Notes on Wait-Free Spans

John Havlicek
University of Texas at Austin
havlicek@cs.utexas.edu

September 22, 1998

0. Introduction

A number of key results in the theory of asynchronous, fault-tolerant computation of generalized consensus problems (known as *decision tasks*) are, at heart, topological. There is a developing body of research in this area that exploits the language and tools of combinatorial and algebraic topology.

The seminal papers [FLP, BMZ] assume an asynchronous message-passing model and require that protocols tolerate a single crash failure. In this context, the geometry of the tasks and protocols is essentially one-dimensional and can be adequately expressed by graphs, such as the "adjacency graphs" in [BMZ]. The important topological property for understanding computability in this setting is connectedness.

More recent work [B, BG1-3, GK, HR1-3, HS1-3, SZ] assumes an asynchronous read-write shared-memory model, with some form of atomic snapshot reads, and expects tolerance of more than just a single crash failure. A protocol that solves a decision task regardless of any pattern of crash failures by up to f processes is called *f-resilient*, and a protocol that is n-resilient in a system of $n+1$ processes is called *wait-free*.[1] By using Borowsky-Gafni simulation [BG2, B], the question of f-resilient solvability of a decision task can in many cases be reduced to the question of wait-free solvability of essentially the same task in a smaller system. Thus, the theory of wait-free computation is of considerable significance.

Herlihy and Shavit [HS1, HS2] introduced a framework for describing and reasoning about solvability of decision tasks that uses simplicial complexes from combinatorial and algebraic topology. A decision task is specified by a triple $(\mathcal{I}, \mathcal{O}, \Delta)$, where $\mathcal{I}$ and $\mathcal{O}$ are the complexes of inputs and outputs, respectively, and Δ is the relation that associates to each input configuration the collection of output configurations that are satisfactory under the task. Strictly speaking, $\mathcal{I}$ and $\mathcal{O}$ are combinatorial objects, but they are naturally associated with topological spaces $|\mathcal{I}|$ and $|\mathcal{O}|$ of dimension one less than the number of processes in the system. The complexes are "dual" to the adjacency graphs in [BMZ] and provide the flexibility

[1] Due to asynchrony, a crash is indistinguishable from a severe delay, so a protocol tolerating failure by up to all but one process is "wait-free" in the sense that a non-faulty process can and will finish without waiting or blocking due to the delay or failure of other processes.

to represent configurations of any positive number of processes.

The Asynchronous Computability Theorem (ACT) of Herlihy and Shavit [HS2] shows that simplicial complexes also provide the "right" topology for studying wait-free solvability of decision tasks. Loosely, the theorem says that $(\mathcal{I}, \mathcal{O}, \Delta)$ admits wait-free solution if and only if there is a suitable continuous map $f \colon |\mathcal{I}| \to |\mathcal{O}|$ that respects the relation Δ. In slightly more precise terms, the necessary and sufficient condition is that there exist a subdivision $\chi(\mathcal{I})$ that respects process identifiers and a simplicial map $\mu \colon \chi(\mathcal{I}) \to \mathcal{O}$ that respects both process identifiers and the relation Δ. The map f can be taken as $|\mu|$ when $|\chi(\mathcal{I})|$ is identified with $|\mathcal{I}|$.[2] While connectedness is still important in understanding computability, the existence of f also depends on higher-dimensional topological properties.

In order to prove the ACT, Herlihy and Shavit introduce an auxiliary complex, the *protocol complex*, whose topology captures the capabilities of the protocol. The protocol complex $\mathcal{P}(\mathcal{I})$ is the simplicial complex of configurations of final views of finishing processes in all possible executions of the protocol with input configurations from $\mathcal{I}$. If the protocol solves the task, then there is a simplicial map $\delta \colon \mathcal{P}(\mathcal{I}) \to \mathcal{O}$ that respects process identifiers and the task relation. The map δ simply maps finishing processes to their output values. The proof of the necessary condition of the ACT assumes, without loss of generality, that the protocol records full information. It is shown that, for a fine enough subdivision $\chi(\mathcal{I})$, there exists a simplicial map $\varphi \colon \chi(\mathcal{I}) \to \mathcal{P}(\mathcal{I})$ that respects both process identifiers and, in an appropriate fashion, carriers of the subdivision. Herlihy and Shavit call such a map φ a *span*. The map μ can then be obtained as the composition $\delta \circ \varphi$. Thus, the necessary condition of the ACT is a corollary of the existence of spans.

The proof of sufficiency relies essentially on Borowsky and Gafni's clever Participating Set Protocol [BG1] for solving the *simplex agreement task*. This protocol serves as a "universal" protocol for wait-free decision tasks, and the existence of the map μ allows specialization of the universal protocol to the task at hand.

The necessary condition of the ACT can be used to prove impossibility results for solution of Chaudhuri's *set consensus task* [Ch] and Attiya et al.'s *renaming task* [At+]. Such proofs can be found in [HS1, HS3], although it should be noted that independent proofs of impossibility results for solution of the set consensus task appear in [BG2, SZ].

These notes prove the existence of spans for wait-free full-information protocols and apply spans to obtain impossibility results for solution of the set consensus task. They are intended to make the theory accessible to the motivated reader having no prior experience with combinatorial or algebraic topology. The proofs are based on the exposition in [HS1, HS2, HS3] and are the result of my effort to understand these papers. I have modified much of the notation and made great efforts to improve the clarity of the arguments. In order to estimate my success, the reader should make an earnest effort to understand [HS1, HS2, HS3] before studying the notes!

Section 1 is a leisurely introduction to some topology that should give the reader enough context to understand the simplicial complexes used to reason about computability. Homotopy is introduced because the proof of the existence of a span uses homotopy criteria to show that the span can be built up inductively. The details

[2] A simplicial map $\varphi \colon \mathcal{K} \to \mathcal{L}$ is associated naturally to a continuous, piecewise-linear map $|\varphi| \colon |\mathcal{K}| \to |\mathcal{L}|$.

of polyhedra and subdivisions can be skipped over as long as the reader gains an intuitive grasp of the concepts, including the notion of carrier.

Section 2 sets down the details of the computation model and the decision tasks. The simplicial complexes associated to tasks and protocols are introduced and some of their elementary properties are proved. This section is probably much too long, but I have tried to leave no gaps in the presentation.

Section 3 covers the heart of the matter: the proof of the existence of spans and application to the set consensus task. Subsections 3.1 through 3.5 tell most of the story. The presentation is simplified by ignoring the chromatic requirement that spans respect process identifiers, and it turns out that non-chromatic spans are enough for the applications to set consensus. Subsections 3.6 through 3.9 cover details to arrange that the spans respect process identifiers. This latter material can safely be omitted on first reading. It would not be unreasonable to claim that Theorem 3.8.2 is original. Subsection 3.8 clarifies and proves a claim whose treatment in [HS1] is confusing, with only brief justification, and that has, to my knowledge, been addressed nowhere else.[3]

1. Some Topology

These notes are intended to be accessible to a motivated reader with little or no experience in topology. This section introduces in a leisurely way some of the topological material that will be used. There are many examples and exercises to help the reader build familiarity and intuition. We do not deal with issues of general topology. Rather, we regard all topological spaces as subsets of $\mathbb{R}^d$, for various d, topologized with the Euclidean metric. Most of the facts imported from topology are cited without proof. The reader wishing to move on quickly can skip all the proofs and read Subsections 1.3, 1.5, and 1.7 lightly.

1.1 Continuity and homeomorphism; disks and spheres.

All topological spaces will be understood to be subsets of $\mathbb{R}^d$, for various d, with topologies induced by the Euclidean metric (i.e., distance function). We write $||x||$ for the *norm* of a point $x \in \mathbb{R}^d$. If $x = (x_1, \ldots, x_d)$, then

$$||x|| = \left(\sum_{i=1}^{d} x_i^2 \right)^{1/2} .$$

The *Euclidean metric* is the distance function

$$\text{dist}(x, y) = ||x - y|| ,$$

which should be familiar from multi-variable calculus.

Let $X \subseteq \mathbb{R}^m$ and $Y \subseteq \mathbb{R}^n$. Continuity of a function $f \colon X \to Y$ is defined as in calculus: f is *continuous* if and only if

$$\lim_{x \to a} f(x) = f(a) \qquad (*)$$

[3] See Section 4 of [HS1], in particular Definition 4.1 and the wait-free case of Theorem 4.8.

for every $a \in X$. Explicitly, $(*)$ means that for any real number $\varepsilon > 0$, there exists a real number $\delta > 0$ such that

$$x \in X \text{ and } ||x - a|| \leq \delta \text{ imply } ||f(x) - f(a)|| \leq \varepsilon \ .$$

The first norm is in $\mathbb{R}^m$, while the second is in $\mathbb{R}^n$. Continuity is usually the minimum requirement for a function to be of topological interest. Continuous functions are often referred to as *continuous maps*, or simply *maps*. Compositions of maps are maps, and the identity function from a space X to itself is a map.

A map $f: X \to Y$ is called a *homeomorphism* if (1) f is bijective, and (2) the inverse $f^{-1}: Y \to X$ is also a map. In this case X and Y are said to be *homeomorphic*. It is easy to see that homeomorphism defines an equivalence relation. Homeomorphism is the standard notion of "topological equivalence."

EXAMPLE 1.1.1: For $a \in \mathbb{R}^d$ and $r > 0$, let

$$B_r^o(a, \mathbb{R}^d) = \{x \in \mathbb{R}^d : ||x - a|| < r\} \ .$$

$B_r^o(a, \mathbb{R}^d)$ is called the *open disk* in $\mathbb{R}^d$ with center a and radius r. (If $\mathbb{R}^d$ is understood, it may be dropped from the notation.) Assume also that $a' \in \mathbb{R}^d$ and $r' > 0$. Consider the function

$$f: B_r^o(a, \mathbb{R}^d) \to B_{r'}^o(a', \mathbb{R}^d)$$

defined by

$$f(x) = (r'/r)(x - a) + a' \ . \tag{$\dagger$}$$

The reader can easily check that f is a homeomorphism. This shows that any two open disks in $\mathbb{R}^d$ are homeomorphic. It is a much more subtle fact, sometimes referred to as the *topological invariance of dimension*, that for $d \neq d'$, an open disk in $\mathbb{R}^d$ is not homeomorphic to any open disk in $\mathbb{R}^{d'}$. □

EXAMPLE 1.1.2: One may also consider *closed disks* in $\mathbb{R}^d$: for $a \in \mathbb{R}^d$ and $r > 0$, let

$$B_r(a, \mathbb{R}^d) = \{x \in \mathbb{R}^d : ||x - a|| \leq r\} \ .$$

Using formula ($\dagger$), one sees that any two closed disks in $\mathbb{R}^d$ are homeomorphic. It is also true that if $d \neq d'$, then $B_r(a, \mathbb{R}^d)$ is not homeomorphic to $B_{r'}(a', \mathbb{R}^{d'})$.

Let

$$S_r(a, \mathbb{R}^d) = \{x \in \mathbb{R}^d : ||x - a|| = r\} \ .$$

$S_r(a, \mathbb{R}^d)$ is the *sphere* of radius r centered at a. Formula ($\dagger$) shows that two such spheres in $\mathbb{R}^d$ are homeomorphic. The "dimension" of $S_r(a, \mathbb{R}^d)$ is $d - 1$, which should be intuitively reasonable from geometric examples when $d = 2$ or 3. Notice that

$$S_r(a, \mathbb{R}^d) = B_r(a, \mathbb{R}^d) - B_r^o(a, \mathbb{R}^d) \ ,$$

and it is natural to think of $S_r(a, \mathbb{R}^d)$ as the *boundary* sphere for the closed disk $B_r(a, \mathbb{R}^d)$. (For the reader who knows some point-set topology, the sphere with center a and radius r is the *point-set boundary* in $\mathbb{R}^d$ of both the open and the closed disks in $\mathbb{R}^d$ with the same center and same radius. For the reader acquainted with manifolds, the sphere is the *boundary manifold* of the closed disk.) □