Combinatorial Topology and Distributed Computing

Maurice Herlihy Dmitry Kozlov Sergio Rajsbaum

February 22, 2011

Contents

Part I: Undergraduate Course

Chapter 2

Elements of Combinatorial Topology

This chapter defines the basic notions of combinatorial algebraic topology needed to start reading this book. These definitions provide the basic language for describing distributed and concurrent computation. We will introduce more advanced concepts as they are needed.

We use $[m : n]$, where $n \geq m$, as shorthand for $\{m, \ldots, n\}$, and we write $[n]$, as shorthand for $[0 : n]$.

2.1 The objects and the maps

Since simplicial complexes and maps between them are the very central objects of this book, we take a little bit of time first to elaborate on this concept. Fundamentally, there are 3 aspects, and accordingly 3 viewpoints on simplicial complexes: combinatorial, topological, and geometric. We shall now outline each one and connections between them.

2.1.1 The Combinatorial View

We start with the combinatorial notion.

Definition 2.1.1. Given a set S, and a family $\mathcal{A}$ of finite subsets of S. We say that $\mathcal{A}$ is an *abstract simplicial complex* on S, if the following are satisfied:

(1) if $X \in \mathcal{A}$, and $Y \subseteq X$, then also $Y \in \mathcal{A}$;

(2) $\{v\} \in \mathcal{A}$, for all $v \in S$.

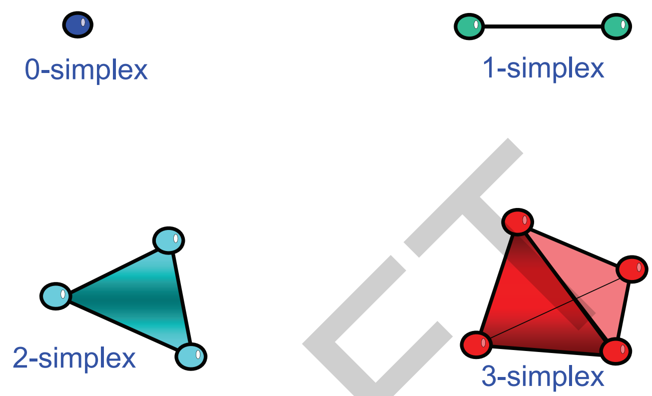

Figure 2.1: Simplices of various dimensions.

Individual elements of S are called *vertices* of $\mathcal{A}$. The set of all vertices of $\mathcal{A}$ is denoted by $V(\mathcal{A})$. More generally, individual elements of $\mathcal{A}$ are called *simplices* of $\mathcal{A}$. A simplex $\sigma \in \mathcal{A}$ is said to have dimension $|\sigma| - 1$. In particular, vertices are 0-dimensional simplices. We sometimes mark a simplex's dimension with a superscript: σ^n. A simplex of dimension n is sometimes called an *n-simplex*.

We usually use vector notation for vertexes $(\vec{x}, \vec{y}, \vec{z}, \ldots)$, Greek letters for simplices $(\sigma, \tau, \ldots)$, and calligraphic font for complexes $(\mathcal{A}, \mathcal{B}, \ldots)$.

A simplex τ is a *face* of σ if $\tau \subseteq \sigma$, and a *proper face* if $\tau \subset \sigma$. If τ has dimension k, the τ is a *k-face* of σ. Let $\sigma = (\vec{s}_0, \ldots \vec{s}_n)$ Define $\text{Face}_i\, \sigma^n$, the i^{th} *face* of σ^n, to be the $(n-1)$-simplex $(\vec{s}_0, \ldots, \hat{s}_i, \ldots, \vec{s}_n)$, where the circumflex denotes omission. If U is a set of processes, define $\text{Face}_U\, \sigma^n = \cap_{i \in U} \text{Face}_i\, \sigma^n$.

maurice: Added Face_U definition. Must check that we use it

A simplex σ in a complex $\mathcal{C}$ is a *facet* if it is maximal in $\mathcal{C}$: σ is not a proper face of any other simplex. The *dimension* of a complex $\mathcal{C}$ is the maximum dimension of any of its facets. A complex is *pure* if all facets have the same dimension. A complex $\mathcal{D}$ is a *subcomplex* of $\mathcal{C}$ if every simplex of $\mathcal{D}$ is also a simplex of $\mathcal{C}$. If $\mathcal{C}$ is a pure complex, the *codimension* of $\sigma \in \mathcal{C}$ is $\dim \mathcal{C} - \dim \sigma$, in particular, any facet has codimension 0. When $\mathcal{C}$ is clear from context, we denote the codimension by $\operatorname{codim} \sigma$.

2.1.2 The Geometric View

We next switch to geometry. We let $\mathbb{R}^d$ denote the d-dimensional Euclidean space. The *standard n-simplex* is the convex hull of the $n+1$ points in $\mathbb{R}^{n+1}$ with coordinates $(1, 0, \ldots, 0)$, $(0, 1, 0, \ldots, 0)$, $\ldots$, $(0, \ldots, 0, 1)$. More generally, a *geometric n-simplex*, or a *geometric simplex of dimension n*, is the convex hull of any set of $n+1$ affinely independent points in $\mathbb{R}^d$, in particular, we must have $d \geq n$. As illustrated in Fig. 2.1, a 0-dimensional simplex is a point, a 1-simplex is an edge linking two points, a 2-simplex is a solid triangle, a 3-simplex a solid tetrahedron, and so on.

In direct analogy with the combinatorial framework we have the following terminology. When $v_0, \ldots, v_n \in \mathbb{R}^d$ are affinely independent, we call them *vertices* of the n-simplex $\sigma = \operatorname{conv}\{v_0, \ldots, v_n\}$. In this case, for any $S \subseteq [n]$, the $(|S|-1)$-simplex $\tau = \operatorname{conv}\{v_s \mid s \in S\}$ is called a *face*, or an $(|S|-1)$-face of σ; it is called a *proper face* if, in addition, $S \neq [n]$. We set $\operatorname{Face}_i \sigma^n := \operatorname{conv}\{v_0, \ldots, \hat{v}_i, \ldots, v_n\}$. Gluing geometric simplices together, along their faces, yields the geometric analog of Definition 2.1.1.

Definition 2.1.2. A *geometric simplicial complex* $\mathcal{K}$ in $\mathbb{R}^d$ is a collection of of geometric simplices, such that

(1) any face of a $\sigma \in \mathcal{K}$ is again in $\mathcal{K}$;

(2) for all $\sigma, \tau \in \mathcal{K}$, their intersection $\sigma \cap \tau$ is a face of each of them.

Other notions, such as codimension and subcomplex are the same as in the combinatorial case.

Given a geometric simplicial complex $\mathcal{K}$, we can define the underlying abstract simplicial complex $\mathcal{C}(\mathcal{K})$ as follows: take the union of all the sets of vertices of the simplices of $\mathcal{K}$ as the vertices of $\mathcal{C}(\mathcal{K})$, then for each simplex $\sigma = \operatorname{conv}\{v_0, \ldots, v_n\}$ of $\mathcal{K}$ take the set $\{v_0, \ldots, v_n\}$ to be the simplex of $\mathcal{C}(\mathcal{K})$. In the opposite direction: given an abstract simplicial complex $\mathcal{A}$ with finitely many vertices, there exist many geometric simplicial complexes $\mathcal{K}$, such that $\mathcal{C}(\mathcal{K}) = \mathcal{A}$. The simplest construction is as follows: assume $\mathcal{A}$ has d vertices, take the standard simplex σ in $\mathbb{R}^d$, and take the subcomplex of σ consisting of the geometric simplices which correspond to the sets in the set family $\mathcal{A}$. Usually, one can find $\mathcal{K}$ of a much lower dimension than d, but then the construction could be quite a bit more complicated.

2.1.3 The Topological View

Finally, we proceed to the topological framework. Given a geometric simplicial complex $\mathcal{K}$ in $\mathbb{R}^d$, we let $|\mathcal{K}|$ denote the union of its simplices. This space has the "usual" topology as the subspace of $\mathbb{R}^d$ (in point-set topological language one would say that the topology on $|\mathcal{K}|$ is induced by the standard topology on $\mathbb{R}^d$). Somewhat confusingly the space $|\mathcal{K}|$ is called the *geometric realization* of $\mathcal{K}$. If $\mathcal{A}$ is an abstract simplicial complex, we can first construct $\mathcal{K}$, such that $\mathcal{C}(\mathcal{K}) = \mathcal{A}$, and then let $|\mathcal{A}| = |\mathcal{K}|$. This will not depend on the particular geometric simplicial complex $\mathcal{K}$, only on the abstract simplicial complex $\mathcal{A}$. One can also construct $|\mathcal{A}|$ by starting with a set of disjoint simplices, and then gluing them together along their boundaries, using the combinatorial data as the gluing schema.

dmitry: Do we want more details here? quotient spaces?

Let us now look at the maps between the objects which we just described. Let $\mathcal{A}$ and $\mathcal{B}$ be abstract simplicial complexes. A *vertex map* $\mu : \mathcal{A} \to \mathcal{B}$ maps each vertex of $\mathcal{A}$ to a vertex of $\mathcal{B}$. The vertex map μ is a *simplicial map* if it also carries simplices to simplices. A vertex map $\mu : \mathcal{A} \to \mathcal{B}$ does not have to induce a continuous map between the geometric realizations $|\mathcal{A}|$ and $|\mathcal{B}|$. For example, if both $\mathcal{A}$ and $\mathcal{B}$ have the vertex set $\{1, 2\}$, and the set $\{1, 2\}$ is a simplex of $\mathcal{A}$, but not of $\mathcal{B}$, then the identity map $\text{id} : \{1, 2\} \to \{1, 2\}$ is a vertex map, but there does not exist a continuous map from an interval to the set of its endpoints, which is identity on the endpoints. However, any simplicial map μ will automatically induce a continuous map $|\mu|$ between geometric resolution. The easiest way to describe $|\mu|$ is by using *barycentric coordinates.*

dmitry: do we want more details on that?

Before proceeding with constructions, we would like to mention that in standard use in algebraic topology the word *simplex* is overloaded. It is used to denote the abstract simplicial complex consisting of *all* subsets of a certain finite set, but it is also used to refer to individual elements of the family of sets constituting and abstract simplicial complex. There is a relation here: to a simplex in the second sense one can associate a subcomplex of the considered abstract simplicial complex which is a simplex in the first sense. We will use simplex in both of these meanings. In some texts, simplex is also used to denote the geometric realization of that abstract simplicial complex; here we shall say *geometric simplex* instead.

2.2 Standard constructions

maurice: Should we define subdivisions here? Carriers?

Let $\mathcal{C}$ be an abstract simplicial complex or a geometric complex. Let ℓ be a nonnegative integer. The set of simplices of $\mathcal{C}$ of dimension at most ℓ is a subcomplex of $\mathcal{C}$, called the *ℓ-skeleton*, and denoted $\mathrm{skel}^{\ell}\,\mathcal{C}$. In particular, the 0^{th} skeleton of a simplicial complex is simply its set of vertices. If σ is an n-simplex, then $\mathrm{skel}^{n-1}\,\sigma$ is called the *boundary complex* of σ, it is obtained from σ by deleting the single facet.

Assime now σ is some simplex of a simplicial complex $\mathcal{C}$. There are three standard constructions, each yielding a subcomplex of $\mathcal{C}$: the star, the link and the deletion. We describe this in the abstract simplicial case, the case of geometric simplicial complexes is completely analogous.

Star. The *star* of a simplex $\sigma \in \mathcal{C}$, written $\mathrm{St}(\sigma, \mathcal{C})$ (or $\mathrm{St}\,\sigma$ when $\mathcal{C}$ is clear from context), is the subcomplex of $\mathcal{C}$ whose facets are the simplices of $\mathcal{C}$ that contain σ. The complex $\mathrm{St}(\sigma, \mathcal{C})$ consists of all the simplices τ which contain σ, and furthermore, all the simplices contained in such a simplex τ. The geometric realization of $\mathrm{St}(\sigma, \mathcal{C})$ is also called the star of σ.

In the topological context there is also the notion of an *open star*, denoted $\mathrm{Ost}(\vec{v})$, which is the union of the interiors of all the simplices which contain σ, (see Fig. 2.2). Note, that the open star is not an abstract or geometric simplicial complex, but just a topological space. To distinguish the two notions, the star is also sometimes called the *closed star*.

Link. The *link* of $\sigma \in \mathcal{C}$, written $\mathrm{Lk}(\sigma, \mathcal{C})$ (or $\mathrm{Lk}\,\sigma$), is the subcomplex of $\mathcal{C}$ consisting of all simplices in $\mathrm{St}(\sigma, \mathcal{C})$ that do not have common vertices with σ (see Fig. 2.2. Again, the geometric realization of $\mathrm{Lk}(\sigma, \mathcal{C})$ is also called the link of σ.

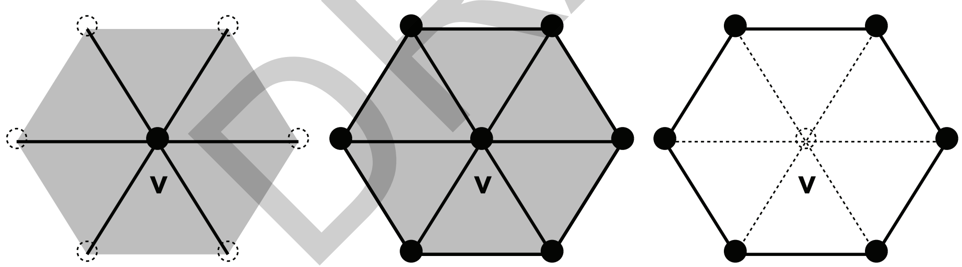

Figure 2.2: The open star $\mathrm{Ost}(\vec{v})$, star $\mathrm{St}(\vec{v})$, and link $\mathrm{Lk}(\vec{v})$ of the vertex v.

dmitry: I could make a picture with examples where σ has higher dimension

Deletion. The *deletion* of $\sigma \in \mathcal{C}$, written $\mathrm{dl}(\sigma, \mathcal{C})$, is the subcomplex of $\mathcal{C}$ consisting of all simplices of $\mathcal{C}$ that do not have common vertices with σ. Clearly, we have the relations:

$$\mathrm{Lk}(\sigma, \mathcal{C}) = \mathrm{dl}(\sigma, \mathcal{C}) \cap \mathrm{St}(\sigma, \mathcal{C}),$$

and

$$\mathcal{C} = \mathrm{dl}(\sigma, \mathcal{C}) \cap \mathrm{St}(\sigma, \mathcal{C}),$$

for all $\sigma \in \mathcal{C}$.

Given two abstract simplicial complexes $\mathcal{A}$ and $\mathcal{B}$ with disjoint sets of vertices $S_{\mathcal{A}}$ and $S_{\mathcal{B}}$, their *join*, or *simplicial join*, $\mathcal{A} * \mathcal{B}$ is the abstract simplicial complex with the set of vertices $S_{\mathcal{A}} \cup S_{\mathcal{B}}$, whose simplices are all the unions $\alpha \cup \beta$, where $\alpha \in \mathcal{A}$, and $\beta \in \mathcal{B}$. Note that it is allowed to take α or β to be empty sets. In particular, both $\mathcal{A}$ and $\mathcal{B}$ are subcomplexes of $\mathcal{A} * \mathcal{B}$.

Assume furthermore, that $\mathcal{K}$ is a geometric simplicial complex in $\mathbb{R}^m$, such that $\mathcal{C}(\mathcal{K}) = \mathcal{A}$, and $\mathcal{L}$ is a geometric simplicial complex in $\mathbb{R}^n$, such that $\mathcal{C}(\mathcal{L}) = \mathcal{B}$. Then there is a standard way to construct a geometric simplicial complex in $\mathbb{R}^{m+n+1}$ whose underlying abstract simplicial complex is $\mathcal{A} * \mathcal{B}$. Consider the following embeddings: $\varphi : \mathbb{R}^m \to \mathbb{R}^{m+n+1}$, given by

$$\varphi(x_1, \dots, x_m) = (x_1, \dots, x_m, 0, \dots, 0),$$

and $\psi : \mathbb{R}^n \to \mathbb{R}^{m+n+1}$ given by

$$\psi(y_1, \dots, y_n) = (0, \dots, 0, y_1, \dots, y_n, 1).$$

The images under these embeddings of $\mathcal{K}$ and $\mathcal{L}$ are geometric simplicial complexes whose geometric realizations are disjoint. We can define a new geometric simplicial complex $\mathcal{K} * \mathcal{L}$ by taking all convex hulls $\operatorname{conv}(\sigma, \tau)$, where σ is a simplex of $\mathcal{K}$, and τ is a simplex of $\mathcal{L}$. It is a matter of simple linear algebra to show that the open intervals (x, y), where $x \in \operatorname{Im}\varphi$ and $y \in \operatorname{Im}\psi$, never intersect, and so $\mathcal{K} * \mathcal{L}$ satisfies the conditions for the geometric simplicial complex. It is easy to see that the topological spaces $|\mathcal{A} * \mathcal{B}|$ and $|\mathcal{K} * \mathcal{L}|$ are homeomorphic.

dmitry: a picture of the join of two intervals giving a tetrahedron would make sense here

As an example, taking join of $\mathcal{K}$ with a single vertex amounts to building a cone over $\mathcal{K}$ for any complex $\mathcal{K}$. Another example is taking join of an m-simplex with an n-simplex, the result of that is an $(m+n+1)$-simplex. One can also show that the join operation is commutative and associative.

There is a purely topological definition of the join of two topological spaces. We refer to [21] for the precise definition. Here we just mention that the simplicial and topological joins commute with the geometric realization, that is for any two abstract simplicial complexes $\mathcal{A}$ and $\mathcal{B}$, the spaces $|\mathcal{A} * \mathcal{B}|$ and $|\mathcal{A}| * |\mathcal{B}|$ are homeomorphic.

dmitry: I am guessing this is enough

Given two abstract simplicial complexes $\mathcal{A}$ and $\mathcal{B}$, a *carrier map* $\mathcal{M}$ from $\mathcal{A}$ to $\mathcal{B}$ maps each simplex $\sigma \in \mathcal{A}$ to a subcomplex $\mathcal{M}(\sigma)$ of $\mathcal{B}$, so that

$$\mathcal{M}(\sigma) \cap \mathcal{M}(\tau) = \mathcal{M}(\sigma \cap \tau), \tag{2.2.1}$$

for all $\sigma, \tau \in \mathcal{A}$. In other words, carrier maps preserve intersections. Consider the special case when σ is a face of τ, then $\sigma \cap \tau = \sigma$, so (2.2.1) yields $\mathcal{M}(\sigma) \subseteq \mathcal{M}(\tau)$. This means that in particular $\mathcal{M}$ is *order preserving*, i.e., the inclusion pattern of the subcomplexes $\mathcal{M}(\sigma)$ should be the same as the inclusion pattern of the simplices of $\mathcal{A}$.

Since the carrier map takes simplices of $\mathcal{A}$ to the subcomplexes of $\mathcal{B}$ we adopt the "power notation" for writing it, i.e., we write $\mathcal{M} : \mathcal{A} \to 2^{\mathcal{B}}$.

maurice: propagate power notion to later chapters

2.3 Chromatic complexes

An *m-labeling*, or just a *labeling*, of an abstract simplicial, or a geometric simplicial complex $\mathcal{A}$ is a map carrying each vertex of $\mathcal{A}$ to an element of some domain of cardinality m. In other words, it is a set map $\varphi : V(\mathcal{A}) \to D$, where $|D| = m$.

An *m-coloring*, or just a *coloring*, of a n-dimensional complex $\mathcal{A}$ is an m-labeling $\chi : V(\mathcal{A}) \to \Pi$, such that restricted to the vertices of any simplex of $\mathcal{A}$, the map χ is injective. A simplicial complex $\mathcal{A}$ together with a coloring χ is called a *chromatic* complex. We shall write it as a pair $(\mathcal{A}, \chi)$.

maurice: Do we/should we use "rigid" in later chapters?

Definition 2.3.1. Given two abstract simplicial complexes $\mathcal{A}$ and $\mathcal{B}$. A simplicial map $\varphi : \mathcal{A} \to \mathcal{B}$ is called *rigid*, if the image of each simplex σ has the same dimension as σ, i.e., $|\varphi(\sigma)| = |\sigma|$.

Rigid maps are much more rare than simplicial maps. For example, for any pair of abstract simplicial complexes $\mathcal{A}$ and $\mathcal{B}$, there will be many simplicial maps from $\mathcal{A}$ to $\mathcal{B}$; at the very least we could simply map everything to a single vertex of $\mathcal{B}$. This is not allowed for the rigid maps, and there are situations where there are no rigid simplicial maps between given abstract simplicial complexes at all. For example, this is the case when $\mathcal{A}$ is the boundary of a triangle, and $\mathcal{B}$ is a single interval. We note that a composition of rigid simplicial maps is again a rigid simplicial map.

We see that requiring that a labeling χ of an abstract simplicial complex $\mathcal{A}$ is injective on the simplices is the same as asking that $|\chi(\sigma)| = |\sigma|$, for all $\sigma \in \mathcal{A}$. In other words, m-colorings are precisely the rigid simplicial maps into an $(m-1)$-simplex $\chi : \mathcal{A} \to \Delta^{m-1}$.

From the point of view of graph theory, we remark that a coloring $\chi : \mathcal{A} \to \Delta^{m-1}$ exists if and only if the 1-skeleton of $\mathcal{A}$, viewed as a graph, is m-colorable in the sense of graph colorings (more precisely vertex-colorings of graphs).

Definition 2.3.2. Given two m-chromatic simplicial complexes $(\mathcal{A}, \chi_{\mathcal{A}})$ and $(\mathcal{B}, \chi_{\mathcal{B}})$, a simplicial map $f : \mathcal{A} \to \mathcal{B}$ is called *color-preserving* if the following diagram commutes: In other words, for every vertex $\vec{v} \in \mathcal{A}$, we

Figure 2.3: Definition of color-preserving simplicial maps.

have $\chi_{\mathcal{A}}(\vec{v}) = \chi_{\mathcal{B}}(f(\vec{v}))$.

maurice: We will need to be consistent about "rigid" versus "color-preserving" versus "chromatic" maps

We define two further rigid notions.

Definition 2.3.3. Given two abstract simplicial complexes $\mathcal{A}$ and $\mathcal{B}$, and a carrier map $\mu : \mathcal{A} \to 2^{\mathcal{B}}$. We call μ a *rigid carrier map* if for every n, and for any n-simplex $\sigma \in \mathcal{A}$, the subcomplex $\mu(\sigma)$ is pure and n-dimensional.

Definition 2.3.4. A chromatic abstract simplicial complex $(\mathcal{A}, \chi)$ is called *rigid chromatic* $\mathcal{A}$ is pure of dimension n, and χ is an $(n+1)$-coloring.

2.4 Simplicial models in Distributed Computing

In recent breakthrough developments in distributed computing, it was discovered that protocols and tasks can be modeled using the language of combinatorial topology.

In this language, a task will consist of an input complex $\mathcal{I}$, an output complex $\mathcal{O}$, and a task specification map μ. Mathematically, both the input complex and the output complex are pure rigid chromatic abstract simplicial complexes of dimension $N-1$, where N is the number of processors. We denote both their coloring functions with id. A task specification map is a rigid carrier map $\mu : \mathcal{I} \to 2^{\mathcal{O}}$.

The connection with the distributed computing is as follows. The target set of the coloring function id is the set of process ID's. Vertices of the complexes $\mathcal{I}$ and $\mathcal{O}$ are pairs consisting of a process ID and a process state, and each k-dimensional *simplex* σ is a set of $k+1$ vertexes labeled with distinct process IDs.

Input complex. Each possible initial state of the system is given by a maximal *input simplex*, assigning an input value to each process. Together, all possible input simplices make up the task's *input complex*.

Output complex. An *output simplex* and the task's *output complex* are defined similarly.

Task specification map. A rigid carrier map μ takes each input simplex σ to the pure subcomplex $\mu(\sigma)$, whose maximal simplices are precisely all global legal final states for the input σ.

Protocol complex. Any protocol has an associated *protocol complex* $\mathcal{P}$, in which each vertex is labeled with a process id and that process's final state (called its *view*). Each simplex thus corresponds to an equivalence class of executions that "look the same" to the processes at its vertexes. For $0 \leq m \leq n$, for a given input m-simplex σ, we understand $\mathcal{P}(\sigma)$ to be the complex generated by all executions starting in σ, in which only the processes in $\text{id}(\sigma)$ take part, while the rest fail before taking any steps. Formally, the protocol complex is a pure rigid chromatic abstract simplicial complex of dimension $N-1$.

It is important to note that while input and output complexes, as well as the task specification carrier map, depend on the task only, the protocol complex crucially depends on the choice of the model of computation.

2.5 Chapter Notes

Most of these definitions are adapted from Munkres [24] and Kozlov [21]. The use of simplicial complexes to model distributed computing model is due to Herlihy and Shavit [18].

2.6 Exercises

[[[[no exercises yet]]]]

Bibliography

[1] Yehuda Afek, Hagit Attiya, Danny Dolev, Eli Gafni, Michael Merritt, and Nir Shavit. Atomic snapshots of shared memory. *Journal of the ACM*, 40(4):873–890, 1993.

[2] Hagit Attiya, Amotz Bar-Noy, Danny Dolev, David Peleg, and Rudiger Reischuk. Renaming in an Asynchronous Environment. *Journal of the ACM*, July 1990.

[3] Hagit Attiya, Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Bounds on the time to reach agreement in the presence of timing uncertainty. *J. ACM*, 41:122–152, January 1994.

[4] Hagit Attiya and Jennifer Welch. *Distributed Computing Fundamentals, Simulations, and Advanced Topics Second Edition*.

[5] Ofer Biran, Shlomo Moran, and Shmuel Zaks. A combinatorial characterization of the distributed tasks which are solvable in the presence of one faulty processor. In *PODC '88: Proceedings of the seventh annual ACM Symposium on Principles of distributed computing*, pages 263–275, New York, NY, USA, 1988. ACM.

[6] Elizabeth Borowsky and Eli Gafni. Generalized FLP impossibility result for t-resilient asynchronous computations. In *STOC '93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 91–100, New York, NY, USA, 1993. ACM.

[7] Elizabeth Borowsky and Eli Gafni. A Simple Algorithmically Reasoned Characterization of Wait-Free Computations (Extended Abstract). In *PODC '97: Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing*, pages 189–198, New York, NY, USA, 1997. ACM.

[8] S. Chaudhuri. Agreement Is Harder Than Consensus: Set Consensus Problems in totally asynchronous systems. In *Proceedings Of The Ninth Annual ACM Symosium On Principles of Distributed Computing*, pages 311–234, August 1990.

[9] Soma Chaudhuri, Maurice Herlihy, Nancy A. Lynch, and Mark R. Tuttle. A Tight Lower Bound for k-Set Agreement. In *In Proceedings of the 34th IEEE Symposium on Foundations of Computer Science*, pages 206–215, 1993.

[10] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Andreas Tielmann. The Disagreement Power of an Adversary. In Idit Keidar, editor, *Distributed Computing*, volume 5805 of *Lecture Notes in Computer Science*, chapter 6, pages 8–21. Springer Berlin / Heidelberg, Berlin, Heidelberg, 2009.

[11] M. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility Of Distributed Commit With One Faulty Process. *Journal of the ACM*, 32(2), April 1985.

[12] Eli Gafni and Elias Koutsoupias. Three-Processor Tasks Are Undecidable. *SIAM J. Comput.*, 28(3):970–983, 1999.

[13] Eli Gafni, Sergio Rajsbaum, and Maurice Herlihy. Subconsensus Tasks: Renaming Is Weaker Than Set Agreement. In *Distributed Computing, 20th International Symposium, Stockholm, Sweden, September 18-20, 2006, Proceedings(DISC)*, volume 4167 of *Lecture Notes in Computer Science*, pages 329–338. Springer, 2006.

[14] Maurice Herlihy and Sergio Rajsbaum. The decidability of distributed decision tasks (extended abstract). In *STOC '97: Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 589–598, New York, NY, USA, 1997. ACM.

[15] Maurice Herlihy and Sergio Rajsbaum. The topology of shared-memory adversaries. In *Proceeding of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, PODC '10, pages 105–113, New York, NY, USA, 2010. ACM.

[16] Maurice Herlihy, Sergio Rajsbaum, and Mark Tuttle. An Axiomatic Approach to Computing the Connectivity of Synchronous and Asynchronous Systems. *Electron. Notes Theor. Comput. Sci.*, 230:79–102, 2009.

[17] Maurice Herlihy, Sergio Rajsbaum, and Mark R. Tuttle. Unifying synchronous and asynchronous message-passing models. In *PODC '98: Proceedings of the seventeenth annual ACM symposium on Principles of distributed computing*, pages 133–142, New York, NY, USA, 1998. ACM.

[18] Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(6):858–923, 1999.

[19] Flavio Junqueira and Keith Marzullo. A framework for the design of dependent-failure algorithms: Research Articles. *Concurr. Comput. : Pract. Exper.*, 19(17):2255–2269, 2007.

[20] Flavio P. Junqueira and Keith Marzullo. Designing Algorithms for Dependent Process Failures. Technical report, 2003.

[21] Dimitry Kozlov. *Combinatorial Algebraic Topology*, volume 21 of *Algorithms and Computation in Mathematics*. Springer, 1 edition, October 2007.

[22] M. C. Loui and H. H. Abu-Amara. *Memory requirements for agreement among unreliable asynchronous processes*, volume 4, pages 163–183. JAI press, 1987.

[23] Yoram Moses and Sergio Rajsbaum. A Layered Analysis of Consensus. *SIAM J. Comput.*, 31:989–1021, April 2002.

[24] James Munkres. *Elements of Algebraic Topology*. Prentice Hall, 2 edition, January 1984.

[25] Michael Saks and Fotios Zaharoglou. Wait-Free k-Set Agreement is Impossible: The Topology of Public Knowledge. *SIAM Journal on Computing*, 29(5):1449–1483, 2000.

[26] Michael Saks and Fotios Zaharoglou. Wait-Free k-Set Agreement is Impossible: The Topology of Public Knowledge. *SIAM J. Comput.*, 29(5):1449–1483, 2000.

[27] Francis Sergeraert. The Computability Problem In Algebraic Topology. *Adv. Math*, 104:1–29, 1994.

[28] Edwin H. Spanier. *Algebraic topology*. Springer-Verlag, New York, 1981.

[29] John Stillwell. *Classical Topology and Combinatorial Group Theory*. Springer, 2nd edition, March 1993.