Γ° are vertex regions, and the orientations of X_Γ° and P_Γ° . The orientation of X_Γ° can be recovered from the orientation of the junctures and the data of which regions of ∂X_Γ° are vertex regions. To encode the orientation of each juncture γ , choose three distinct points on γ . An orientation of γ is then the same as a cyclic ordering of these points. From one of the points p_1 extend one arc into the corresponding vertex region R_v . From the next point p_2 in the cyclic ordering, extend two arcs into R_v , and three arcs from the third point p_3 . With these added arcs, the boundary pattern encodes which regions are vertex regions as well as the orientations of junctures, hence also the orientation of X_Γ° .

To account for decorations of Γ we further modify this boundary pattern. To encode the colour $f(v)$ of a vertex v , rather than extending three arcs from p_3 , extend $f(v) + 3$ arcs instead. Similarly, to encode the colour $g(e)$ of the edge e at γ , extend $g(e)$ arcs from p_3 into the edge region R_e . We are left to encode the orientations of the edges. Note that there are two junctures around e , corresponding to its positive and negative ends (with respect to any PL embedding $[-1, 1] \rightarrow e \setminus V$ orienting e). If γ corresponds to the positive end of e , we extend a single arc from p_2 into R_e . After these modifications, the resulting boundary pattern still encodes which regions of ∂X_Γ° are vertex regions, as every vertex region receives arcs from three points in each of its junctures, and edge regions receive arcs from at most two. The number of arcs also clearly encodes the colourings f, g and/or the edge orientation.

To make the above construction rigorous, we define a (k, l, m) -**model disc**, where $k, l \in \mathbb{N}$ and $m \in \{0, 1\}$, as the standard oriented ball pair $(D, I) := ([-1, 1]^2, [-1, 1] \times \{0\})$ together with a few line segments in D emanating from points in I as in the left hand side of Figure 6.2. Explicitly, consider the points $p_1 := (-\frac{1}{2}, 0)$, $p_2 := (0, 0)$ and $p_3 := (\frac{1}{2}, 0)$ of I . From p_1 , extend one line segment into the upper region of D (which induces the orientation of I). From p_2 , extend two line segments into the upper region and m line segments into the lower region. From p_3 , extend $k + 3$ line segments into the upper region and l line segments into the lower region of D .




Figure 6.2. Left: a $(2, 4, 1)$ -model disc. Right: model disc for P_Γ at the region of an isolated vertex v with colour $f(v) = 1$. Arrows indicate orientations.




Figure 6.3. The marked exterior of a decorated spatial graph Γ . Numbers indicate the colours of the vertices and edge.

Choose a regular neighbourhood of P_Γ° in ∂X_Γ° , which is comprised of an annulus A_γ for each juncture γ . If γ bounds the vertex region R_v and the edge region R_e , choose an orientation-preserving PL embedding of the $(f(v), g(e), m)$ -model disc $\Phi: (D, I) \rightarrow (A_\gamma, \gamma)$, where $f: V \rightarrow \mathbb{N}$ and $g: E \rightarrow \mathbb{N}$ are the colourings and $m \in \{0, 1\}$ is 1 if and only if γ corresponds to the positive end of e . Then, add the images of the arcs in the model disc to P_Γ° ; see Figure 6.3 for an example.

We still have to take care of isolated vertices, whose colours have not yet been encoded in the boundary pattern (and if Γ has only isolated vertices, the boundary pattern does not encode the orientation of ∂X_Γ°). We use as a local model the disc on the right hand side of Figure 6.2, again with 1, 2 and $f(v) + 3$ lines extending from the triangle in the middle. We push forward this pattern to each region R_v corresponding to an isolated vertex v via any orientation-preserving PL embedding.

The boundary pattern P_Γ° together with the added patterns near the junctures and at vertex regions of isolated vertices makes up the new boundary pattern P_Γ .

Definition 6.5. A *marked exterior* of a decorated spatial graph Γ is a manifold with boundary pattern (X_Γ, P_Γ) , where $X_\Gamma = X_\Gamma^\circ$ as unoriented PL manifolds and P_Γ is obtained from P_Γ° as described above.

We stress that whereas for an oriented marked exterior $(X_\Gamma^\circ, P_\Gamma^\circ)$ both X_Γ° and P_Γ° are oriented, the manifold X_Γ is not oriented (and P_Γ is not even a manifold). Our construction was designed to allow for recovering the orientation data of $(X_\Gamma^\circ, P_\Gamma^\circ)$ from (X_Γ, P_Γ) .

Proposition 6.6 (Marked exteriors are well-defined). *Let Γ be a decorated spatial graph with two marked exteriors (X_k, P_k) , $k \in \{1, 2\}$. Then there exists a homeomorphism of manifolds with boundary pattern $\Phi: (X_1, P_1) \rightarrow (X_2, P_2)$ such that, for each vertex v of Γ , Φ sends the corresponding vertex regions $R_{v,1}$ to $R_{v,2}$ and similarly for edge regions.*

Proof. By Proposition 6.2, the oriented marked exteriors (X_k°, P_k°) used to construct the (X_k, P_k) are homeomorphic as manifolds with boundary pattern via a homeomorphism $\Phi: (X_1^\circ, P_1^\circ) \rightarrow (X_2^\circ, P_2^\circ)$ respecting vertex and edge regions of the boundary, as well as the orientation of each factor. By the Regular Neighbourhood Theorem [15, Theorem 3.24] there is a PL isotopy H_A of ∂X_2 fixing P_2° and pushing the Φ -image of the chosen regular neighbourhood of P_1° onto the chosen regular neighbourhood of P_2° . Denote the final homeomorphism of this isotopy by Ψ .

For each juncture $\gamma \subset P_1^\circ$, apply the Disc Theorem for pairs (Theorem 4.4) to get a PL isotopy H_γ of $A_{\Phi(\gamma)}$ that fixes $\partial A_{\Phi(\gamma)}$, preserves the juncture $\Phi(\gamma)$, and isotopes the postcomposition with $\Psi \circ \Phi$ of the embedding of the model disc at γ to the embedding of the model disc at $\Phi(\gamma)$. Using all the PL isotopies H_γ and extending to the whole boundary ∂X_2 as the identity, we obtain a PL isotopy H_D of ∂X_2 carrying $\Psi \circ \Phi(P_1)$ to P_2 .

The isotopy of ∂X_2 obtained by the concatenation of H_A and H_D extends to a PL isotopy of X_2 [15, Proposition 3.22(ii)]. Its final homeomorphism, when precomposed with Φ , yields a homeomorphism $(X_1, P_1) \rightarrow (X_2, P_2)$ respecting the vertex and edge regions. □

Proposition 6.7 (Faithfulness of marked exteriors). *Let (X_1, P_1) and (X_2, P_2) be marked exteriors for two non-empty decorated spatial graphs Γ_1 and Γ_2 , and $\Phi_X: (X_1, P_1) \rightarrow (X_2, P_2)$ a PL homeomorphism. Then Φ_X extends to an isomorphism $\Phi: \Gamma_1 \rightarrow \Gamma_2$.*

Proof. For $k \in \{1, 2\}$ consider the boundary patterns $P_k^\circ \subset P_k$ consisting only of the junctures. As an unoriented manifold, P_k° can be intrinsically characterized as the union of all embedded circles in P_k that are not the only circle in their component of ∂X_k . Hence, Φ_X maps P_1° to P_2° .

Observe that Φ_X preserves vertex regions and edge regions: indeed, components of ∂X_k with a single component of P_k correspond to isolated vertices and, on all other connected components, vertex regions are those receiving arcs from 3 distinct points of each juncture, while edge regions only receive arcs from at most two distinct points. As the number of arcs extended into the vertex region encodes the orientation of each juncture, Φ_X has to preserve the orientation of P_k° as well as the orientation of the triangles at isolated vertices. As the orientation of X_k° is determined by the orientation of the junctures and which side is the vertex region, Φ_X also preserves the orientation of the X_k° .

Overall, we conclude that Φ_X is a PL homeomorphism $(X_1^\circ, P_1^\circ) \rightarrow (X_2^\circ, P_2^\circ)$ preserving the orientation of each factor. Thus, Proposition 6.3 can be applied to conclude that, up to decorations, Φ_X extends to an isomorphism $\Phi: \Gamma_1 \rightarrow \Gamma_2$. From the way the decorations got encoded into the boundary patterns, it is clear that Φ respects decorations. □

6.2. Properties of the marked exterior

Here we explain how being non-split and having no cut vertices translates into features of marked exteriors.

Definition 6.8. *Let (M, P) be a manifold with boundary pattern.*

- *A properly embedded PL 2-sphere that does not bound a PL 3-ball in M is called a **reducing sphere**. We call M **reducible** if it admits a reducing sphere, and **irreducible** otherwise. We apply the same terminology to (M, P) .*
- *A subspace $X \subseteq M$ is called **clean** if $X \cap P = \emptyset$. Let $D \subset M$ be a clean properly embedded PL disc. If ∂D does not bound a clean disc in ∂M , then D is called a **reducing disc** for (M, P) . We say that (M, P) is **boundary-reducible** if it has a reducing disc; otherwise it is **boundary-irreducible**.⁴*

Proposition 6.9 (Splitting and reducibility). *Let Γ be a decorated spatial graph and (X_Γ, P_Γ) its marked exterior. Then Γ is split if and only if X_Γ is reducible.*

Proof. (\implies) Suppose S is a splitting sphere for Γ . Building X_Γ out of small enough regular neighbourhoods N_V and N_E as to avoid S , we see S is a reducing 2-sphere for X_Γ since no component of $X_\Gamma \setminus S$ is an open 3-ball, as both have non-empty boundary.

(\impliedby) Denote by \mathcal{S} the ambient 3-sphere of Γ , assume S is a reducing sphere for a marked exterior $X_\Gamma \subset \mathcal{S}$, and let $B_1, B_2 \subset \mathcal{S}$ be the 3-balls into which S splits \mathcal{S} . If for some $i \in \{1, 2\}$ the intersection $|\Gamma| \cap B_i$ were empty, then we would have $B_i \subset X_\Gamma$, in contradiction with S being a reducing sphere. Hence, if S decomposes Γ as $\Gamma_1 \sqcup \Gamma_2$, then none of the Γ_i is empty. □

The second part of the proof actually shows a finer statement:

Corollary 6.10 (Reducing spheres split). *Every reducing sphere for a marked exterior (X_Γ, P_Γ) is a splitting sphere for Γ .*

The relationship between cut vertices of a spatial graph and boundary-reducibility of its marked exterior is more subtle, so we study each direction of the correspondence separately, but the general idea is depicted in [Figure 6.4](#).

Proposition 6.11 (Boundary-reducibility from cut vertices). *Let Γ be a non-split decorated spatial graph. If Γ has a cut vertex, then its marked exterior (X_Γ, P_Γ) is boundary-reducible.*

⁴ The definition of boundary-irreducibility for manifolds with boundary pattern is given in Matveev’s book simply as the ‘straightforward generalization’ of the notion for 3-manifolds without boundary pattern [14, p. 126], leaving unclear whether the definition of a reducing disc D for (M, P) allows ∂D to bound a non-clean disc in ∂M . It is stated on p. 127 that, if M is a solid torus, then (M, P) is boundary-reducible if and only if $\partial M \setminus P$ contains a meridian of M . This is only true if ∂D is not allowed to bound a disc on ∂M , even when that disc intersects P , contrary to our definition. We believe this example is due to an oversight. Indeed, Matveev’s usage of the term (e.g., in the proofs of Lemmas 4.1.33 and 4.1.35) relies on the boundary of a clean properly embedded disc in a boundary-irreducible (M, P) bounding a clean disc in ∂M . This is compatible with the definition we present and with the usage elsewhere in the literature (e.g., [12, p. 16]).




Figure 6.4. A cut sphere S in Γ gives rise to a reducing disc D in (X_Γ, P_Γ) .

Proof. Let $\Gamma = (\mathcal{S}, V, E)$, let v be a cut vertex for Γ , and S a cut sphere through v . We construct a marked exterior (X_Γ, P_Γ) using a small enough regular neighbourhood N_V of V so that N_V is in fact a regular neighbourhood of V in $(\mathcal{S}, |\Gamma| \cup S)$, and also, we use N_E small enough to be disjoint from S . Additionally, we ensure that P_Γ is built from disc embeddings with image small enough to be disjoint from S , so that $S \cap P_\Gamma = \emptyset$. As $S \cap N_v$ is a regular neighbourhood of $\{v\}$ in S , it is a disc. The other side $D := S \cap X_\Gamma$ is thus a clean properly embedded disc in (X_Γ, P_Γ) .

We claim that D is a reducing disc for (X_Γ, P_Γ) . To see this, consider the two balls B_1, B_2 into which S separates \mathcal{S} . The curve ∂D separates the component C of ∂X_Γ containing the vertex region R_v into the two regions $C_i := C \cap B_i$, for $i \in \{1, 2\}$. We need to show that no C_i is a clean disc. Since S is a cut sphere, there is at least one edge e_i incident to v on each B_i . As the corresponding component N_{e_i} of N_E is disjoint from S , we have $N_{e_i} \subset B_i$, and in particular $R_{e_i} \subset C_i$. The juncture between R_{e_i} and R_v is thus contained in C_i , whence C_i is not clean. \square

A converse statement also holds, except for one particular (isomorphism type of) spatial graph: a spatial graph is called a **one-edge graph** if it has exactly two vertices and one edge, with the edge being incident to both vertices.

Lemma 6.12. (Uniqueness of one-edge graphs). *Let Λ_1, Λ_2 be decorated one-edge graphs, and let $F: \langle \Lambda_1 \rangle \rightarrow \langle \Lambda_2 \rangle$ be an isomorphism of their underlying decorated abstract graphs. Then there is an isomorphism $\Phi: \Lambda_1 \rightarrow \Lambda_2$ with $\langle \Phi \rangle = F$.*

This is a straightforward consequence of the following general statement, taking M to be a 3-sphere.

Proposition 6.13. (Arcs in the interior of connected manifolds). *Let M be a connected PL manifold of dimension at least 2. For $k \in \{1, 2\}$, let I_k be a PL-embedded arc in $\text{int}(M)$ with endpoints v_k, u_k . Then there is a PL isotopy of M carrying I_1 onto I_2 , v_1 to v_2 and u_1 to u_2 .*

It is not true in general that any two PL-embedded n -balls in the interior of a PL manifold of dimension at least $n + 1$ are ambient-isotopic. For example, consider the cone D of a trefoil knot in $\partial([-1, 1]^4)$, with the origin as cone point. Then D cannot be

ambiently isotoped in \mathbb{R}^4 onto the disc $[-1, 1]^2 \times \{0\}^2$. This follows from the fact that links of pairs of polyhedra are PL invariants [15, pp. 50–51].

Proof of Proposition 6.13. We need the following:

Claim. For every PL embedded arc $I \subset \text{int}(M)$ with endpoints v, u , and for every $u' \in I \setminus \{v\}$, there is a PL isotopy of M fixing v and carrying I to the sub-arc of I with endpoints v, u' .

Before justifying this claim we use it to prove the proposition. By homogeneity of manifolds [15, Lemma 3.33], there is a PL isotopy of M carrying v_1 to v_2 , so we may assume $v_1 = v_2 =: v$. Choose a star neighbourhood N_v of v in the pair (M, I) , and denote by u'_1, u'_2 , respectively, the points of intersection of the $(n - 1)$ -sphere ∂N_v with each arc I_1, I_2 . Using the above claim on both arcs reduces the problem to showing that there is a PL isotopy of M carrying the straight line segment $[v, u'_1]$ onto $[v, u'_2]$, with v being carried to itself.

Since ∂N_v is connected, again by homogeneity of manifolds, there is a PL isotopy of ∂N_v carrying u'_1 to u'_2 . By coning at v , this isotopy extends to N , taking $[v, u'_1]$ to $[v, u'_2]$ as required. To extend it to all of M , we use the general fact every PL isotopy of the boundary of a manifold (in this case $M \setminus \text{int}(N_v)$) extends to the interior [15, Proposition 3.22(ii)]. □

Proof of the Claim. It suffices to show that the subspace $Q \subset I \setminus \{v\}$ of points u' for which the claim holds is non-empty, open and closed in $I \setminus \{v\}$. Clearly $u \in Q$.

Let us verify that Q is open, beginning with u . A regular neighbourhood N_u of $\{u\}$ in the pair (M, I) is PL-homeomorphic to the standard n -ball $[-1, 1]^n$, with $N_u \cap I$ corresponding to the straight line segment from 0 to a point p_0 in $\partial([-1, 1]^n)$. Let q_0 be in the interior of this line segment. Since $[-1, 1]^n$ is a cone with base $\partial([-1, 1]^n)$ over any of its interior points, the formula $tp \mapsto (1 - t)q_0 + tp$ with $p \in \partial([-1, 1]^n)$ defines a PL homeomorphism of $[-1, 1]^n$ fixing the boundary and taking $[0, p_0]$ to $[q_0, p_0]$. By Alexander’s trick [15, Proposition 3.22(i)], such a map is PL-isotopic to the identity on $[-1, 1]^n$ keeping the boundary fixed. This isotopy can then be transferred to a PL isotopy of N_u and extended as the constant isotopy in all of M . This shows that the point $q \in N$ corresponding to q_0 is in Q , and so Q contains the half-open interval $\text{int}(N_u) \cap I$.

To verify the openness condition at points $u' \in Q \setminus \{u\}$, proceed similarly: choose a regular neighbourhood $N_{u'}$ of u' in (M, I) , and model $(N_{u'}, N_{u'} \cap I)$ as the standard ball pair $([-1, 1]^n, [-1, 1] \times \{0\}^{n-1})$. The previous construction shows that $\text{int}(N_{u'}) \cap I \subset Q$. The same argument shows Q is closed in $I \setminus \{v\}$. □

With the claim established, the proposition is proved.

Proposition 6.14 (Cut vertices from boundary-reducibility). *Let Γ be a non-split decorated spatial graph that is not a one-edge graph. If its marked exterior (X_Γ, P_Γ) is boundary-reducible, then Γ has a cut vertex. Moreover, there is an algorithm to produce a cut sphere for Γ from any reducing disc for (X_Γ, P_Γ) .*

The second statement is meant to be used in tandem with the fact that one can algorithmically find a reducing disc for (X_Γ, P_Γ) , given by Theorem 7.9 below. Thus, to be precise, one should interpret the input reducing disc to be specified as a normal surface in the triangulated X_Γ .

Proof. Let D be a reducing disc for a marked exterior (X_Γ, P_Γ) in the ambient sphere \mathcal{S} of Γ . Since ∂D is disjoint from P_Γ , it is contained in a vertex region R_v or an edge region R_e of ∂X_Γ .

Let us first treat the case where $\partial D \subset \partial R_v$. Consider the 3-ball N_v containing v , which is a component of the vertex set neighbourhood N_V used in constructing (X_Γ, P_Γ) . Since N_v is a regular neighbourhood of $\{v\}$ in $(\mathcal{S}, |\Gamma|)$, the pair $(N_v, N_v \cap |\Gamma|)$ is PL-homeomorphic to a cone of $(\partial N_v, \partial N_v \cap |\Gamma|)$, with v corresponding to the cone point. Let $D_v \subset N_v$ be the disc corresponding to the cone of ∂D , and consider the 2-sphere $S := D \cup D_v$, which intersects $|\Gamma|$ precisely at v . We claim that S is a cut sphere.

Let $B_1, B_2 \subset \mathcal{S}$ be the 3-balls into which S separates \mathcal{S} , let C be the component of ∂X_Γ containing ∂D , and consider the two surfaces $C_i = B_i \cap C$ into which ∂D cuts C . Since D is a reducing disc, none of the C_i is a clean disc. This implies that there are edges of Γ incident to v on both sides of S , and so none of the summands in the decomposition $\Gamma = \Gamma_1 \bullet_v \Gamma_2$ induced by S is a one-point graph. Hence S is a cut sphere for Γ , and v a cut vertex.

Now we treat the case where $\partial D \subset R_e$ for some edge e of Γ . Observe that some vertex incident to e has degree at least 2: for otherwise the component of ∂X_Γ containing R_e would be a 2-sphere in \mathcal{S} with only the edge e in one of its sides, and no other vertices besides its endpoints. Since Γ is non-split, Γ would be a one-edge graph, contrary to assumption. So let v be a vertex incident to e of degree at least 2.

As no component of P_Γ is contained in R_e , and ∂D , being a reducing disc, does not bound a clean disc in R_e , we conclude ∂D does not bound a disc in R_e . It thus cuts R_e into two annuli. Let R'_e be one such annulus having one of its boundary components in R_v , and consider the enlarged disc $D' := D \cup R'_e$. Its boundary $\partial D'$ is contained P_Γ , being the juncture between R_e and R_v . As before, let D_v be the disc obtained by coning $\partial D'$ at v , and define $S := D' \cup D_v$.

We now show S is a cut sphere for Γ . Clearly, $S \cap |\Gamma| = \{v\}$. From the description of $S \cap N_v$ as a cone of the juncture between R_e and R_v , we see that one side of S contains e , and the other side contains all other edges of Γ that are incident to v . Since v has degree at least two, it follows that there are edges incident to v on both sides of S . Hence S induces a non-trivial vertex sum decomposition of Γ . □

We finish this section by discussing the relation between degree-1 vertices in a spatial graph and one-edge graphs. Uniqueness of one-edge graphs yields two propositions about spatial forests, which in turn imply Theorem 1.3.

Lemma 6.15 (One-edge graph summands from leaves). *Let $\Gamma = (\mathcal{S}, V, E)$ be a decorated spatial graph, u a leaf of Γ , e the edge incident to u , and v be the other vertex incident to e . For the sub-graph $\Gamma_0 := (\mathcal{S}, V \setminus \{u\}, E \setminus \{e\})$ and the one-edge sub-graph $\Lambda := (\mathcal{S}, \{u, v\}, \{e\})$, one has $\Gamma = \Gamma_0 \bullet_v \Lambda$.*

Proof. By Lemma 4.17, we need only find a 2-sphere S intersecting $|\Gamma|$ exactly at v such that $|\Gamma_0|$ is in one side of S , and $|\Lambda|$ is in the other.

Let $(X_\Gamma^\circ, P_\Gamma^\circ)$ be an oriented marked exterior for Γ , and denote by γ the juncture between the regions R_e, R_v of ∂X_Γ° . The component N_v of the vertex set neighbourhood used in constructing X_Γ° is PL-homeomorphic to a cone of the pair $(\partial N_v, \partial N_v \cap |\Gamma|)$, with v corresponding to the cone point. We denote by D the disc properly embedded in N_v

that corresponds to the cone of γ , and by C the β -ball that corresponds to the cone over the disc $N_v \cap N_e$. Recalling that R_e is a cylinder and that u is a leaf, the region R_u of ∂X_Γ° corresponding to u is a disc. We claim the 2-sphere $S := R_u \cup R_e \cup D$ decomposes Γ as $\Gamma_0 \bullet_v \Lambda$. Indeed, it is clear that $S \cap |\Gamma| = \{v\}$. Moreover, one of the sides of S is the β -ball $N_u \cup N_e \cup C$, which intersects $|\Gamma|$ precisely at $|\Lambda|$. The other side must then contain $|\Gamma_0|$. □

A **spatial forest** is a spatial graph whose underlying graph is a forest. For spatial forests, the isomorphism problem is reduced to a search for an isomorphism of the underlying graphs, bypassing the machinery in Matveev’s book:

Theorem 6.16 (Uniqueness of spatial forests). *Let Γ, Γ' be decorated spatial forests, and let $F: \langle \Gamma \rangle \rightarrow \langle \Gamma' \rangle$ be an isomorphism of their underlying decorated graphs. Then there is an isomorphism $\Phi: \Gamma \rightarrow \Gamma'$ with $\langle \Phi \rangle = F$.*

Proof. We induct on the number of vertices of Γ . The case $\Gamma \cong \mathbf{0}$ is trivial.

Suppose first that Γ has an isolated vertex v , denote by Λ the one-point sub-graph consisting of v , and by Γ_0 the sub-graph of Γ obtained by suppressing v . Since v is isolated, one has $\Gamma = \Gamma_0 \sqcup \Lambda$ (e.g., by Lemma 3.8). As isolated vertices are determined by their underlying graphs, the vertex $F(v)$ of Γ' is also isolated and we have a similar decomposition $\Gamma' = \Gamma'_0 \sqcup \Lambda'$, with $\langle \Gamma'_0 \rangle = F(\langle \Gamma_0 \rangle)$. By induction, there is an isomorphism $\Phi_0: \Gamma_0 \rightarrow \Gamma'_0$ inducing $F|_{\Gamma_0}$. The fact that all one-point graphs (of the same colour) are isomorphic then yields an isomorphism $\Phi_\Lambda: \Lambda \rightarrow \Lambda'$. By Lemma 3.6 these combine to an isomorphism $\Phi := \Phi_0 \sqcup \Phi_\Lambda$ inducing F .

Assume now that each component of $\langle \Gamma \rangle$ is a tree with at least two vertices. By a standard graph-theoretic argument (see, e.g., [9, Exercise 1.2.5]), finite trees with at least two vertices always have leaves, so $\langle \Gamma \rangle$, and thus also Γ , has a leaf u . Denote by e the edge of Γ incident to u , and by v the other vertex incident to e . Let Λ be the one-edge sub-graph of Γ comprised of u, v and e , and Γ_0 be the sub-graph of Γ obtained by excluding e and u . By Lemma 6.15, we have $\Gamma = \Gamma_0 \bullet_v \Lambda$. Similarly, let Γ'_0 be the sub-graph of Γ' obtained by excluding the edge $F(e)$ and the leaf $F(u)$, and let Λ' be the one-edge sub-graph of Γ' comprised of $F(u), F(v)$ and $F(e)$. As before, we have $\Gamma' = \Gamma'_0 \bullet_{F(v)} \Lambda'$. By induction, the isomorphism $F|_{\Gamma_0}: \langle \Gamma_0 \rangle \rightarrow \langle \Gamma'_0 \rangle$ is induced by an isomorphism $\Phi_0: \Gamma_0 \rightarrow \Gamma'_0$. On the other hand, Lemma 6.12 gives an isomorphism $\Phi_\Lambda: \Lambda \rightarrow \Lambda'$ inducing $F|_\Lambda: \langle \Lambda \rangle \rightarrow \langle \Lambda' \rangle$. By Lemma 4.9 these assemble to the desired $\Phi := \Phi_0 \bullet_v \Phi_\Lambda$. □

Spatial forests are fully characterized by Theorem 6.16, as the following shows:

Proposition 6.17 (Non-uniqueness of non-forests). *Suppose Γ is a spatial graph that is not a forest. Then there exists a spatial graph $\Gamma' \not\cong \Gamma$ such that $\langle \Gamma' \rangle \cong \langle \Gamma \rangle$.*

Proof. A **circuit** of Γ is a sub-graph whose support is a PL circle. If \mathcal{S} is the ambient sphere of Γ , each of the (finitely many) circuits of Γ is a knot in \mathcal{S} . As there is, up to PL isotopy, a unique orientation-preserving homeomorphism $\mathcal{S} \rightarrow \mathbb{S}^3$ to the standard 3-sphere, one can consider the finite set \mathcal{K}_Γ of equivalence classes of knots in \mathbb{S}^3 represented by some circuit of Γ . The set \mathcal{K}_Γ is then an invariant of Γ under isomorphism.

Take P to be a prime knot that is not a connect-summand of any knot in \mathcal{K}_Γ (which exists as there are infinitely many prime knots [11]), and choose an edge e of Γ that

is part of some circuit. Let $N \subset \mathcal{S}$ be a PL 3-ball intersecting e in such a way that $(N, N \cap e)$ is an unknotted ball pair, but with N otherwise disjoint from $|\Gamma|$ (for example, we may take $N = N_e$ as defined in §6.1). Let Γ' be obtained from Γ by performing a surgery of P with e inside N . Then $\langle \Gamma' \rangle \cong \langle \Gamma \rangle$, but P is now a connect-summand of the support of every circuit of Γ containing e . In particular, $\mathcal{K}_\Gamma \not\cong \mathcal{K}_{\Gamma'}$, whence $\Gamma \not\cong \Gamma'$. \square

7. Algorithmic theory of spatial graphs

We now import the final pieces of terminology needed to state the Recognition Theorem and assemble the theory developed so far into a proof of Theorem 1.1.

Given a compact PL 3-manifold M and a properly PL-embedded surface $\Sigma \subset M$, recall that Σ is **incompressible** if for every PL-embedded disc $D \subset M$ such that $D \cap \Sigma = \partial D$, the boundary ∂D bounds a disc in Σ . We say Σ is **two-sided** if its normal bundle is trivial [14, p. 124].

Definition 7.1 ([14, Definition 4.1.20]). *A compact PL 3-manifold M is **sufficiently large** if there exists a PL-embedded closed connected surface $\Sigma \subset M$ that is incompressible, two-sided, and not a 2-sphere or a real projective plane.*

Definition 7.2 ([14, Definition 6.1.5]). *A manifold with boundary pattern (M, P) with M, P compact is called **Haken** if it is irreducible, boundary-irreducible and either:*

- M is sufficiently large, or
- $P \neq \emptyset$ and M is a handlebody of positive genus.

Proposition 7.3 (Non-triviality of the boundary [14, Corollary 4.1.27]). *Every irreducible PL 3-manifold with non-empty boundary is either a handlebody or sufficiently large.*

Theorem 7.4 (Haken–Matveev Recognition Theorem [14, Theorem 6.1.6]). *There is an algorithm to decide whether two given Haken 3-manifolds with boundary pattern are PL-homeomorphic (as manifolds with boundary pattern).*

We shall apply the Recognition Theorem to marked exteriors of blocks. To ensure the ‘Haken’ condition is met, we need Proposition 7.3. However, that proposition leaves room for the exterior of a block to be a genus-0 handlebody. We control this case with the following:

Lemma 7.5 (Blocks with 3-ball exteriors). *Let Λ be a block with marked exterior (X_Λ, P_Λ) . Then X_Λ is a 3-ball if and only if Λ is a one-edge graph.*

Proof. (\Leftarrow) Clearly for *some* one-edge graph, the marked exterior is a 3-ball. But by Lemma 6.12 all one-edge graphs are isomorphic.

(\Rightarrow) Suppose X_Λ is a 3-ball and consider the oriented marked exterior $(X_\Lambda^\circ, P_\Lambda^\circ)$. Clearly X_Λ° is also a 3-ball. Note that P_Λ° is a collection of circles (non-empty, since Λ has at least one edge). The disc bounded by an innermost such circle is then the region R_u corresponding to some leaf u . By Lemma 6.15, we obtain a vertex sum decomposition

$\Lambda = \Lambda_0 \bullet_v \Lambda'$, where Λ' is the one-edge sub-graph containing u . Since Λ is a block, it follows that $\Lambda_0 \cong \mathbf{1}$ and so $\Lambda = \Lambda'$. □

Proposition 7.6 (Algorithmic recognition of blocks). *There is an algorithm to decide whether two blocks with decorations of the same type are isomorphic.*

Proof. For $k \in \{1, 2\}$, let Λ_k be the decorated blocks we wish to compare. First check whether the underlying graphs $\langle \Lambda_k \rangle$ are one-edge graphs. If exactly one of them is so, then $\Lambda_1 \not\cong \Lambda_2$. If both Λ_k are one-edge graphs, Lemma 6.12 reduces the problem to testing whether $\langle \Lambda_1 \rangle \cong \langle \Lambda_2 \rangle$, which is straightforward.

So suppose none of the Λ_k is a one-edge graph, construct marked exteriors (X_k, P_k) , and note that they are Haken:

- The X_k are irreducible by the ‘if’ direction in Proposition 6.9.
- The (X_k, P_k) are boundary-irreducible by Proposition 6.14.
- Since $\partial X_k \neq \emptyset$, Proposition 7.3 tells us they are either sufficiently large or handlebodies. In the latter case, genus-0 is excluded by the ‘only if’ direction in Lemma 7.5. Since blocks have edges, the condition $P_k \neq \emptyset$ is satisfied.

We can thus apply Theorem 7.4 to test whether the (X_k, P_k) are homeomorphic. By Propositions 6.6 and 6.7, this is equivalent to the Λ_k being isomorphic. □

We actually make use of a refined version of Proposition 7.6:

Lemma 7.7 (Algorithmic recognition of multi-pointed blocks). *There is an algorithm that takes as input*

- two blocks Λ_1, Λ_2 with decorations of the same type,
- a tuple (v_1^1, \dots, v_r^1) of distinct vertices of Λ_1 , and
- a tuple (v_2^1, \dots, v_r^2) of distinct vertices of Λ_2 of the same size,

and decides whether there is an isomorphism $\Phi: \Lambda_1 \rightarrow \Lambda_2$ with $\Phi(v_1^l) = v_2^l$ for every $l \in \{1, \dots, r\}$.

Proof. Since having no vertex colouring is the same as having a vertex colouring where all vertices are 0-coloured, we may assume that vertex colourings are part of the decorations.

Let $n \in \mathbb{N}$ be such that both f_k have range contained in $\{0, \dots, n - 1\}$, and let Λ_k^+ be the same block as Λ_k , but with the modified vertex colouring:

$$f_k^+(v) = \begin{cases} f_k(v) & \text{if } v \text{ is not one of the } v_k^l, \\ nl + f_k(v) & \text{if } v = v_k^l. \end{cases}$$

The colouring f_k^+ encodes, for each vertex v , the original colouring $f_k(v)$ as the mod- n residue. Moreover, the division with remainder of $f_k^+(v)$ by n returns l if v is one of the v_k^l , and otherwise it returns 0.

Thus, finding an isomorphism $\Lambda_1 \rightarrow \Lambda_2$ as in the statement is equivalent to finding an isomorphism $\Lambda_1^+ \rightarrow \Lambda_2^+$, which can be algorithmically determined by Proposition 7.6. \square

We can now bootstrap our algorithm for recognition of spatial graphs. We develop it first for pieces (Proposition 7.10) and then in full generality (Theorem 7.13). We shall make no further explicit usage of the Recognition Theorem.

Lemma 7.8 (Algorithmic decomposition into a tree of blocks). *There is an algorithm that, given a non-split spatial graph $\Gamma \not\cong \mathbf{1}$, produces a tree of blocks \mathcal{T} such that $\Gamma = [\mathcal{T}]$.*

For the proof we need the following:

Theorem 7.9 (Algorithmic detection of boundary-reducibility [14, Theorem 4.1.13]). *There exists an algorithm to decide whether a given irreducible 3-manifold with boundary pattern is boundary-irreducible. In case it is boundary-reducible, the algorithm constructs a reducing disc.*

Proof of Lemma 7.8. The proof mimics that of Proposition 4.21, but when constructing the tree of blocks \mathcal{T} we must ensure all steps can be done algorithmically.

We induct on the number of edges of Γ . The case $\Gamma = \mathbf{0}$ is trivial and the case where Γ is a one-point graph is excluded by assumption. Otherwise, we need to determine whether Γ is a block and, in case it is not, we need to find a cut sphere for Γ . To that end, construct the marked exterior (X_Γ, P_Γ) . By Proposition 6.9, X_Γ is irreducible and we can apply Theorem 7.9 to check whether (X_Γ, P_Γ) is boundary-reducible. If it is not, it follows from Proposition 6.11 that Γ has no cut vertices, hence Γ is itself a block and we are done.

In case (X_Γ, P_Γ) is boundary-reducible, Theorem 7.9 assures we can algorithmically construct a reducing disc, which Proposition 6.14 converts into a cut sphere S . The two vertex-summands in the induced decomposition $\Gamma = \Gamma_1 \bullet_v \Gamma_2$ are the sub-graphs of Γ supported on each side of S . The induction hypothesis applies, yielding algorithmically constructed trees of blocks for them, which can be (algorithmically) assembled into \mathcal{T} as in the proof of Proposition 4.21. \square

Proposition 7.10 (Algorithmic recognition of pieces). *There is an algorithm to decide whether two pieces with decorations of the same type are isomorphic.*

Proof. Let Γ_1, Γ_2 be the decorated pieces to be compared. The case where some Γ_k is a one-point graph is trivial.

Use Lemma 7.8 to algorithmically decompose Γ_k as the realization of a tree of blocks $\mathcal{T}_k = (T_k, I_k, J_k, L_k, (\Lambda_i)_{i \in I_k}, (v(l))_{l \in L_k})$. Then list the isomorphisms of abstract trees $f: T_1 \rightarrow T_2$ satisfying $f(I_1) = I_2$, which is a finite combinatorial problem. If no such isomorphism exists, then $\Gamma_1 \not\cong \Gamma_2$ by Proposition 4.22. What is more, if $\Gamma_1 \cong \Gamma_2$, then for some such f there is a family of isomorphisms $(\Phi_i: \Lambda_i \rightarrow \Lambda_{f(i)})_{i \in I_1}$ compatible with the assignments $l \mapsto v(l)$.

For each f in our list, use the algorithm of Lemma 7.7 at every $i \in I_1$ to determine whether there is an isomorphism $\Phi_i: \Lambda_i \rightarrow \Lambda_{f(i)}$ mapping the tuple $(v(l))_l$ to $(v(f(l)))_l$, where l ranges over the edges in L_1 incident to i . If for some $f: T_1 \rightarrow T_2$ we find such

Φ_i , they assemble into an isomorphism $\Phi: \Gamma_1 \rightarrow \Gamma_2$ (see Lemma 4.14). Otherwise we conclude $\Gamma_1 \not\cong \Gamma_2$, again by Proposition 4.22. \square

Lemma 7.11 (Algorithmic decomposition as the disjoint union of pieces). *There is an algorithm that, given a decorated spatial graph Γ , produces a finite collection of pieces $(\Lambda_i)_{i \in I}$ such that $\Gamma = \bigsqcup_{i \in I} \Lambda_i$.*

The lemma above requires the following:

Theorem 7.12 (Algorithmic detection of reducibility [14, pp. 161–162]). *There exists an algorithm to decide whether a given compact PL 3-manifold is irreducible. In case it is reducible, the algorithm constructs a reducing sphere.*

Proof of Lemma 7.11. Construct a marked exterior (X_Γ, P_Γ) and use the algorithm of Theorem 7.12 to test whether X_Γ is reducible. If it is not, Proposition 6.9 implies that Γ is non-split. Here, either $\Gamma = \mathbf{0}$, in which case we take $I = \emptyset$, or Γ is a piece, so we can take I to be a singleton.

If X_Γ is reducible, the algorithm of Theorem 7.12 produces a reducing sphere S for X_Γ , which by Corollary 6.10 is a splitting sphere for Γ . Hence, $\Gamma = \Gamma_1 \sqcup \Gamma_2$, where the disjoint union summands are the non-empty graphs supported on each side of S . Since Γ_1, Γ_2 both have strictly fewer vertices than Γ , we may assume by induction that Γ_1, Γ_2 are algorithmically decomposable into pieces, and these decompositions assemble into one for Γ . \square

Finally, we obtain our main result in full generality:

Theorem 7.13 (Algorithmic recognition of spatial graphs). *There is an algorithm to decide whether two spatial graphs with decorations of the same type are isomorphic.*

Proof. Let Γ_1, Γ_2 be the decorated graphs to be compared. For each $k \in \{1, 2\}$, use Lemma 7.11 to decompose Γ_k as a disjoint union of pieces $\bigsqcup_{i \in I_k} \Lambda_i$. By Proposition 3.11, if the indexing sets I_k have different cardinalities, then $\Gamma_1 \not\cong \Gamma_2$. On the other hand, should an isomorphism $\Gamma_1 \rightarrow \Gamma_2$ exist, there should be a bijection $f: I_1 \rightarrow I_2$ such that, for every $i \in I_1$, there is an isomorphism $\Phi_i: \Lambda_i \rightarrow \Lambda_{f(i)}$.

We thus run through all bijections $f: I_1 \rightarrow I_2$ and apply, for each one, the algorithm of Proposition 7.10 to test whether every Λ_i is isomorphic to $\Lambda_{f(i)}$. If no f has such a compatible family of isomorphisms, then $\Gamma_1 \not\cong \Gamma_2$. In case some f does admit a suitable family $(\Phi_i: \Lambda_i \rightarrow \Lambda_{f(i)})_{i \in I_1}$, an iterated application of Lemma 3.6 allows us to assemble the Φ_i into an isomorphism $\Phi: \Gamma_1 \rightarrow \Gamma_2$. \square

Acknowledgements. We are grateful to Benjamin Ruppik and Claudius Zibrowius for helpful conversations, to Sofia Amontova for bringing our attention to some relevant literature, and to David Gabai for his comments regarding attribution of the Recognition Theorem. We also thank an anonymous referee, whose comments have improved the quality of the text.

Funding statement. The first, second and third authors were supported by the CRC 1085 *Higher Invariants* (Universität Regensburg, funded by the DFG).

References

- (1) J. W. Alexander, An example of a simply connected surface bounding a region which is not simply connected, *Proc. Natl. Acad. Sci. USA* **10**(1) (1924), 8–10. doi: 10.1073/pnas.10.1.8.
- (2) J. W. Alexander, On the subdivision of 3-space by a polyhedron, *Proc. Natl. Acad. Sci. USA* **10**(1) (1924), 6–8. doi: 10.1073/pnas.10.1.6.
- (3) S. Barthel, There exist no minimally knotted planar spatial graphs on the torus, *J. Knot Theory Ramif.* **24**(7) (2015), Id/No 1550035. doi: 10.1142/S0218216515500352.
- (4) S. Friedl and G. Herrmann, Spatial graphs, In: *Encyclopedia of Knot Theory*, Ed. by C. Adams, E. Flapan, A. Henrich, L. H. Kauffman, L. D. Ludwig, and S. Nelson, Chap. 49, pp. 461–466 (New York: CRC Press, 2021). doi: 10.1201/9781138298217
- (5) S. Friedl and G. Herrmann, Graphical neighborhoods of spatial graphs, In: *2019-20 MATRIX Annals*, Ed. by D. R. Wood, J. de Gier, C. E. Praeger and T. Tao, Vol. 4, MATRIX Book Series, pp. 627–646 (Springer, Cham, 2021). doi: 10.1007/978-3-030-62497-2_38.
- (6) C. McA. Gordon and J. Luecke, Knots are determined by their complements, *J. Amer. Math. Soc.* **2**(2) (1989), 371–415. doi: 10.2307/301990979.
- (7) W. Haken, Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I, *Math. Z.* **80** (1962), 89–120. doi: 10.1007/BF01162369.
- (8) K. Johannson, *Homotopy equivalences of 3-manifolds with boundaries*, Vol. 761, Lecture Notes in Mathematics, pp. ii+303 (Springer, Berlin, 1979). doi: 10.1007/BFb0085406.
- (9) D. Jungnickel, *Graphs, Networks and Algorithms*, 2nd edn., Vol. 5, Algorithms and Computation in Mathematics (Springer-Verlag, Berlin, 2005). doi: 10.1007/978-3-642-32278-5.
- (10) L. H. Kauffman, Invariants of graphs in three-space, *Trans. Am. Math. Soc.* **311**(2) (1989), 697–710. doi: 10.1090/S0002-9947-1989-0946218-0.
- (11) L. H. Kauffman and P. Lopes, Infinitely many prime knots with the same Alexander invariants, *J. Knot Theory Ramifications* **26**(9) (2017), Id/No 1743009, p. 13. doi: 10.1142/S021821651743009X.
- (12) L. H. Kauffman and V. O. Manturov, Virtual knots and links, *Tr. Mat. Inst. Steklova* **252**(Geom. Topol., Diskret. Geom. i Teor. Mnoz.) (2006), 114–133. doi: 10.1134/s0081543806010111.
- (13) W. B. R. Lickorish, Simplicial moves on complexes and manifolds, In: *Proceedings of the Kirbyfest (Berkeley, CA, 1998)*, Vol. 2, Geometry & Topology Monographs, 299–320 (Geometry and Topology Publications, Coventry, 1999). doi: 10.2140/gtm.1999.2.299.
- (14) S. Matveev, *Algorithmic Topology and Classification of 3-Manifolds*, 2nd edn., Vol. 9, Algorithms and Computation in Mathematics (Springer, Berlin, 2007). doi: 10.1007/978-3-540-45899-9.
- (15) C. P. Rourke and B. J. Sanderson, *Introduction to Piecewise-Linear Topology*, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 69 (Springer-Verlag, New York-Heidelberg, 1972). doi: 10.1007/978-3-642-81735-9.
- (16) S. Suzuki, On linear graphs in 3-sphere, *Osaka Math. J.* **7** (1970), 375–396. <http://projecteuclid.org/euclid.ojm/1200692930>.
- (17) S. Yamada, An invariant of spatial graphs, *J. Graph Theory* **13**(5) (1989), 537–551. doi: 10.1002/jgt.3190130503.