PUBLICATIONS DE L'INSTITUT MATHÉMATIQUE
Nouvelle série tome 66(80) (1999), 23–45

TOPOLOGY PRESERVING EDGE CONTRACTION

TAMAL K. DEY, HERBERT EDELSBRUNNER,
SUMANTA GUHA, AND DMITRY V. NEKHAYEV

ABSTRACT. We study edge contractions in simplicial complexes and local conditions under which they preserve the topological type. The conditions are based on a generalized notion of boundary, which lends itself to defining a nested hierarchy of triangulable spaces measuring the distance to being a manifold.

1. INTRODUCTION

This paper studies the operation of shrinking or contracting an edge in a simplicial complex. The repeated application of this operation eventually reduces any connected complex to a single vertex. During that process, the complex loses all non-trivial topological properties. We are interested in recognizing edges that can be contracted without changing the topological type. The repeated application of such contractions simplifies the complex while preserving its type. This means there is a homeomorphism that connects the underlying space of the original with that of the the simplified complex, and we are also interested in constructing such a homeomorphism.

Motivation. Edge contractions are used in computer graphics to simplify surfaces for fast rendering. A surface consists of triangles in $\mathbb{R}^3$ connected to each other along shared edges and vertices. In mathematical language it is a 2-complex, and papers in computer graphics generally restrict themselves to 2-manifolds with or without boundary. The pioneering publication in this context is Hoppe et al. [5], but see also Garland and Heckbert [3] for an effective numerical prioritization of edge contractions.

1991 *Mathematics Subject Classification.* 57N80, 57Q15, 68Q25, 68R05.

Research by the first two authors is partially supported by the U.S. Department of Energy, under Subcontract B341494. Research by the second author is also supported by NSF under grant CCR-96-19542, and by ARO under grant DAAG55-98-1-0177. Research by the third author is also supported by NSF under grant CCR-97-10771.

A related topic is parametrization, where a surface is to be covered by possibly few and large patches, each the homeomorphic image of an open disk in $\mathbb{R}^2$. The first workable solution is described in a recent paper by Lee et al. [6]. The algorithm simplifies the surface incrementally and maintains a piecewise linear homeomorphism. The simplification is generated through repeated vertex removal, but the same algorithmic idea can also be based on edge contraction [2]. For this to work it is essential that each simplification step preserves the topological type.

Summary of Results. Edge contractions that preserve the topological type are recognized by local criteria phrased in terms of a generalized concept of boundary. Roughly but not exactly, the order of a simplex σ in a simplicial complex is the smallest integer i such that the underlying space of the star of σ is homeomorphic to $\mathbb{R}^h \times \mathbb{X}$, for some topological space $\mathbb{X}$ of dimension i. We could let $\mathbb{X}$ equal to the underlying space of the star and let h be zero, so the order is well-defined and at most the dimension of the star. The j-th boundary of a simplicial complex consists of all simplices of order $i \geq j$. We introduce the Link Conditions for an edge ab considered for contraction. They require that within each boundary the link of ab is equal to the intersection of the links of a and of b. More precisely, for each j the relation between the links of a, b, ab is required within the j-th boundary extended by cones connecting the $(j+1)$-st boundary to a dummy vertex. For 2-complexes we prove that the Link Conditions characterize edge contractions that permit a homeomorphic modification limited to the stars of a and b. We prove the same for 3-manifolds. For general 3-complexes we only prove that the Link Conditions imply the existence of homeomorphic modifications.

Outline. Section 2 introduces definitions from combinatorial topology. Section 3 defines boundary; its basic properties are established in Appendix A. Section 4 introduces edge contractions and unfoldings. Sections 5 and 6 prove the mentioned results for contractions in 2-complexes and in 3-complexes. Section 7 concludes the paper.

2. Basic Definitions

We use concepts and terminology from combinatorial topology and discuss simplicial complexes, topological spaces, and maps between spaces. Most but not all definitions are standard, and the standard ones can also be found in textbooks such as Munkres [9].

Simplicial complexes. A *k-simplex*, σ, is the convex hull of $k+1 \geq 1$ affinely independent points. Its *dimension* is $\dim \sigma = k$. A *face* of σ is a simplex, τ, defined by a non-empty subset of the $k+1$ points, and τ is *proper* if the subset is proper. We call σ a *coface* of τ and write $\tau \leq \sigma$. The *interior*, $\operatorname{int} \sigma$, is the set of points contained in σ but not in any proper face of σ.

A *simplicial complex*, K, is a finite collection of simplices so $\sigma \in K$ and $\tau \leq \sigma$ implies $\tau \in K$, and $\sigma, \sigma' \in K$ implies $\sigma \cap \sigma'$ is either empty of a face of both. All complexes in this paper are simplicial. The *dimension* of K is $\dim K = \max\{\dim \sigma \mid$

$\sigma \in K\}$. A *d-complex* is a simplicial complex of dimension d. A *subcomplex* is a simplicial complex $L \subseteq K$. A simplex in K is *principal* if it has no coface in K. The *vertex set* contains all 0-simplices: $\mathrm{Vert}\, K = \{\sigma \in K \mid \dim \sigma = 0\}$. The *underlying space* is the union of simplex interiors: $|\, K \,| = \bigcup_{\sigma \in K} \mathrm{int}\, \sigma$.

Let B be a subset of K that is not necessarily a subcomplex. The *closure* of B is the set of all faces of simplices in B. The *star* of B is the set of all cofaces of simplices in B. The *link* of B is the set of all faces of cofaces of simplices in B that are disjoint from simplices in B. In more formal notation,

$$\begin{aligned} \overline{B} &= \{\tau \in K \mid \tau \leq \sigma \in B\}, \\ \mathrm{St}\, B &= \{\sigma \in K \mid \sigma \geq \tau \in B\}, \\ \mathrm{Lk}\, B &= \overline{\mathrm{St}\, B} - \mathrm{St}\, \overline{B}. \end{aligned}$$

The closure is the smallest subcomplex that contains B. The link is always a complex while the star is generally not a complex. The closure of the star is always a complex and denoted as $\overline{\mathrm{St}}\, B = \overline{\mathrm{St}\, B}$. The concepts of dimension and of underlying space extend immediately to subsets of a complex: $\dim B = \dim \overline{B}$ and $|\, B \,| = \bigcup_{\sigma \in B} \mathrm{int}\, \sigma$.

We introduce operations that create new simplices and complexes from old ones. The *cone* from a point x to a simplex σ is defined if x is not an affine combination of the vertices of σ, and in this case it is the simplex $x \cdot \sigma = \mathrm{conv}\,(x \cup \sigma)$ of dimension $\dim \sigma + 1$. A *subdivision* of K is a complex $\mathrm{Sd}\, K$ so $|\, K \,| = |\, \mathrm{Sd}\, K \,|$ and every simplex in $\mathrm{Sd}\, K$ is contained in a simplex in K. Subdivisions can be created by various operations. One such operation is *starring* from a point $x \in |\, K \,|$: remove all simplices that contain x and add x together with the cones from x to the faces of the removed simplices that do not contain x.

Topological spaces. A d-dimensional *point* is a d-tuple of real numbers. The *norm* of a point $x = (x_1, x_2, \ldots, x_d)$ is $\|x\| = (\sum x_i^2)^{1/2}$. The d-dimensional *Euclidean space*, $\mathbb{R}^d$, is the set of d-dimensional points together with the Euclidean distance function that maps each pair of points x, y to the non-negative real $\|x - y\|$. In addition to $\mathbb{R}^d$ we need names for three other standard topological spaces: the $(d-1)$*-sphere*, the *d-ball*, and the *d-halfspace*:

$$\begin{aligned} \mathbb{S}^{d-1} &= \{x \in \mathbb{R}^d \mid \|x\| = 1\}, \\ \mathbb{B}^d &= \{x \in \mathbb{R}^d \mid \|x\| \leq 1\}, \\ \mathbb{H}^d &= \{x \in \mathbb{R}^d \mid x_1 \geq 0\}. \end{aligned}$$

Most spaces in this paper are underlying spaces of complexes K. The space $|\, K \,|$ is subset of some Euclidean space $\mathbb{R}^e$, and it is equipped with the subspace topology inherited from the Euclidean topology of $\mathbb{R}^e$.

A *d-manifold* is a non-empty topological space, $\mathbb{M}$, so every point $x \in \mathbb{M}$ has an open neighborhood homeomorphic to $\mathbb{R}^d$. For complexes it suffices to check the defining condition at the vertices: $|\, K \,|$ is a d-manifold iff the underlying space of every vertex star is homeomorphic to $\mathbb{R}^d$. A *d-manifold with boundary* is a non-empty topological space, $\mathbb{N}$, so every $x \in \mathbb{N}$ has an open neighborhood homeomorphic to

$\mathbb{R}^d$ or to $\mathbb{H}^d$. The *boundary* of $\mathbb{N}$ is the set of points with neighborhoods homeomorphic to $\mathbb{H}^d$. The boundary is either empty or a $(d-1)$-manifold. For complexes it again suffices to check the defining condition at the vertices: $|K|$ is a d-manifold with boundary iff the underlying space of every vertex star is homeomorphic to $\mathbb{R}^d$ or to $\mathbb{H}^d$. Note that every d-manifold is a d-manifold with boundary, namely empty boundary, but a d-manifold with non-empty boundary is not a d-manifold.

Maps. A *homeomorphism* between two topological spaces $\mathbb{X}$ and $\mathbb{Y}$ is a bijection $h : \mathbb{X} \to \mathbb{Y}$ so h and h^{-1} are both continuous. If such an h exists then $\mathbb{X}$ and $\mathbb{Y}$ are *homeomorphic*, denoted as $\mathbb{X} \approx \mathbb{Y}$, and they are said to have the same *topological type*. A *triangulation* of $\mathbb{X}$ is a simplicial complex K with $\mathbb{X} \approx |K|$. $\mathbb{X}$ is *triangulable* if it has a triangulation.

We need some definitions to introduce the combinatorial counterpart of a homeomorphism. A *vertex map* for two complexes K and L is a function $f : \mathrm{Vert}\, K \to \mathrm{Vert}\, L$ so the vertices of a simplex in K are mapped to the vertices of a simplex in L. The *barycentric coordinates* of a point $x \in \sigma$, $\sigma \in K$, are the unique reals $\mathsf{b}_u(x)$, $u \in \mathrm{Vert}\, K$, so $\mathsf{b}_u(x) \neq 0$ only if $u \leq \sigma$ and

$$x = \sum_{u \in \mathrm{Vert}\, K} \mathsf{b}_u(x) \cdot u,$$

$$1 = \sum_{u \in \mathrm{Vert}\, K} \mathsf{b}_u(x).$$

We use barycentric coordinates to extend f in a piecewise linear fashion. The *simplicial map* $\phi : |K| \to |L|$ is defined by

$$\phi(x) = \sum_{u \in \mathrm{Vert}\, K} \mathsf{b}_u(x) \cdot f(u)$$

for every $x \in |K|$. The map ϕ is continuous by construction, but it is neither necessarily injective nor necessarily surjective. It is a homeomorphism iff f is bijective and f^{-1} is also a vertex map. In this case ϕ is an *isomorphism* and K and L are *isomorphic*, which is denoted as $K \sim L$. K and L are *combinatorially equivalent* if they have isomorphic subdivisions, which is denoted as $K \simeq L$. We need the concept of combinatorial equivalence also for subsets of complexes: $B \simeq C$ if there is an isomorphism $|\overline{B}| \to |\overline{C}|$ that maps $|B|$ to $|C|$.

We comment that there is a subtle difference between the piecewise linear and the topological categories. This was first discovered by Milnor [8] who exhibited two homeomorphic triangulations that are not combinatorially equivalent. To avoid related difficulties we stay within the piecewise linear category by basing further definitions on the notion of combinatorial equivalence. All applications of combinatorial equivalence in this paper are to complexes and to sets of simplices whose complement in the closure are complexes. We therefore do not need a combinatorial theory of non-compact spaces.

3. Order and Boundary

The results of this paper rest on the fundamental concept of a stratification defined as a nested sequence of boundaries. For 2-complexes these boundaries have been defined earlier by Whittlesey [13]. We begin by defining the order of a simplex and then proceed to introduce the stratification.

Order of a simplex. Let σ be a simplex in a complex K, and let k be the dimension of the star: $k = \dim \mathrm{St}\, \sigma$. The *order* of σ is the smallest integer $i = \mathrm{ord}\, \sigma$ for which there is a $(k-i)$-simplex η with combinatorially equivalent star: $\mathrm{St}\, \sigma \simeq \mathrm{St}\, \eta$. We assume η belongs to some suitable other complex so its star is defined. Since $\mathrm{int}\, \eta$ is homeomorphic to $\mathbb{R}^{k-i}$ the star of σ is homeomorphic to $\mathbb{R}^{k-i} \times \mathbb{X}$, for some topological space $\mathbb{X}$ of dimension i. Recall that i is chosen as small as possible. The order cannot exceed the difference between the dimension of σ and the dimension of its star:

Order Bound. $\mathrm{ord}\, \sigma \;\leq\; \dim \mathrm{St}\, \sigma - \dim \sigma.$

Proof. For $i = k - \dim \sigma$ we have $\dim \eta = k - i = \dim \sigma$ and can therefore choose $\eta = \sigma$. The stars of η and σ are the same and therefore certainly combinatorially equivalent. □

Shark-fin example. The shark-fin complex in Figure 1 illustrates some of the definitions. It is constructed by gluing two closed disks (triangulations of $\mathbb{B}^2$) along a simple path. That path is a contiguous piece of the boundary of one disk (the fin) and it lies in the interior of the other disk.

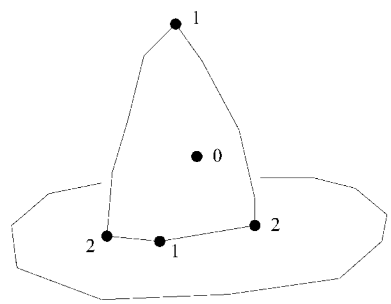

Figure 1. The shark-fin complex has dimension 2 and vertices of all orders: 0, 1, 2.

The dimension of the shark-fin complex is 2 so every triangle has order 0. Each edge belongs to one, two, or three triangles and we call this number the *degree* of the edge. All degree-2 edges have order 0 and the others have order 1. The degree-3 edges are witnesses of the fact that the shark-fin complex is not a manifold with boundary. The violation of the manifold property by degree-3 edges is less severe than the violation that can be found at the two endpoints of their path. The star of an order-0 vertex is a disk, and that of an order-1 vertex is a cycle of half-disks glued along a path of two edges. There are two order-2 vertices, and their stars are more complicated than disks or glued half-disks.

Boundary of a complex. The *j-th boundary* of a simplicial complex K is the set of simplices with order no less than j:

$$\mathrm{Bd}_j K \quad = \quad \{\sigma \in K \mid \mathrm{ord}\,\sigma \geq j\}.$$

By the Order Bound, the j-th boundary contains only simplices of dimension $\dim K - j$ or less. Consider again the shark-fin complex in Figure 1. The 1-st boundary consists of two circles. Even though the 1-st boundary of the two circles is empty, the 2-nd boundary of the shark-fin is non-empty and consists of the two endpoints of the path along which the two disks are glued, see also Property 5 in Appendix A. Note that both boundaries of the shark-fin are complexes. Property 3 in Appendix A shows that this is not a coincidence and that every boundary of a complex is again a complex. Property 2 asserts that the j-th boundary is a topological concept and does not depend on the triangulation. More precisely, the restriction of a simplicial homeomorphism to the j-th boundaries of two complexes is again a simplicial homeomorphism.

The j-th boundary contains the $(j+1)$-st boundary. Hence, if the j-th boundary is empty then all later boundaries are also empty. Underlying spaces of complexes with empty 1-st boundary are manifolds, but the reverse is not true. Underlying spaces of complexes with empty 2-nd boundary can be manifolds with boundary but can also be different. For example, the 2-nd boundary of a 2-complex that triangulates a sphere with equator disk is empty, but the complex is not a manifold with boundary.

Hierarchy of complexes. The boundary concept can be used to define a hierarchy of progressively more complicated complexes. Let M_j be the class of simplicial complexes with empty boundaries beyond index j. Since succeeding boundaries are contained in preceding ones, we have

$$\mathsf{M}_j \quad = \quad \{K \mid \mathrm{Bd}_{j+1} K = \emptyset\}.$$

The only member of M_{-1} is the empty complex. The classes form a nested hierarchy:

$$\{\emptyset\} = \mathsf{M}_{-1} \subset \mathsf{M}_0 \subset \mathsf{M}_1 \subset \mathsf{M}_2 \subset \ldots ,$$

and all inclusions are proper. For a complex we use the minimum index i with $K \in \mathsf{M}_i$ as a measure of how complicated it can get locally. It is plausible but also true that the i-th boundary is at least i classes simpler than the original set.

NESTING LEMMA. *If $K \in \mathsf{M}_j$ then $\mathrm{Bd}_i K \in \mathsf{M}_{j-i}$.*

Proof. $K \in \mathsf{M}_j$ iff $\mathrm{Bd}_{j+1} K = \emptyset$. By Property 5 in Appendix A we have $\mathrm{Bd}_\ell \mathrm{Bd}_i K \subseteq \mathrm{Bd}_{i+\ell} K$ for every index $\ell \geq 0$. This implies $\mathrm{Bd}_{j-i+1} \mathrm{Bd}_i K = \emptyset$, which is equivalent to $\mathrm{Bd}_i K \in \mathsf{M}_{j-i}$. □

Note, however, that $\mathrm{Bd}_i K$ can be more than i classes simpler than K. Consider for example two tetrahedra that meet at a common vertex, u, and let K be the

2-complex of proper faces. Vertex u has order 2 and all other simplices have order 0. It follows that $K \in \mathsf{M}_2$ and $\mathrm{Bd}_1 K = \{u\} \in \mathsf{M}_0$, which is one index stronger than claimed by the Nesting Lemma.

4. Edge Contraction

The motivation for defining j-th boundary is its role in characterizing edge contractions that preserve the topological type. We begin by defining edge contractions and then proceed to discussing conditions under which they permit the construction of isomorphic subdivisions.

Edge contractions. The *contraction* of an edge ab in a complex K replaces $\mathrm{St}\,\overline{ab} = \mathrm{St}\,a \cup \mathrm{St}\,b$ by the star of a new vertex, $\mathrm{St}\,c$. Let $E = \overline{\mathrm{St}}\,\overline{ab}$ and $C = \overline{\mathrm{St}}\,c$ be the closures of the two stars. E and C connect to the rest of K at the common link of $\overline{ab}$ and c, which is $X = E - \mathrm{St}\,\overline{ab} = C - \mathrm{St}\,c$. We can think of the contraction as a surjective simplicial map $\varphi_{ab} : |K| \to |L|$ defined by the surjective vertex map

$$f(u) = \begin{cases} u & \text{if } u \in \mathrm{Vert}\,K - \{a,b\}, \\ c & \text{if } u \in \{a,b\}. \end{cases}$$

Outside $|E|$, φ_{ab} is the identity, but in the interior it is not even injective. We are interested in ways to make edge contractions homeomorphic. An *unfolding* of φ_{ab} is a simplicial homeomorphism $\iota : |K| \to |L|$. It is *local* if ι differs from φ_{ab} only inside $|E|$, and it is *relaxed* if ι differs from φ_{ab} only inside $|\overline{\mathrm{St}}\,E|$. Clearly every local unfolding is also relaxed, but not every relaxed unfolding is local.

Isomorphic subdivisions. Each unfolding ι of φ_{ab} corresponds to a pair of isomorphic subdivisions of K and L. If the subdivisions affect only $\mathrm{St}\,\overline{ab}$ and $\mathrm{St}\,c$ then ι is local, and if they only affect $\mathrm{St}\,E$ and $\mathrm{St}\,C$ then ι is relaxed. Subdivisions of both kinds can be generated from isomorphic subdivisions $\mathrm{Sd}\,E$ of E and $\mathrm{Sd}\,C$ of C. Subdivisions that exploit symmetry need to be avoided since they cannot be combined with the identity. We therefore say the isomorphism $\iota : |\mathrm{Sd}\,E| \to |\mathrm{Sd}\,C|$ *preserves the connection* if $\iota(x) \in \sigma$ for every point $x \in \sigma \in X$, where $X = E \cap C$ as before. We call $\mathrm{Sd}\,E$ *transparent* if $X \subseteq \mathrm{Sd}\,E$, and similar for C. The restriction to $|X|$ of any connection preserving isomorphism defined by transparent subdivisions is necessarily the identity.

Isomorphic Subdivision Lemma. *If E and C have subdivisions $\mathrm{Sd}\,E$ and $\mathrm{Sd}\,C$ admitting a connection preserving isomorphism then φ_{ab} has a relaxed unfolding. If furthermore $\mathrm{Sd}\,E$ and $\mathrm{Sd}\,C$ are transparent then φ_{ab} has a local unfolding.*

Proof. We first show the second claim. Since $\mathrm{Sd}\,E$ is transparent, we can replace E by $\mathrm{Sd}\,E$ and get a subdivision of the entire complex, and similar for C. Let these subdivisions be $K' = (K - E) \cup \mathrm{Sd}\,E$ and $L' = (L - C) \cup \mathrm{Sd}\,C$. To see that K' and L' are isomorphic note that they share $X = E \cap C$ by assumption of transparency. On one side of X we have an isomorphism $|\mathrm{Sd}\,E| \to |\mathrm{Sd}\,C|$ whose restriction to $|X|$ is the identity. On the other side of X we have the identity because $K' - \mathrm{Sd}\,E = L' - \mathrm{Sd}\,C$.

To show the first claim we form K' and L' as before, but because X may have been subdivided, K' and L' may not be complexes. Let $\sigma \in K - E$ be a simplex with a face $\tau \in E$ that has been subdivided in $\operatorname{Sd} E$. In this case $\tau \notin K'$, and to locally repair the complex property we only need to subdivide σ by starring from an interior point. The starring is done inductively in the order of non-decreasing dimension, and it effects only simplices in $\operatorname{St} E - E$. Whenever σ is subdivided within K' it is similarly subdivided within L'. The result are isomorphic subdivisions of K' and L' defining a relaxed unfolding of φ_{ab}. □

By definition, if φ_{ab} has a local or a relaxed unfolding then $K \simeq L$. We will see in Section 5 that the reverse is not true: there are edge contractions with unfoldings that are necessarily global.

Link conditions. We formulate a general condition, which we show implies edge contractions with local and relaxed unfoldings in some cases. For each i we extend the i-th boundary by adding a dummy vertex, ω, and cones from ω to all simplices in the $(i+1)$-st boundary:

$$\operatorname{Bd}_i^\omega K \quad = \quad \operatorname{Bd}_i K \ \cup \ \omega \cdot \operatorname{Bd}_{i+1} K.$$

If $\operatorname{Bd}_{i+1} K = \emptyset$ then $\operatorname{Bd}_i^\omega K = \operatorname{Bd}_i K$. We are only interested in the topology of the extended complex and do not worry about the location of ω and the geometric shape of the cones. For a simplex $\sigma \in \operatorname{Bd}_i^\omega K$ we denote the link within $\operatorname{Bd}_i^\omega K$ as $\operatorname{Lk}_i^\omega \sigma$.

Link Conditions.: $\operatorname{Lk}_i^\omega a \cap \operatorname{Lk}_i^\omega b = \operatorname{Lk}_i^\omega ab$, $\forall i \geq 0$.

Refer to the two portions of a surface triangulation in Figure 2 as examples. In both cases only the links within K are relevant, that is, we only consider the case $i = 0$.

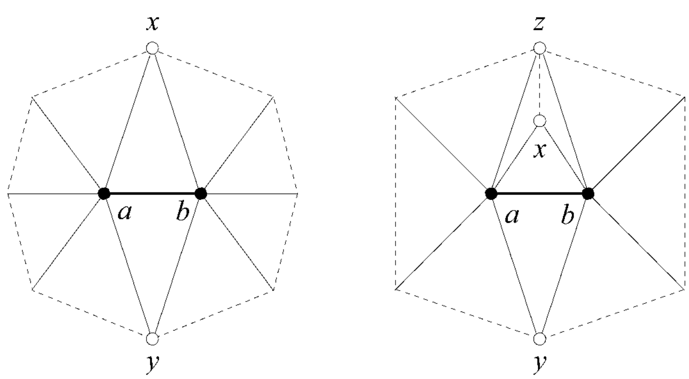

Figure 2. To the left we have $\operatorname{Lk} a \cap \operatorname{Lk} b = \{x, y\} = \operatorname{Lk} ab$, and the contraction of ab has a local unfolding. To the right we have $\operatorname{Lk} a \cap \operatorname{Lk} b = \{x, y, z, xz\} \neq \operatorname{Lk} ab$, and the contraction of ab has no unfolding.

Redundancy. Table 1 unwinds the Link Conditions for a d-complex in M_j. The conditions simplify for large values of i and j. For $i = j$ in the diagonal of Table 1 the extension with ω is redundant. For $i = j = d$ the i-th boundary is a 0-complex so all links are empty and the condition is void. For $i = j = d - 1 \geq 1$

	M_0	M_1	$\dots$	M_{d-1}	M_d
0	Lk_0	Lk_0^ω	$\dots$	Lk_0^ω	Lk_0^ω
1		Lk_1	$\dots$	Lk_1^ω	Lk_1^ω
$\vdots$			$\ddots$	$\vdots$	$\vdots$
$d-1$				$\emptyset$	Lk_{d-1}^ω
d					$\emptyset$

TABLE 1. For a complex in M_j there are $j + 1$ conditions, some may be void ($i = j = d$) and some may be subsumed by others ($i = j = d - 1$).

the condition is subsumed by the condition for $d-2$. To see this note that $\mathrm{Bd}_{d-1} K$ has dimension at most 1. The condition thus simplifies to $\mathrm{Lk}_{d-1}\, a \cap \mathrm{Lk}_{d-1}\, b = \emptyset$, which is violated iff a and b belong to a cycle of three edges. Let x be the third vertex. Then the edge $x\omega \in \mathrm{Lk}_{d-2}^\omega\, a$ because ax belongs to the $(d-1)$-st boundary and thus $ax\omega \in \mathrm{Bd}_{d-2}^\omega K$. Similarly, $x\omega \in \mathrm{Lk}_{d-2}^\omega\, b$, but $x\omega \notin \mathrm{Lk}_{d-2}^\omega\, ab$ because the extended $(d-2)$-nd boundary is a 2-complex and thus contains no tetrahedra.

1-complexes. It is instructive to consider the fairly straightforward case of a 1-complex or graph G. The contraction of an edge $ab \in G$ changes the topological type iff a and b have a common neighbor or both have degree different from 2, see Figure 3.

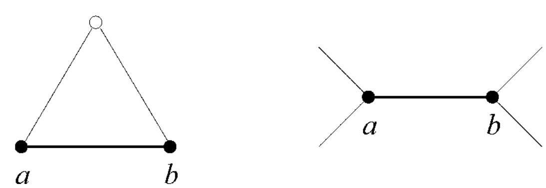

FIGURE 3. The contraction of ab removes a loop to the left and a vertex of the 1-st boundary to the right.

The two cases are captured by the Link Conditions for $d = j = 1$. Indeed, $\mathrm{Lk}_0^\omega\, a \cap \mathrm{Lk}_0^\omega\, b = \emptyset$ iff ab is different from the two cases illustrated in Figure 3. Suppose now that ab satisfies the link condition and assume without loss of generality that a has order 0. In this case we get a local unfolding by subdividing xc into xu, uc, where $x \neq b$ is the other neighbor of a. If ab violates the link condition then there is no unfolding, not even a non-local one.

THEOREM A. *If $G \in \mathsf{M}_1$ is a 1-complex then the following statements are equivalent:*

(i) $\mathrm{Lk}_0^\omega\, a \cap \mathrm{Lk}_0^\omega\, b = \emptyset$.
(ii) φ_{ab} *has a local unfolding.*
(iii) φ_{ab} *has an unfolding.*

5. 2-COMPLEXES

This section proves that for 2-complexes the Link Conditions characterize edge contractions with local unfoldings. This result is strengthened for 2-manifolds where the Link Conditions characterize edge contractions that have any unfolding at all.

Stars and half-stars. To prove that an edge contraction φ_{ab} has a local unfolding we establish transparent subdivisions of $E = \overline{\mathrm{St}}\,\overline{ab}$ and $C = \overline{\mathrm{St}}\,c$ that permit a connection preserving isomorphism. This task is simplified by assuming $\mathrm{ord}\,a = \mathrm{ord}\,ab$ and choosing the new vertex c equal to b. The contraction can then be visualized by sliding a towards and eventually merging into b, see Figure 4. The operation only affects simplices in $A = \overline{\mathrm{St}}\,a$ and leaves simplices in $E - A$ unchanged. Call $R = \mathrm{Lk}\,a = A - \mathrm{St}\,a$ the *rim* of A. Of the simplices in A the ones in $\overline{\mathrm{St}}\,ab$ disappear or merge into the rim, and the others remain but assume different geometric shape and position. We call $A' = A - \mathrm{St}\,b$ the *half-star* of a and $R' = A' - (\mathrm{St}\,a - \overline{\mathrm{St}}\,b)$ the *rim* of A'. The image of A' under the contraction is an isomorphic subcomplex C' of C. To establish transparent subdivisions of E and C that permit a connection preserving isomorphism it suffices to construct such subdivisions of A and C'. This is equivalent to constructing isomorphic subdivisions of A and A' that permit an isomorphism whose restriction to the intersection of the two rims is the identity. It is therefore essential that $R \sim R'$, which will always be the case when we apply the construction.

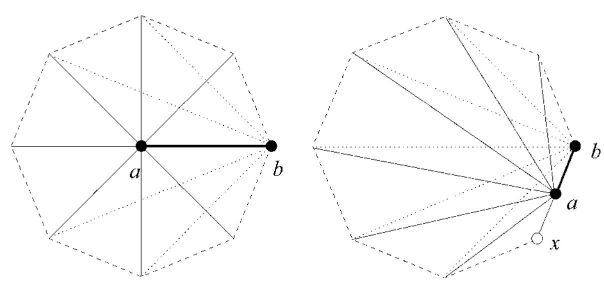

FIGURE 4. The overlay of two subdivisions of a regular k-gon. To the left the subdivisions are obtained by starring from the center and from a vertex. To the right they are obtained by starring from the midpoint and from an endpoint of an edge.

Figure 4 illustrates the construction of isomorphic transparent subdivisions. Both A and A' are mapped isomorphically to subdivisions of the same regular

k-gon. To the left, A and A' are disks of k and $k-2$ triangles. To the right, A and A' are half-disks of $k-1$ and $k-2$ triangles.

General 2-complexes. Let K be a 2-complex and $ab \in K$. Recall that the contraction of ab is a simplicial map $\varphi_{ab} : |K| \to |L|$. There are three Link Conditions and Table 1 indicates that the last one is void.

THEOREM B. *If $K \in \mathsf{M}_2$ is a 2-complex then the following statements are equivalent:*

(i) *ab satisfies the Link Conditions for $j = 2$:*
  (i.0) $\mathrm{Lk}_0^{\omega}\, a \cap \mathrm{Lk}_0^{\omega}\, b = \mathrm{Lk}_0^{\omega}\, ab$, *and*
  (i.1) $\mathrm{Lk}_1^{\omega}\, a \cap \mathrm{Lk}_1^{\omega}\, b = \emptyset$.
(ii) *φ_{ab} has a local unfolding.*

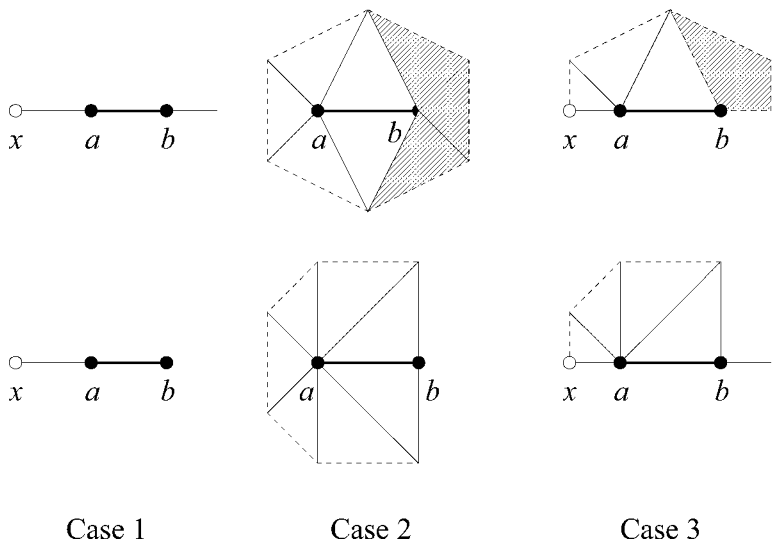

FIGURE 5. From left to right, ab is principal, ab is not principal and does not belong to the 1-st boundary, ab belongs to the 1-st boundary. In each case $\mathrm{ord}\, a = \mathrm{ord}\, ab$ and in the upper row b has the same and in the lower row it has higher order.

Proof. (i) $\Longrightarrow$ (ii). The argument for sufficiency of the Link Conditions distinguishes three cases all of which are illustrated in Figure 5. We note that the order of an endpoint of edge ab is at least as large as the order of ab, see Property 3 of Appendix A. The Link Conditions imply that not both endpoints can exceed the order of ab, and we assume without loss of generality that $\mathrm{ord}\, a = \mathrm{ord}\, ab$.

`Case 1.:` ab is principal. Thus $\mathrm{ord}\, ab = \mathrm{ord}\, a = 0$, and a belongs to exactly two edges: ab and ax. L is obtained from K by removing xa, a, ab and adding xb. The isomorphic transparent subdivisions of $\{xa, a, ab\}$ and $\{xb\}$ are obvious.

`Case 2.:` ab is not principal and $\operatorname{ord} ab = \operatorname{ord} a = 0$. Map $A = \overline{\mathrm{St}}\, a$ isomorphically to a subdivision of the regular k-gon as in Figure 4 to the left, where k is the number of triangles in A. Similarly, map the half-star $A' = A - \mathrm{St}\, b$ isomorphically to a subdivision of the regular k-gon into $k-2$ triangles. The two rims map isomorphically to the boundary of the k-gon and the maps agree at $R \cap R'$. We can therefore form a common subdivision that leaves the boundary of the k-gon unchanged. This subdivision maps back to transparent subdivisions $\mathrm{Sd}\, A$, $\mathrm{Sd}\, A'$, and $\mathrm{Sd}\, C'$. We add the remaining simplices of E to $\mathrm{Sd}\, A$ and the remaining simplices of C to $\mathrm{Sd}\, C'$ and get transparent subdivisions of E and C that permit an isomorphism whose restriction to the underlying space of $X = E \cap C$ is the identity. In other words, the isomorphism preserves the connection. The Isomorphic Subdivision Lemma implies (ii).

`Case 3.:` $\operatorname{ord} ab = \operatorname{ord} a = 1$. Let ax be the other edge with order 1 in the star of a. A is a cycle of half-disks, and each half-disk is a fan of triangles starting at ax and ending at ab. Let D be the closure of such a half-disk and let $k-1$ be the number of triangles in D. Map D isomorphically to a subdivision of the regular k-gon with the image of a at the midpoint of a k-gon edge, as in Figure 4 to the right. A' is another cycle of half-disks, this time around ax. Let D' be the closure of the half-disk that corresponds to D. Map D' isomorphically to a subdivision of the regular k-gon, with the image of a at the vertex that is the image of b under the earlier map. We form a common subdivision that leaves the k-gon boundary unchanged, except for one edge which is cut into two by the image of a under the first map. We denote this edge as yz. That subdivision maps back to isomorphic subdivisions of D and D'. Both subdivisions are transparent except at the preimages of yz. The subdivisions of the half-disks of A meet along xa, ab and together they form $\mathrm{Sd}\, A$, which is transparent. Similarly, the subdivisions of the half-disks of A' meet along xa. The images of these subdivisions together form $\mathrm{Sd}\, C'$, which is also transparent. By construction $\mathrm{Sd}\, A$ and $\mathrm{Sd}\, C'$ are isomorphic. After adding the remaining simplices of E to $\mathrm{Sd}\, A$ and those of C to $\mathrm{Sd}\, C'$ we have transparent subdivisions of E and C that permit a connection preserving isomorphism. The Isomorphic Subdivision Lemma implies (ii).

$\neg$(i) $\Longrightarrow$ $\neg$(ii). The argument for necessity of the Link Conditions distinguishes between the violation of (i.0) and that of (i.1). We will show that in either case the topological type of the link of at least one simplex in $\mathrm{Lk}\, \overline{ab}$ changes. Since a local unfolding is the identity outside $\mathrm{St}\, \overline{ab}$, this contradicts its existence. Note that the link of an edge is contained in the link of its endpoints and define $L_0 = (\mathrm{Lk}_0^\omega\, a \cap \mathrm{Lk}_0^\omega\, b) - \mathrm{Lk}_0^\omega\, ab$ and $L_1 = \mathrm{Lk}_1^\omega\, a \cap \mathrm{Lk}_1^\omega\, b$.

`Case 1.:` (i.0) is violated, or equivalently $L_0 \neq \emptyset$. First suppose that L_0 contains an edge xy. The link of xy is a finite set of vertices that contains both a and b. The contraction of ab decreases the cardinality of $\mathrm{Lk}\, xy$ and thus changes its type. Next suppose that L_0 contains no edge but it contains a vertex x. $\mathrm{Lk}\, x$ is a 1- or 0-complex that contains both a and b as vertices.

If a and b belong to different components then the contraction of ab merges these components. Otherwise, every path from a to b in $\mathrm{Lk}\, x$ has at least three edges, for $ay, by \in \mathrm{Lk}\, x$ implies $ayx, byx \in K$ and therefore $yx \in L_0$, which contradicts the assumption. The contraction of ab thus forms one or more cycles. In both cases the topological type of $\mathrm{Lk}\, x$ changes.

`Case 2.:` (i.0) is satisfied and (i.1) is violated, or equivalently $L_0 = \emptyset$ and $L_1 \neq \emptyset$. Then $a, b \in \mathrm{Bd}_1 K$ for else L_1 would be empty, and $ab \in \mathrm{Bd}_1 K$ for else $\omega \in L_0$ would violate (i.0). Let x be a vertex in L_1. If $x = \omega$ then $a, b \in \mathrm{Bd}_2 K$. $\mathrm{St}\, \overline{ab}$ thus contains two vertices of order 2, while $\mathrm{St}\, c$ contains at most one such vertex, namely c. This contradicts $E \simeq C$. Finally suppose $x \neq \omega$. We have $a, b, ab \in \mathrm{Lk}\, x$, and because $xa, xb \in \mathrm{Bd}_1 K$ the degrees of a and b in $\mathrm{Lk}\, x$ are both different from 2. In other words, a and b both belong to $\mathrm{Bd}_1 \mathrm{Lk}\, x$, which is a set of vertices. The contraction of ab removes a vertex from the 1-st boundary and thus changes the type of $\mathrm{Lk}\, x$. □

Non-local isomorphism. An edge contraction with local unfolding preserves the topological type, but there are type preserving edge contractions that have neither a local nor a relaxed unfolding. An example is the folding chair complex illustrated in Figure 6. Before the contraction of ab the complex consists of 5 triangles in the star of x and 4 disks U, V, Y, Z glued to the link of x as shown. Vertices a and b belong to the 1-st boundary, but ab does not. The dummy vertex ω thus violates the Link Condition for $i = 0$ and so does x.

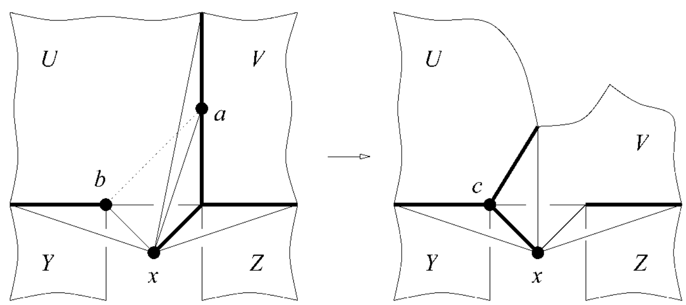

FIGURE 6. The 2-dimensional folding chair complex. Fat edges belong to three triangles each.

After the contraction there is one less triangle in the star of x, U loses two triangles, and V, Y, Z are unchanged. The contraction exchanges left and right in the asymmetry of the complex. We can find a homeomorphism $|K| \to |L|$ that acts like a mirror, mapping U to V, V to U, Y to Z, Z to Y. Indeed, every homeomorphism must act this way and differ from the identity almost everywhere.

2-manifolds. The 2-complex K belongs to M_0 iff $|K|$ is a 2-manifold. In this case condition (i.1) in Theorem B is void and condition (i.0) simplifies because the extension with ω is redundant. We strengthen the result implied by Theorem B by proving that the violation of the Link Condition contradicts the existence of any unfolding, whether local or not.

THEOREM B_0. *If* $K \in \mathsf{M}_0$ *is a 2-complex then the following statements are equivalent:*

(i) $\mathrm{Lk}\, a \cap \mathrm{Lk}\, b = \mathrm{Lk}\, ab$.
(ii) φ_{ab} *has a local unfolding.*
(iii) φ_{ab} *has an unfolding.*

Proof. (i) $\Longrightarrow$ (ii) follows from Theorem B and (ii) $\Longrightarrow$ (iii) from the definitions. To prove $\neg$(i) $\Longrightarrow$ $\neg$(iii) we distinguish two cases depending on the dimension of the violating simplex. Let $L_0 = (\mathrm{Lk}\, a \cap \mathrm{Lk}\, b) - \mathrm{Lk}\, ab$. $\neg$(i) is equivalent to $L_0 \neq \emptyset$.

`Case 1.:` L_0 contains an edge xy. Then axy and bxy are replaced by a single triangle cxy. Hence xy belongs to only one triangle in L, which contradicts $L \in \mathsf{M}_0$.

`Case 2.:` L_0 contains no edge but it contains a vertex x. Then ax and bx are edges in K. Each belongs to two triangles: $apx \neq aqx$, $brx \neq bsx$. The four triangles are pairwise different because $abx \notin K$. The four vertices p, q, r, s are pairwise different because L_0 contains no edge. Hence cx belongs to four triangles in L, which contradicts $L \in \mathsf{M}_0$. □

6. 3-COMPLEXES

This section extends Theorems B and B_0 to complexes of dimension 3. We begin with the main geometric tool, which is Steinitz' classical theorem for convex 3-polytopes [11].

Steinitz' theorem. A *convex* 3*-polytope* is the convex hull of finitely many points in $\mathbb{R}^3$ that do not all lie in a common plane. Its boundary is a complex of vertices, edges, and (2-dimensional) facets. If the points are in general position then all facets are triangles. The 1*-skeleton* is the subcomplex of all vertices and edges. A *graph* G is a 1-complex. It is *connected* if for every partition of $\mathrm{Vert}\, G$ into two non-empty sets, G contains an edge with one endpoint in each set. A connected graph is *three-connected* if the deletion of any two vertices together with their edges leaves the graph connected. G is *planar* if it is isomorphic to a 1-complex in $\mathbb{R}^2$. For example, the 1-skeleton of every convex 3-polytope is planar and three-connected. A fundamental result by Steinitz asserts that these 1-skeletons exhaust all three-connected planar graphs [11].

STEINITZ' THEOREM. *For every three-connected planar graph there is a convex 3-polytope with an isomorphic 1-skeleton.*

Planar graphs that are not three-connected arise by removing edges and vertices. Let X be a triangulation of $\mathbb{S}^2$. Let $a \in X$ be a vertex and consider the graph G consisting of all edges and vertices in X other than the ones in the star of a. A drawing of G in $\mathbb{R}^2$ has one k-gon and otherwise only triangles. The boundary of the k-gon is the link of a within X, and we denote it by U. In Figure 7 we have $k = 8$ and the k-gon is the unbounded outer region of the drawing. G is three-connected iff no edge $xy \in G - U$ has both endpoints in U. If there is such an edge xy we repair three-connectedness by cutting xy at an interior point z and connecting z to the opposite vertices of the two adjacent triangles. This is done in any opportune sequence over all such edges xy, as in Figure 7. The operation corresponds to subdividing X by starring from the points z in the same sequence.

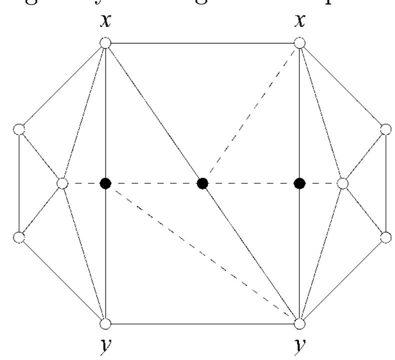

FIGURE 7. A triangulation of the 2-sphere after removing a vertex of degree $k = 8$. There are three edges xy that are cut into two each to restore three-connectedness.

General 3-complexes. Let K be a 3-complex and let $ab \in K$. There are four Link Conditions, and Table 1 indicates that the last one is void.

THEOREM C. *If $K \in \mathsf{M}_3$ is a 3-complex then the first statement implies the second:*

(i) *ab satisfies the Link Conditions for $j = 3$:*
  (i.0) $\mathrm{Lk}_0^\omega\, a \cap \mathrm{Lk}_0^\omega\, b = \mathrm{Lk}_0^\omega\, ab$,
  (i.1) $\mathrm{Lk}_1^\omega\, a \cap \mathrm{Lk}_1^\omega\, b = \mathrm{Lk}_1^\omega\, ab$, *and*
  (i.2) $\mathrm{Lk}_2^\omega\, a \cap \mathrm{Lk}_2^\omega\, b = \emptyset$.
(ii) *φ_{ab} has a relaxed unfolding.*

Proof. We distinguish four cases, the first of which has been treated in the proof of Theorem B. Recall that the order of an endpoint of ab is at least as large as the order of ab. The Link Conditions imply that not both endpoints can exceed the order of ab, and we assume $\mathrm{ord}\, a = \mathrm{ord}\, ab$.

`Case 1.:` $\dim \mathrm{St}\, ab \leq 2$. The Order Bound implies $\mathrm{ord}\, ab = \mathrm{ord}\, a \leq 1$. The presence of a tetrahedron in $\mathrm{St}\, a$ would imply one in $\mathrm{St}\, ab$, hence $\dim \mathrm{St}\, a \leq 2$. In other words, the neighborhood of a is as in Theorem B. Let K' be the 2-complex obtained from K by removing all simplices σ that satisfy $\dim \mathrm{St}\, \sigma = 3$ and $\mathrm{ord}\, \sigma = 0$. For a simplex $\sigma \in K'$ the order in K' is either the same or less than the order in K. Specifically, if $\dim \mathrm{St}\, \sigma \leq 2$ in K then the star remains unchanged and so does $\mathrm{ord}\, \sigma$. If $\dim \mathrm{St}\, \sigma = 3$ in K then the dimension of the star drops and so does $\mathrm{ord}\, \sigma$. We will argue shortly that the 3-dimensional Link Conditions for K imply the 2-dimensional Link Conditions for K'. Theorem B therefore applies and we get a local unfolding $\varphi'_{ab} : |K'| \to |L'|$ which differs from the identity only inside $\|\mathrm{St}\, a\|$. A local unfolding of φ_{ab} is obtained by extending φ'_{ab} with the identity inside all simplices in $K - K'$.

We now argue that if ab satisfies (i) then $ab \in K'$ satisfies the Link Conditions for $j = 2$. Since $\dim \mathrm{St}\, ab \leq 2$ in K, the star of ab is the same in K and in K'. It follows that $\mathrm{Lk}_0^\omega\, ab$ is the same in K and K'. In general, stars and boundaries cannot increase from K to K'. If follows that $\mathrm{Lk}_0^\omega\, a$ and $\mathrm{Lk}_0^\omega\, b$ do not increase. Since the link of an edge is always contained in the links of its vertices, $\mathrm{Lk}_0^\omega\, a \cap \mathrm{Lk}_0^\omega\, b = \mathrm{Lk}_0^\omega\, ab$ in K implies the same in K'. The second Link Condition applies only if the order of both a and b is at least 1 in K'. Then $\mathrm{ord}\, ab \geq 1$ because of (i.0). But ab is principal in $\mathrm{Bd}_1\, K$ so its link is empty. The links of a and b can again not increase from $\mathrm{Bd}_1^\omega\, K$ to $\mathrm{Bd}_1^\omega\, K'$. Hence $\mathrm{Lk}_1^\omega\, a \cap \mathrm{Lk}_1^\omega\, b = \emptyset$ in K implies the same in K'.

`Case 2.:` $\dim \mathrm{St}\, ab = 3$ and $\mathrm{ord}\, ab = \mathrm{ord}\, a = 0$. Define $A = \overline{\mathrm{St}}\, a$ as usual. Using Steinitz' Theorem we map A isomorphically to a subdivision S of a convex 3-polytope P, with the image of a in the interior, the image of b at a vertex v of P, and the rim $R = \mathrm{Lk}\, a$ isomorphic to the boundary complex of P. Similarly, we map the half-star $A' = A - \mathrm{St}\, b$ isomorphically to another subdivision S' of P, with the image of a at v, and the rim R' isomorphic to the boundary complex of P. We construct a common subdivision T of S and S' that keeps the boundary complex unchanged. T maps back to $\mathrm{Sd}\, A$, $\mathrm{Sd}\, A'$, and $\mathrm{Sd}\, C'$, all transparent and isomorphic. We add the remaining simplices of E to $\mathrm{Sd}\, A$ and the remaining simplices of C to $\mathrm{Sd}\, C'$ and obtain transparent subdivisions of E and C. By construction there is an isomorphism that preserves the connection. The Isomorphic Subdivision Lemma implies that φ_{ab} has a local unfolding.

`Case 3.:` $\dim \mathrm{St}\, ab = 3$ and $\mathrm{ord}\, ab = \mathrm{ord}\, a = 1$. By definition there is a triangle η with a 3-dimensional star $\mathrm{St}\, \eta \simeq \mathrm{St}\, a$ consisting of one, three, or more tetrahedra that share η. The number of tetrahedra cannot be two, else we would have $\mathrm{ord}\, a = 0$. Let U be the set of simplices with order 1 in $\mathrm{St}\, a$. It corresponds to η in $\mathrm{St}\, \eta$ and therefore forms an open disk that decomposes $\mathrm{St}\, a$ into one, three, or more components. A component B_1 of $\mathrm{St}\, a - U$ has only order 0 simplices.

The closure of a component, $A_1 = \overline{B}_1$, is a triangulation of $\mathbb{B}^3$, and the boundary, $X_1 = \mathrm{Bd}_1\, A_1$, is a triangulation of $\mathbb{S}^2$. We use Steinitz' Theorem to map A_1 to a subdivision S_1 of a convex 3-polytope P_1. Except for a all vertices of A_1 map to vertices of P_1, and a maps to a point in the interior of a facet of P_1. To accomplish this proceed as described earlier: construct the 1-complex of edges and vertices in X_1, remove a and its edges, cut and add edges to restore three-connectedness, and let G_1 be the isomorphic 1-skeleton of a convex 3-polytope P_1. Finally, subdivide P_1 by starring from the image of a. Because edges were cut and added, the boundary complex of P_1 is isomorphic to a subdivision of X_1 but not necessarily to X_1 itself. Similarly, S_1 is isomorphic to a subdivision of A_1 but not necessarily to A_1 itself. The image of b is a vertex of the k-gon facet. We form a second subdivision S_1' of

P_1 by starring from this vertex. S_1' is isomorphic to a subdivision of the half-star $A_1' = A_1 - \mathrm{St}\, b$. The right picture in Figure 4 illustrates S_1 and S_1' in the 2-dimensional case where P_1 is a convex polygon. The common subdivision T_1 of S_1 and S_1' is obtained as usual, by intersection and starring. T_1 maps back to a subdivision of A_1 and a subdivision of A_1'. These subdivisions are not necessarily transparent.

To finish the argument we repeat the construction for all other components B_ℓ of $\mathrm{St}\, a - U$. The subdivisions of the A_ℓ are glued along the preimage of U, which is subdivided as a result of mapping back the subdivision of the k-gon facet, see left picture in Figure 4. Here it is important that the subdivisions of the facets be isomorphic, but this can easily be achieved because the image of a can be freely chosen anywhere in the interior of the facet. The result is a subdivision $\mathrm{Sd}\, A$ of A. We add the remaining simplices of E, possibly after subdivision because $\mathrm{Sd}\, A$ is not necessarily transparent, and obtain $\mathrm{Sd}\, E$. Similarly, the subdivisions of the A_ℓ' are glued along the preimage of U and mapped to $\mathrm{Sd}\, C'$. We add the remaining simplices of C, possibly after subdivision, and obtain $\mathrm{Sd}\, C$. Since the boundaries of $\mathrm{Sd}\, A$ and $\mathrm{Sd}\, A'$ are isomorphic, the subdivision of the remaining simplices in C can be done such that $\mathrm{Sd}\, E \sim \mathrm{Sd}\, C$. By construction, $\mathrm{Sd}\, E$ and $\mathrm{Sd}\, C$ permit an isomorphism that preserves the connection. The Isomorphic Subdivision Lemma implies (ii).

`Case 4.:` $\mathrm{ord}\, ab = \mathrm{ord}\, a = 2$. By the Order Bound, the dimension of the stars is $\dim \mathrm{St}\, ab = \dim \mathrm{St}\, a = 3$. It follows that a belongs to exactly two edges of order 2, ab and xa. The argument is similar to Case 3. The disk U is replaced by a ring U of half-disks glued along xa, ab. Again, U decomposes $\mathrm{St}\, a$ into one or more components. The closure A_ℓ of each such component B_ℓ is a triangulation of $\mathbb{B}^3$. The boundary $X_\ell = \mathrm{Bd}_1\, A_\ell$ is a triangulation of $\mathbb{S}^2$, xa, ab are edges of X_ℓ, and the closed star of a within X_ℓ is a disk consisting of two half-disks in the ring U.

The use of Steinitz' Theorem is similar to Case 3 except that now we map a to the interior point of an edge. To accomplish this, we modify the construction of the graph by adding the edge xb after removing a and its edges. The drawing in the plane has a k-gon adjacent to an m-gon and otherwise only triangles. Three-connectedness is recovered by cutting and adding edges that neither belong to the k-gon nor to the m-gon. Let G_ℓ be the isomorphic 1-skeleton of a convex 3-polytope P_ℓ. A subdivision S_ℓ of P_ℓ is obtained by starring from the image of a. Another subdivision S_ℓ' of P_ℓ is obtained by starring from the image of b, which is an endpoint of the edge common to the k-gon and the m-gon. T_ℓ is again a common subdivision of S_ℓ and S_ℓ' and is mapped back to isomorphic subdivisions of A_ℓ and A_ℓ'. The subdivisions of the A_ℓ are glued to form $\mathrm{Sd}\, A$ and the subdivision of the A_ℓ' are glued and mapped to form $\mathrm{Sd}\, C'$. Finally, the remaining simplices of E and C are added, possibly after subdivision, to obtain subdivisions $\mathrm{Sd}\, E$ and $\mathrm{Sd}\, C$ with a connection preserving isomorphism. The Isomorphic Subdivision Lemma implies (ii). □

3-manifolds. Steinitz' theorem can be applied to the vertex links of a 3-manifold K to prove $K \in \mathsf{M}_0$. For 3-manifolds, the Link Conditions consolidate to a single relation. We strengthen the result implied by Theorem C in two respects: we construct *local* unfoldings, and we show the Link Condition is *equivalent* to the existence of an unfolding.

THEOREM C_0. *If $K \in \mathsf{M}_0$ is a 3-complex then the following statements are equivalent:*

(i) $\mathrm{Lk}\, a \cap \mathrm{Lk}\, b = \mathrm{Lk}\, ab$.
(ii) φ_{ab} *has a local unfolding.*
(iii) φ_{ab} *has an unfolding.*

Proof. Only Case 2 of the proof of Theorem C arises for 3-manifolds. In this case the conclusion is that φ_{ab} has a local unfolding, which shows (i) $\Longrightarrow$ (ii). (ii) $\Longrightarrow$ (iii) follows from the definitions. To prove $\neg$(i) $\Longrightarrow$ $\neg$(iii) we distinguish three cases depending on the dimension of the violating simplex. Let $L_0 = (\mathrm{Lk}\, a \cap \mathrm{Lk}\, b) - \mathrm{Lk}\, ab$. $\neg$(i) is equivalent to $L_0 \neq \emptyset$.

`Case 1.:` L_0 contains a triangle xyz. Then $axyz, bxyz$ are replaced by a single tetrahedron $cxzy$. If follows that xyz belongs to only one tetrahedron in the complex L obtained from K by contracting ab. This contradicts $L \in \mathsf{M}_0$.

`Case 2.:` L_0 contains no triangle but it contains an edge xy. Then axy and bxy are triangles in K. Each belongs to two tetrahedra: $apxy \neq aqxy$, $brxy \neq bsxy$. The four tetrahedra are pairwise different because $abxy \notin K$, which follows from $xy \notin \mathrm{Lk}\, ab$. The four vertices p, q, r, s are pairwise different because L_0 contains no triangle. Hence, cxy belongs to four tetrahedra in L, which contradicts $L \in \mathsf{M}_0$.

`Case 3.:` L_0 contains no edge and no triangle, but it contains a vertex x. Then ax and bx are edges in K. Their links are two circles. These circles are disjoint because $y \in \mathrm{Lk}\, ax \cap \mathrm{Lk}\, bx$ would imply that L_0 contains an edge, namely xy. We have $b \notin \mathrm{Lk}\, ax$ because else $abx \in K$ and hence $x \notin L_0$. Similarly, we have $a \notin \mathrm{Lk}\, bx$. It follows that after the contraction of ab to c both circles belong to $\mathrm{Lk}\, cx$, which contradicts $L \in \mathsf{M}_0$. □

7. Discussion

This section concludes the paper with an open problem and a comment on simplifying manifolds using edge contractions.

Link conditions. The most important remaining problem is the extension of the link condition results to complexes of dimension beyond 3. At this time the limitation of Steinitz' Theorem to convex polytopes of dimension at most 3 is an obstacle in extending the proofs of this paper. Do Theorems B_0 and C_0 extend to combinatorial d-manifolds for $d \geq 4$? Specifically, is it true that for every $K \in \mathsf{M}_0$ the contraction of an edge $ab \in K$ has an unfolding iff $\mathrm{Lk}\, a \cap \mathrm{Lk}\, b = \mathrm{Lk}\, ab$? Is there a general result that relates the Link Conditions with topology preserving edge contractions for simplicial complexes of any fixed dimension? Specifically, do the Link Conditions for d and j imply that the contraction of ab in a d-complex $K \in \mathsf{M}_j$ has an unfolding?

Irreducible triangulations. A simplicial complex is *irreducible* if it has no edge whose contraction preserves the topological type. It is not difficult to prove that the boundary complex of the tetrahedron is the only irreducible triangulation of the 2-sphere. It is also known that every compact 2-manifold has only finitely many irreducible triangulations [1]. In other words, topology preserving edge contractions can be used to quickly classify a triangulated 2-manifold.

The classification problem for 3-manifolds is considerably more difficult [4], and for 4-manifolds it is known to be undecidable [7]. Algorithms that recognize the 3-sphere have been found only recently [10, 12]. In view of the apparent difficulties, it is not surprising that even the 3-sphere has infinitely many irreducible triangulations. To construct an infinite family we use knots made up of only three edges each. Take a 3-cube decomposed into n^3 little cubes, and let n be large enough so we can drill a one cube wide tunnel in the form of a non-trivial knot. Instead of completing the drilling we leave the last cube of the tunnel so the complex is still homeomorphic to $\mathbb{B}^3$. Let ab be one of the edges of the retained last cube that connects the end of the incomplete tunnel with the outside. A triangulation of $\mathbb{B}^3$ is formed by decomposing each little cube into tetrahedra, which is done without adding new vertices. Finally, a triangulation of $\mathbb{S}^3$ is obtained by adding the cone from a new vertex, x, over the boundary complex of the triangulation of $\mathbb{B}^3$. The cycle of three edges ab, bx, xa forms a knot of the type of the tunnel. The triangle abx is not part of the triangulation because ab is not part of the boundary complex. In fact, this triangle cannot be embedded in $\mathbb{S}^3$ because ab, bx, xa is a non-trivial knot. Hence none of the three edges can be contracted without changing the topological type of the triangulation. We get an infinite family by drilling tunnels of different knot types. Indeed, if we had a finite set of irreducible triangulations we would have only finitely many cycles of three edges and thus only finitely many knot types. This contradicts the existence of infinitely many different knot types.

Acknowledgements

The second author thanks Wolfgang Haken and Min Yan for interesting discussions and Günter Ziegler for suggesting the knot construction in the triangulation of the 3-sphere mentioned in Section 7.

References

[1] D. W. Barnette and A. L. Edelson, *All 2-manifolds have finitely many minimal triangulations,* Israel J. Math. **67** (1989), 123–128.

[2] H. Edelsbrunner, M. A. Facello and D. V. Nekhayev, *Surface remeshing and parametrization,* Report rgi-tech-98-019, Raindrop Geomagic, Champaign, Illinois, 1998.

[3] M. Garland and P. S. Heckbert, *Surface simplification using quadratic error metrics,* Computer Graphics, Proc. siggraph 1997, 209–216.

[4] W. Haken, *Über das Homöomorphie Problem der 3-Mannigfaltigkeiten I,* Math. Z. **80** (1962), 89–120.
[5] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald and W. Stützle, *Mesh optimization,* Computer Graphics, Proc. siggraph 1993, 19–26.
[6] A. W. F. Lee, W. Sweldens, P. Schröder, L. Cowsar and D. P. Dobkin, *MAPS: multiresolution adaptive parameterization of surfaces,* Computer Graphics, Proc. siggraph 1998, 95–104.
[7] A. A. Markov, *Insolubility of the problem of homeomorphy,* In: "Proc. Internat. Congress Math., 1960", Cambridge Univ. Press.
[8] J. Milnor, *Two complexes which are homeomorphic but combinatorially distinct,* Ann. of Math. **74** (1961), 575–590.
[9] J. R. Munkres, *Elements of Algebraic Topology,* Addison-Wesley, Redwood City, 1984.
[10] J. H. Rubinstein, *An algorithm to recognise* $\mathbb{S}^3$*,* In: "Proc. Internat. Congress Math., 1994", Zürich, Switzerland.
[11] E. Steinitz, *Polyeder und Raumeinteilungen,* Enzykl. Math. Wiss. **3** (Geometrie), Part 3AB12 (1922), 1–139.
[12] A. Thompson, *Thin position and the recognition problem for* $\mathbb{S}^3$*,* Math. Res. Lett. **1** (1994), 613–630.
[13] E. F. Whittlesey, *Finite surfaces: a study of finite 2-complexes,* Math. Mag. **34** (1960), 11–22 and 67–80.

Appendix A

This appendix proves basic properties of boundary as defined in Section 3. All properties are intuitively clear but the proofs are somewhat technical.

Invariance of order. Isomorphic subdivisions are constructed by just one method also described in the proof of the Isomorphic Subdivisions Lemma: map the complexes to a common underlying space, intersect simplices, and subdivide by starring. We use this construction to establish that simplices with combinatorially equivalent stars indeed have the same order.

Property 1. *If* $\mathrm{St}\,\tau \simeq \mathrm{St}\,\sigma$ *then* $\mathrm{ord}\,\tau = \mathrm{ord}\,\sigma$.

Proof. Let $\ell = \mathrm{ord}\,\sigma$ and $k = \dim \mathrm{St}\,\sigma$. By definition of order there is a $(k-\ell)$-simplex η in some hypothetical complex with $\mathrm{St}\,\sigma \simeq \mathrm{St}\,\eta$. Construct a subdivision L of $\overline{\mathrm{St}}\,\sigma$ by combining the subdivision isomorphic to $\mathrm{Sd}\,\overline{\mathrm{St}}\,\tau$ with the one isomorphic to $\mathrm{Sd}\,\overline{\mathrm{St}}\,\eta$. Map L to new and finer subdivisions of $\overline{\mathrm{St}}\,\tau$ and $\overline{\mathrm{St}}\,\eta$ using the simplicial homeomorphisms $\phi : \|\overline{\mathrm{St}}\,\sigma\| \to \|\overline{\mathrm{St}}\,\tau\|$ and $\psi : \|\overline{\mathrm{St}}\,\sigma\| \to \|\overline{\mathrm{St}}\,\eta\|$. By definition of combinatorial equivalence for subsets of complexes we have $\phi(\|\mathrm{St}\,\sigma\|) = \|\mathrm{St}\,\tau\|$ and $\psi(\|\mathrm{St}\,\sigma\|) = \|\mathrm{St}\,\eta\|$. Then $\psi \circ \phi^{-1}$ is an isomorphism between these finer subdivisions and $\psi \circ \phi^{-1}(\|\mathrm{St}\,\tau\|) = \|\mathrm{St}\,\eta\|$. This implies $\mathrm{St}\,\tau \simeq \mathrm{St}\,\eta$ and $\mathrm{ord}\,\tau \le \dim \mathrm{St}\,\tau - \dim \eta = \ell = \mathrm{ord}\,\sigma$. By the symmetric argument we get $\mathrm{ord}\,\sigma \le \mathrm{ord}\,\tau$. ☐

Boundary commutes with subdivision. We show that the boundary of a subdivided simplicial complex K is the same as the boundary of K subdivided. This is intuitively what one expects as the subdivision operation does not change the geometric neighborhood of any point in the underlying space of K.

PROPERTY 2. $\mathrm{Bd}_i\,\mathrm{Sd}\,K = \mathrm{Sd}\,\mathrm{Bd}_i\,K$.

Proof. For each $\tau \in \mathrm{Sd}\,K$ there is a unique simplex $\sigma \in K$ with $\mathrm{int}\,\tau \subseteq \mathrm{int}\,\sigma$. We prove that τ and σ have the same order by showing that their closed stars are simplicially equivalent.

Choose a point $x \in \mathrm{int}\,\tau \subseteq \mathrm{int}\,\sigma$. Let $\mathbb{C}^{e-1}$ be the boundary of a cube with center at x, where e is the dimension of the ambient Euclidean space. Subdivide both $\overline{\mathrm{St}}\,\tau$ and $\overline{\mathrm{St}}\,\sigma$ by starring from x. In the first subdivision the closed star of x has the same underlying space as $\overline{\mathrm{St}}\,\tau$ and the link of x is the difference between the closed star and the star. Similarly, in the second subdivision the closed star of x has the same underlying space as $\overline{\mathrm{St}}\,\sigma$ and the link is the difference between the closed star and the star:

$$\begin{aligned} \mathrm{Lk}_1\,x &= \overline{\mathrm{St}}\,\tau - \mathrm{St}\,\tau, \\ \mathrm{Lk}_2\,x &= \overline{\mathrm{St}}\,\sigma - \mathrm{St}\,\sigma. \end{aligned}$$

Let $\mathbb{X}$ be the central projection of $|\,\mathrm{Lk}_1\,x\,|$ to $\mathbb{C}^{e-1}$. Since the subdivision operation does not change the geometric neighborhoods of $x \in |\,K\,| = |\,\mathrm{Sd}\,K\,|$, $\mathbb{X}$ is also the central projection of $|\,\mathrm{Lk}_2\,x\,|$ to $\mathbb{C}^{e-1}$. We construct a common subdivision of the two projected links. By projecting this subdivision back to the two links and subdividing the two stars accordingly we get isomorphic subdivisions of $\overline{\mathrm{St}}\,\tau$ and $\overline{\mathrm{St}}\,\sigma$. Property 1 implies $\mathrm{ord}\,\tau = \mathrm{ord}\,\sigma$.

In words, the subdivision operation preserves orders so the i-th boundaries of K and $\mathrm{Sd}\,K$ have the same underlying space. It follows that the i-th boundary of $\mathrm{Sd}\,K$ is a subdivision of the i-th boundary of K. □

Boundary is closed. A fairly straightforward consequence of Property 2 is that faces of a simplex in the boundary also belong to the boundary. Together with the Order Bound this implies that the i-th boundary of a d-complex is a complex of dimension at most $d - i$.

PROPERTY 3. $\mathrm{Bd}_i\,K$ *is a simplicial complex.*

Proof. It suffices to show that the order of a simplex cannot exceed that of its faces:

$$\tau \leq \sigma \quad \Longrightarrow \quad \mathrm{ord}\,\sigma \leq \mathrm{ord}\,\tau.$$

Let $\ell = \mathrm{ord}\,\tau$ and $k = \dim \mathrm{St}\,\tau$. By definition, there is a $(k-\ell)$-simplex η with $\mathrm{St}\,\tau \simeq \mathrm{St}\,\eta$. Let $\mathrm{Sd}\,\overline{\mathrm{St}}\,\tau$ and $\mathrm{Sd}\,\overline{\mathrm{St}}\,\eta$ be isomorphic subdivisions of the two closed stars so the defined simplicial homeomorphism maps $|\,\mathrm{St}\,\tau\,|$ to $|\,\mathrm{St}\,\eta\,|$. Let σ' be a highest-dimensional simplex in $\mathrm{Sd}\,\overline{\mathrm{St}}\,\tau$ with $\mathrm{int}\,\sigma' \subseteq \mathrm{int}\,\sigma$, and let $\xi' \in \mathrm{Sd}\,\overline{\mathrm{St}}\,\eta$ be the isomorphic image of σ'. Finally, let ξ be the simplex in $\mathrm{St}\,\eta$ with $\mathrm{int}\,\xi' \subseteq \mathrm{int}\,\xi$. Clearly, $\dim\sigma = \dim\sigma' = \dim\xi' \leq \dim\xi$. Using Property 2, Property 1, Property 2, and the Order Bound, in this sequence, we get

$$\begin{aligned} \mathrm{ord}\,\sigma &= \mathrm{ord}\,\sigma' \\ &= \mathrm{ord}\,\xi' \\ &= \mathrm{ord}\,\xi \\ &\leq \dim \mathrm{St}\,\xi - \dim\xi. \end{aligned}$$

The result follows because $\dim \mathrm{St}\, \xi \leq k$ and $\dim \xi \geq \dim \eta = k - \ell$, so $\mathrm{ord}\, \sigma \leq \ell = \mathrm{ord}\, \tau$. □

Boundary decreases order. It is intuitively clear that in the i-th boundary the order of a simplex is at least i less than in the original complex. The reason is that the i-th boundary reduces the dimension of the star of any simplex by i. We write $\mathrm{ord}_i\, \sigma$ for the order of σ in $\mathrm{Bd}_i\, K$.

PROPERTY 4. $\mathrm{ord}_i\, \sigma \;\leq\; \mathrm{ord}\, \sigma - i$ *for all simplices* $\sigma \in \mathrm{Bd}_i\, K$.

Proof. We have $\sigma \in \mathrm{Bd}_i\, K$ iff $\ell = \mathrm{ord}\, \sigma \geq i$. Let $k = \dim \mathrm{St}\, \sigma$. By definition of order there is a $(k - \ell)$-simplex η in some complex Y with $\mathrm{St}\, \sigma \simeq \mathrm{St}\, \eta$. Let $\mathrm{Sd}\, \overline{\mathrm{St}}\, \sigma$ and $\mathrm{Sd}\, \overline{\mathrm{St}}\, \eta$ be isomorphic subdivisions of the two closed stars so the defined simplicial homeomorphism maps $\|\mathrm{St}\, \sigma\|$ to $\|\mathrm{St}\, \eta\|$. By Property 1, the orders of corresponding simplices in the two subdivisions are the same. We formally rephrase this observation in the first line below and derive the second using Property 2:

$$\begin{aligned} \mathrm{Bd}_i\, \mathrm{Sd}\, \overline{\mathrm{St}}\, \sigma \quad &\sim \quad \mathrm{Bd}_i\, \mathrm{Sd}\, \overline{\mathrm{St}}\, \eta, \\ \mathrm{Sd}\, \mathrm{Bd}_i\, \overline{\mathrm{St}}\, \sigma \quad &\sim \quad \mathrm{Sd}\, \mathrm{Bd}_i\, \overline{\mathrm{St}}\, \eta. \end{aligned}$$

When we reverse the order of the boundary and the closure operations we get the same results within the stars. This is because the stars of all simplices $\tau \in \mathrm{St}\, \sigma$ are the same within K and within $\mathrm{St}\, \sigma$. Hence τ has the same order in K and in $\overline{\mathrm{St}}\, \sigma$ and belongs to the i-th boundary of K iff it belongs to the i-th boundary of $\overline{\mathrm{St}}\, \sigma$. The same is true for simplices in the star of η. In other words, the star of σ in $\mathrm{Bd}_i\, K$ and the star of η in $\mathrm{Bd}_i\, Y$ are combinatorially equivalent. Using Property 1 again we conclude that the order of σ in the i-th boundary of K is the same as the order of η in the i-th boundary of Y. By the Order Bound, the latter is bounded from above by $\dim \mathrm{Bd}_i\, Y - \dim \eta$. Because $\dim \mathrm{Bd}_i\, Y \leq k - i$ and $\dim \eta = k - \ell$, the order of σ in the i-th boundary is bounded from above by $\ell - i$, and the claim follows. □

Taking boundary simplifies. Taking 1-st boundaries i times does not necessarily produce the same result as taking the i-th boundary once. Indeed, the former operation tends to produce smaller complexes than the latter. The intuitive reason is that taking boundary eliminates context and blurs the topological properties of the remaining neighborhood. We prove a slightly stronger result.

PROPERTY 5. $\mathrm{Bd}_\ell\, \mathrm{Bd}_i\, K \subseteq \mathrm{Bd}_{i+\ell}\, K$.

Proof. We have $\mathrm{Bd}_{i+\ell}\, K \subseteq \mathrm{Bd}_i\, K$. A simplex σ belongs to $\mathrm{Bd}_i\, K - \mathrm{Bd}_{i+\ell}\, K$ iff $i \leq \mathrm{ord}\, \sigma < i + \ell$. By Property 4, the order of σ in $\mathrm{Bd}_i\, K$ is at least i less than in K, which implies $0 \leq \mathrm{ord}_i\, \sigma < \ell$, or equivalently $\sigma \in \mathrm{Bd}_i\, K - \mathrm{Bd}_\ell\, \mathrm{Bd}_i\, K$. In words, if σ in the i-th boundary does not belong to the $(i+\ell)$-th boundary then it also does not belong to the ℓ-th boundary of the i-th boundary. $\square$

Tamal K. Day
Department of Computer Science and Engineering
Indian Institute of Technology
Kharagpur 721302, India

Herbert Edelsbruner
Department of Computer Science
Duke University
Durham, North Carolina 27708
and
Raindrop Geomagic
Research Triangle Park
North Carolina 27709

Sumanta Guha
Department of Electrical Engineering and Computer Science
University of Wisconsin at Milwaukee
Milwaukee
Wisconsin 53201

Dmitry V. Nekhayev
Raindrop Geomagic
Research Triangle Park
North Carolina 27709