

TOPOLOGY OF THE VIEW COMPLEX

DMITRY N. KOZLOV

(communicated by Donald M. Davis)

Abstract

In this paper we consider a family of simplicial complexes, which we call the *view complexes*. Our choice of objects of study is motivated by theoretical distributed computing, since the view complex is a key simplicial construction used for protocol complexes in the snapshot computational model. We show that the view complex View^n can be collapsed to the well-known complex $\chi(\Delta^n)$, called standard chromatic subdivision of a simplex, and that $\chi(\Delta^n)$ is itself collapsible. Furthermore, we show that the collapses can be performed simultaneously in entire orbits of the natural symmetric group action. Our results yield a purely combinatorial and constructive understanding of the topology of view complexes, at the same time as they enhance our knowledge about the standard chromatic subdivision of a simplex.

1. Introduction

Although a close connection between theoretical distributed computing and algebraic topology has by now been securely established (see, e.g., [AR02, HS99] [HKR14, Ko12, Ko14a, Ko14b]), many questions pertaining to arising simplicial structures have remained unanswered. In this introduction we make a short excursion to the distributed context, before proceeding with a purely topological study in the following sections.

In theoretical distributed computing one considers n processes which are trying to solve a task by executing a distributed protocol. To specify the task, one simply fixes the set of possible input configurations, and for each input one fixes the set of legal output configurations. A distributed protocol is a sequence of instructions where the type of instructions we are allowed to use depends on the choice of the communication model.

Since one considers the asynchronous model, there is no such thing as the unique execution of the distributed protocol; instead, one has a number of possible executions. One says that the distributed protocol solves the given task if it yields a correct output for any input configuration and any possible execution.

It has been realized, at least since [HS99], that it is fruitful to summarize the totality of all possible executions as a single simplicial complex. This complex is called the

Received June 26, 2014, revised December 15, 2014; published on May 11, 2015.

2010 Mathematics Subject Classification: 57Q05, 68Q85.

Key words and phrases: collapses, distributed computing, combinatorial algebraic topology, immediate snapshot, read-write protocols.

Article available at <http://dx.doi.org/10.4310/HHA.2015.v17.n1.a15>

Copyright © 2015, International Press. Permission to copy for private use granted.

protocol complex and depends on the distributed protocol and on the chosen communication model. Specifically, the protocol complex is a pure simplicial complex whose dimension is one less than the number of processes. The top-dimensional simplices of the protocol complex correspond to all possible executions of a given protocol, and its vertices correspond to all possible views of processes at the end of an execution.

Even though we focus on the shared-memory communication models, there are still many options to choose from some are more natural than others. Curiously, the protocol complexes corresponding to the most natural model, the one using write and snapshot read operations with no restrictions on their interweaving during executions, have not been studied much, since their simplicial structure is rather complicated. Still, it was shown that these complexes are always contractible; see [HKR14, Chapter 10], [Ha04].

In contrast, the topology of the protocol complexes of the layered immediate snapshot wait-free protocols is very easy. Namely, one can show that each such protocol complex is a subdivision of a simplex whose vertices are indexed by the processes; see [Ko12, HKR14]. A crucial construction in understanding the topology of these protocol complexes for $n + 1$ processes is the so-called *standard chromatic subdivision* $\chi(\Delta^n)$; see [HKR14, Subsections 3.6.3, 8.4.1, Chapter 16], [AR02, BG93, HS99, Ko12, SZ00].

In this paper, we study the analog of this construction, which is derived from the snapshot model. We call the corresponding simplicial complex the *view complex*; see Definition 3.1 for a completely combinatorial description. From the point of view of distributed computing, the view complex is a very central object, since it is the protocol complex for the snapshot protocol in which each process executes exactly one round. However, for us this is just a motivation, and we study the family of simplicial complexes $(\text{View}^n)_{n=1}^\infty$ from a purely topological point of view.

As the lower-dimensional examples already show, starting with $n = 2$ the simplicial complex View^n does not have to be a subdivision of a simplex. As a matter of fact, it is not a manifold—not even a pseudomanifold, since its simplices of codimension 1 may belong to more than two top-dimensional simplices. Yet we show that it is possible to understand the topology of the complex View^n rather completely.

To start with, it is easy to see directly that the simplicial complex View^n contains $\chi(\Delta^n)$ as a subcomplex. This makes sense in the distributed computing context since every immediate snapshot execution is also an execution in the snapshot model. Our main theorem, Theorem 4.4, states that View^n can be collapsed to $\chi(\Delta^n)$, and that $\chi(\Delta^n)$ is itself collapsible. This yields a constructive and purely combinatorial proof of contractibility of View^n . However, it is stronger than the mere contractibility, being rather a statement about the involved simplicial structures. We remark also that the fact that the standard chromatic subdivision $\chi(\Delta^n)$ is collapsible is new.

The simplicial complexes View^n and $\chi(\Delta^n)$ are equipped with a canonical simplicial action of the permutation group $\mathcal{S}_{[n]}$. This is the reflection of the fact that the considered protocols are symmetric with respect to the renaming of processors. The statement which we actually prove in Theorem 4.4 is stronger than just collapsibility. We show that our collapses can be done in an $\mathcal{S}_{[n]}$ -equivariant way, meaning that entire $\mathcal{S}_{[n]}$ -orbits of collapses can be performed simultaneously; see Section 4, and specifically Definition 4.1, for the precise meaning of $\mathcal{S}_{[n]}$ -collapses.

We recall that in theoretical distributed computing it is well-known that the snapshot and the immediate snapshot models are computationally equivalent; see, e.g., [HKR14, Chapter 14].

2. Notation

In order to make this paper self-contained, we shall now fix notation and define several standard notions. To start with, for an arbitrary positive integer n , we let $[n]$ denote the set $\{0, \dots, n\}$. Furthermore, we shall use the symbols \subset and \supset to denote the strict set inclusion.

Definition 2.1. Let S be a finite set. A family of subsets $K \subseteq 2^S$ is called an **abstract simplicial complex** if

- for all $x \in S$, we have $\{x\} \in K$;
- if $A \subset B$, and $B \in K$, then $A \in K$.

In the context of Definition 2.1, the set S is called the *vertex set* of K . Each $\sigma \in K$ is called a *simplex* of K . The number $|\sigma| - 1$ is called the *dimension* of σ and is denoted by $\dim \sigma$. For brevity, and following the standard practice, when $K = 2^S$ we shall simply call the corresponding abstract simplicial complex a *simplex*.

There are two simplicial complexes whose vertex set is an empty set, i.e., $S = \emptyset$: namely $K = \emptyset$, which we call the *void complex*, and $K = \{\emptyset\}$, which we call the *empty complex*. The two complexes may appear similar, but this impression is misleading, as they have different topological properties.

Given a simplicial complex, its simplices can be ordered by inclusion; the obtained partially ordered set $\mathcal{F}(K)$ is called the *face poset* of K . For the void complex, the face poset is empty; in all other cases the face poset has a single minimal element, which corresponds to the empty set. The face poset of a simplicial complex with at least one vertex has a single maximal element if and only if this complex is a simplex. In this case, the face poset is also called a *boolean lattice*, and is denoted by \mathcal{B}_n , where n is the number of vertices of the simplex. Another example is shown in Figure 2.1; for further details on the face poset of a simplicial complex we refer to [Ko07, Chapter 2].




Figure 2.1: A collapsible simplicial complex, and its face poset, with free simplices marked with solid dots.

Assume now that K is a simplicial complex, and σ is one of its simplices. The *link* of σ in K is a simplicial subcomplex of K defined as follows:

$$\text{lk}_K(\sigma) = \{\tau \in K \mid \tau \cap \sigma = \emptyset, \tau \cup \sigma \in K\}.$$

The next definition gives a purely combinatorial description of the notions of collapse and collapsible complex.

Definition 2.2. Assume K is an abstract simplicial complex and $\sigma \in K$. Set $\mathcal{F}(K)_{\geq \sigma} := \{\tau \supseteq \sigma \mid \tau \in K\}$ viewed as a subposet of the face poset of $\mathcal{F}(K)$. We call σ **free**¹ if $\mathcal{F}(K)_{\geq \sigma} \simeq \mathcal{B}_t$, for some $t \geq 1$, where \simeq denotes the poset isomorphism. Alternatively, a simplex σ is free if its link is a nonempty simplex.

For a free simplex σ , a **collapse** of K associated to σ is the process of deleting from K all the simplices in $\mathcal{F}(K)_{\geq \sigma}$. We denote the obtained complex by $K \downarrow \sigma$. A collapse is called **elementary** if $\mathcal{F}(K)_{\geq \sigma} \simeq \mathcal{B}_1$.

When M is a subcomplex of K , we say that K is **collapsible to M** if there exists a sequence of collapses leading from K to M . We say that K is **collapsible** if it is collapsible to the void simplicial complex.

Figure 2.1 illustrates the notions of free simplices and collapsibility. Note that the void simplicial complex is collapsible, as is any simplex, while the empty simplicial complex is not collapsible. For a collapsible simplicial complex the total number of simplices in odd dimensions (including the empty simplex) must be equal to the total number of simplices in even dimensions.

When K is collapsible to M , we shall use the notation $K \searrow M$. Note that in this case, there could be many different collapsing sequences, as can be seen in the example when K is a simplex and M is void.

Topologically, each collapse can be viewed as a strong deformation retraction. In particular, if K is collapsible to M , then one can easily construct an explicit strong deformation retraction from (the geometric realization of) K to M . This means, of course, that K and M have the same homotopy type, that collapsible simplicial complexes are also contractible, and that contraction can be explicitly described. We refer the reader who wishes to gain deeper insight into the notion of collapse to consult [Co73] or [Ko07, Chapter 6].

3. Combinatorial description of the view complex

We now proceed to defining the main objects of study of this paper. As mentioned in the introduction, our definition is strongly dictated by the context of theoretical distributed computing. However, we choose to give an abstract description using only combinatorial topology.

Definition 3.1. Assume we are given a natural number n . An n -**view** is a $2 \times t$ -matrix of subsets of $[n]$

$$\begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & I_t \end{pmatrix}, \quad (3.1)$$

where $t \geq 1$, such that the following properties are satisfied:

¹We note that there are various notions of free simplex used in the literature; see, e.g., [St66].

- (1) $\emptyset \neq V_1 \subset \dots \subset V_{t-1} \subset [n]$;
- (2) the sets I_1, \dots, I_t are disjoint;
- (3) $\emptyset \neq I_k \subseteq V_k$, for $k = 1, \dots, t - 1$.

When n is fixed or clear from the context, we shall simply call such a $2 \times t$ -matrix a *view*. We shall also use the convention $V_t = [n]$.

Definition 3.2. Assume we have a natural number n , and an n -view

$$W = \begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & I_t \end{pmatrix}.$$

We set $\dim W := |I_1| + \dots + |I_t| - 1$, and call it the **dimension** of the view W .

Clearly, there is exactly one n -view of dimension -1 , namely

$$W = \begin{pmatrix} [n] \\ \emptyset \end{pmatrix}.$$

The n -views of dimension 0 are of the form

$$\begin{pmatrix} \{x\} \cup A & [n] \\ \{x\} & \emptyset \end{pmatrix} \text{ and } \begin{pmatrix} [n] \\ \{x\} \end{pmatrix}, \text{ where } \{x\} \cup A \subset [n].$$

In any case we see that the n -views of dimension 0 are indexed by pairs (V, x) , where $V \subseteq [n]$ and $x \in V$. We call such a pair a *local view*, or sometimes more specifically a *local view of x* .

Definition 3.3. Assume we are given an n -view

$$W = \begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & I_t \end{pmatrix},$$

and a local view $L = (V, x)$. We say that L **belongs to** W , writing $L \in W$, if there exists $1 \leq k \leq t$, such that $V = V_k$ and $x \in I_k$.

For an arbitrary view W , we let $V(W)$ denote the set of all local views belonging to W . Clearly, $|V(W)| = \dim W + 1$.

Definition 3.4. For an arbitrary natural number n , we define an abstract simplicial complex View^n as follows:

- the set of vertices is the set of all local views

$$V(\text{View}^n) := \{(V, x) \mid x \in V \subseteq [n]\};$$

- a subset $S \subseteq V(\text{View}^n)$ forms a simplex if and only if $S = V(W)$ for some n -view W .

We shall identify n -views with simplices of View^n .

Proposition 3.5. *The simplicial complex View^n is well-defined.*

Proof. Given a simplex W of dimension d , one obtains all of its boundary simplices of dimension $d - 1$ by deleting an element from one of the sets I_1, \dots, I_t . If after this the set becomes empty, one deletes the corresponding column in the $2 \times t$ -matrix, unless

it is the last column. Clearly, what we get is again a view whose set of local views is obtained from $V(W)$ by deleting one of the elements. Iterating this argument, we see that the conditions of Definition 2.1 are satisfied, and the simplicial complex View^n is well-defined. \square

We shall say that a view W contains a view U , and write $U \subseteq W$, if the simplex indexed by W contains the simplex indexed by U .

Some facts about the simplicial complex View^n are immediate. It is a *pure* simplicial complex of dimension n , meaning that all of its maximal simplices have dimension n . It is easily seen to have $(n + 1) \cdot 2^n$ vertices. With a little more effort one can see that View^n has $(n + 1) \cdot n \cdot (2 \cdot 3^{n-1} - 2^{n-2})$ edges. The examples of View^n , for $n = 1$ and $n = 2$, are shown on Figure 3.1.




Figure 3.1: The complexes View^n for $n = 1$ and $n = 2$.

For an arbitrary set A , let \mathcal{S}_A denote the permutation group of the set A ; in particular, let $\mathcal{S}_{[n]}$ denote the permutation group of the set $[n]$. Clearly, this group is isomorphic to the symmetric group \mathcal{S}_{n+1} . Furthermore, there is a natural simplicial group action of $\mathcal{S}_{[n]}$ on View^n induced by the permutation action on the ground set $[n]$.

We now define a distinguished subcomplex of View^n .

Definition 3.6. Assume n is a natural number.

- (1) We call an n -view

$$W = \begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & I_t \end{pmatrix}$$

an **immediate snapshot view** if we have $I_k \subseteq V_k \setminus V_{k-1}$ for all $k = 2, \dots, t$.

- (2) If W is an immediate snapshot view and $U \subset W$, then U is also an immediate snapshot view. Therefore the immediate snapshot views form a simplicial subcomplex of View^n , which we denote by $\chi(\Delta^n)$.

Since the condition of being an immediate snapshot view is formulated using the set operations only, we see that $\chi(\Delta^n)$ is also invariant under the $\mathcal{S}_{[n]}$ -action above.

It is furthermore clearly seen from the condition in Definition 3.6(1) that the difference between View^n and $\chi(\Delta^n)$ is first visible when $n = 2$. When $n = 1$, these two complexes are the same. When $n = 2$ the complex View^n is obtained from $\chi(\Delta^n)$ by adding 6 triangles of the form

$$\begin{pmatrix} \{a, b\} & [2] \\ \{a\} & \{b, c\} \end{pmatrix}$$

where $\{a, b, c\} = [2]$; see Figure 3.1.

In general, the simplicial complex $\chi(\Delta^n)$ is known as the *standard chromatic subdivision* of an n -simplex; see [HKR14]. Its topology has been studied in [Ko12], where the author showed that $\chi(\Delta^n)$ is a simplicial subdivision of an n -simplex. The reader is invited to see how our description here is equivalent to the one given in [Ko12, Proposition 2.3].

It is easy to see that the $\mathcal{S}_{[n]}$ -action above is the coordinate permutation action on the standard n -simplex. In the context of distributed computing it is particularly important to know that $\chi(\Delta^n)$ is a pseudomanifold. This is well-known due to the work of Herlihy and others; see [HKR14, Chapter 9].

In contrast, even though the distributed computing interpretation of the complex View^n is simpler than that of $\chi(\Delta^n)$, understanding its simplicial structure is harder. As a matter of fact, it is easy to use the nerve lemma, [Ko07], in the same way as in [HKR14, Chapter 10], to show that View^n is contractible. We do not detail this argument here, since a much stronger result will be shown in the next section. Namely, we show that View^n is equivariantly collapsible, and provide an explicit sequence of such equivariant collapses.

4. Collapsing procedure

As the main result of this paper, we shall see that View^n can be collapsed to $\chi(\Delta^n)$. As a matter of fact, the collapses can be done in an $\mathcal{S}_{[n]}$ -equivariant way. The next definition formalizes this concept.

Definition 4.1. Assume K is an abstract simplicial complex with a simplicial action of finite group G .

- (a) A simplex σ is called *G -free* if it is free, and for all $g \in G$, such that $g(\sigma) \neq \sigma$, we have

$$\mathcal{F}(K)_{\geq \sigma} \cap \mathcal{F}(K)_{\geq g(\sigma)} = \emptyset. \tag{4.1}$$

- (b) If σ is G -free, we call the procedure of deleting all the simplices from the union $\bigcup_{g \in G} \mathcal{F}(K)_{\geq g(\sigma)}$ the *G -collapse* of K .

Note that when σ is free, each $g(\sigma)$ is automatically free as well. Therefore, deleting the simplices from $\mathcal{F}(K)_{\geq g(\sigma)}$ is also a collapse. The condition (4.1) guarantees that all these collapses can be done simultaneously and independently of each other. This

is because for all $g, h \in G$, whenever τ is a simplex of K , the identity

$$\mathcal{F}(K)_{\geq h_1(\tau)} \cap \mathcal{F}(K)_{\geq h_2(\tau)} = \emptyset$$

follows from (4.1) by substituting $\sigma := h_1(\tau)$ and $g := h_2 \circ h_1^{-1}$. In particular, we see that geometrically a G -collapse yields a G -equivariant strong deformation retraction.

The simplest example of a free simplex which is not G -free is given by taking K to be a 1-simplex, and letting $G = \mathbb{Z}_2$ act on K by swapping the vertices; see Figure 4.1. Each vertex is free, and leads to an (elementary) collapse, but they are not \mathbb{Z}_2 -free, and the collapses cannot be performed simultaneously. As shown in the same figure, subdividing the interval in the middle leads to a \mathbb{Z}_2 -complex in which both end vertices are \mathbb{Z}_2 -free. A more complicated example of an \mathcal{S}_3 -collapsing sequence is shown in Figure 4.3.



Figure 4.1: Two simplicial complexes with \mathbb{Z}_2 -action. The end points are \mathbb{Z}_2 -free in the second, but not in the first complex.

In analogy with the usual collapses we shall say that K is G -collapsible to a subcomplex M if there exists a sequence of G -collapses leading from K to M . In this case M must be G -invariant, and we use the notation $K \searrow_G M$. In the same way, we say that K is G -collapsible if it is G -collapsible to a void complex.

An example of an $\mathcal{S}_{[n]}$ -collapsible complex is the n -simplex Δ^n . However, in contrast to the situation with regular collapses, there is a unique $\mathcal{S}_{[n]}$ -collapsing sequence from Δ^n to the void complex. Namely, we must collapse Δ^n to the void complex in one single step.

We now proceed to define two functions on the set of views, which will be crucial for constructing our collapsing sequence.

Definition 4.2. For an arbitrary n -view

$$W = \begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & I_t \end{pmatrix}, \tag{4.2}$$

we set


$$\Phi(W) := \begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & I_t \cap V_{t-1} \end{pmatrix},$$

and

$$\Psi(W) := \begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & I_t \cup ([n] \setminus V_{t-1}) \end{pmatrix}.$$

Furthermore, we let $I(W)$ denote the closed interval $[\Phi(W), \Psi(W)]$ in the face poset of View^n .

In particular, we allow $\Phi(W)$ to be the empty simplex. This is the case when $I_1 = \dots = I_{t-1} = I_t \cap V_{t-1} = \emptyset$, which is equivalent to saying that $t = 1$, in other words $W = \begin{pmatrix} [n] \\ I \end{pmatrix}$, for some $I \subseteq [n]$. Definition 4.2 is illustrated on Figure 4.2.


 Figure 4.2: The functions Φ and Ψ .

Proposition 4.3. *The maps Φ and Ψ enjoy the following properties:*

- (1) $\Phi(W) \subseteq W \subseteq \Psi(W)$, for all views W ;
- (2) $\Phi(W) \subseteq U \subseteq \Psi(W)$ implies $\Phi(U) = \Phi(W)$ and $\Psi(U) = \Psi(W)$;
- (3) $\Phi(W) \neq \Psi(W)$, for all views W ;
- (4) if $U \subseteq W$, then $\Phi(U) \subseteq \Phi(W)$;
- (5) for all views W we have the following implication: if $\Phi(W) \in \mathcal{F}(\chi(\Delta^n))$ then $W \in \mathcal{F}(\chi(\Delta^n))$;
- (6) for all $W \in \text{View}^n$, and all $\pi \in \mathcal{S}_{[n]}$, we have $\pi(I(W)) = I(\pi(W))$.

Proof. To see (1), note that the view $\Phi(W)$ is obtained from W by removing the local views of the form $([n], x)$, such that $x \in I_t \setminus V_{t-1}$, hence $\Phi(W) \subseteq W$. The view $\Psi(W)$ is obtained from W by adding local views of the form $([n], x)$, such that $x \notin I_t \cup V_{t-1}$, hence $\Psi(W) \supseteq W$.

Let us now show (2). Assume $\Phi(W) \subseteq U \subseteq \Psi(W)$, where W is a view given in (4.2). The difference between $\Phi(W)$ and $\Psi(W)$ consists of all local views $([n], x)$, such that $x \notin V_{t-1}$; see Figure 4.2. This means that there exists a set S satisfying $S \cap V_{t-1} = \emptyset$, such that

$$U = \begin{pmatrix} V_1 & \dots & V_{t-1} & [n] \\ I_1 & \dots & I_{t-1} & (I_t \cap V_{t-1}) \cup S \end{pmatrix}.$$

Clearly,

$$V_{t-1} \cap ((I_t \cap V_{t-1}) \cup S) = (I_t \cap V_{t-1}) \cup (V_{t-1} \cap S) = I_t \cap V_{t-1},$$

so $\Phi(U) = \Phi(W)$. Furthermore,

$$(I_t \cap V_{t-1}) \cup S \cup ([n] \setminus V_{t-1}) = (I_t \cap V_{t-1}) \cup ([n] \setminus V_{t-1}),$$

so $\Psi(U) = \Psi(W)$.

To see (3), we note that we always have $V_{t-1} \neq [n]$; hence there exists at least one local view $([n], x) \in V(W)$ such that $x \notin V_{t-1}$.

To show (4), assume $U \subseteq W$. If U is the empty simplex, then the statement is obvious, so let us assume U is not the empty simplex. Consider the presentations

$$W = \begin{pmatrix} V_1 & \cdots & V_{t-1} & [n] \\ I_1 & \cdots & I_{t-1} & I_t \end{pmatrix}, \quad U = \begin{pmatrix} V_1 & \cdots & V_{t-1} & [n] \\ J_1 & \cdots & J_{t-1} & J_t \end{pmatrix},$$

where $J_k \subseteq I_k$, for all $k = 1, \dots, t$. For the view W the presentation above is standard, but not necessarily for U . Namely, we allow some of J_k 's to be empty. If they are, we can simply delete the corresponding column (unless it is the last column) to obtain the standard presentation on an n -view that we have used so far.

Let $1 \leq k \leq t-1$ be the largest index less than t for which $J_k \neq \emptyset$, and set $S := V_k$. If $J_1 = \dots = J_{t-1} = \emptyset$, we set $S := \emptyset$. We have

$$\Phi(U) = \begin{pmatrix} V_1 & \cdots & V_{t-1} & [n] \\ J_1 & \cdots & J_{t-1} & J_t \cap S \end{pmatrix},$$

where again we use the presentation where empty sets in the second row are allowed. In any case, we have $J_t \subseteq I_t$ and $S \subseteq V_{t-1}$, hence $J_t \cap S \subseteq I_t \cap V_{t-1}$. Together with the fact that $J_k \subseteq I_k$, for $k = 1, \dots, t-1$ this implies $\Phi(U) \subseteq \Phi(W)$.

To see (5), note that, since W differs from $\Phi(W)$ only in I_t , we only need to check the condition from Definition 3.6(1) for this set. As noticed earlier, when going from $\Phi(W)$ to W , the set I_t is enlarged by some elements $x \notin V_{t-1}$. Hence the condition $I_t \subseteq [n] \setminus V_{t-1}$ will be satisfied for W as well, and we may conclude that W is an immediate snapshot view, if $\Phi(W)$ is one.

To see (6), note that the definitions of $\Phi(W)$ and $\Psi(W)$, and hence also the definition of the interval $I(W)$, are formulated completely in terms of set operations; see also Figure 4.2. The $\mathcal{S}_{[n]}$ -action is simply renaming of the elements of the underlying set $[n]$, so $\pi(I(W)) = I(\pi(W))$. \square

Note that Proposition 4.3(6) implies that in particular $\Phi(\pi(W)) = \pi(\Phi(W))$ and $\Psi(\pi(W)) = \pi(\Psi(W))$. Furthermore, since the intervals $\pi(I(W))$ and $I(W)$ are either equal or disjoint, we see that if $\Phi(W) \neq \pi(\Phi(W))$, then $I(W)$ is disjoint from $I(\pi(W))$.

We are now ready to prove our main theorem.

Theorem 4.4. *For every natural number n , the following statements are true.*

- (1) *The simplicial complex View^n is $\mathcal{S}_{[n]}$ -collapsible to $\chi(\Delta^n)$.*
- (2) *The simplicial complex $\chi(\Delta^n)$ is $\mathcal{S}_{[n]}$ -collapsible.*

Proof. Proposition 4.3(1) implies that $\Phi(W) \subseteq \Psi(W)$, for all views W . Note that as a poset $I(W) \simeq \mathcal{B}_{|\Psi(W)|-|\Phi(W)|}$. Since $W \in I(W)$, for all views W , the intervals $I(W)$ cover the poset $\mathcal{F}(\text{View}^n)$.

On the other hand, Proposition 4.3(2) shows that either $I(W) = I(U)$, or $I(W) \cap I(U) = \emptyset$, for all views U, W . So $\mathcal{F}(\text{View}^n)$ can, in fact, be decomposed as a disjoint union of intervals $I(W_1), \dots, I(W_c)$, for some choice of views W_1, \dots, W_c . Without

loss of generality we can assume that $\Phi(W_1) = W_1, \dots, \Phi(W_c) = W_c$, and $|W_1| \geq |W_2| \geq \dots \geq |W_c|$. By Proposition 4.3(2) we see that for all $U \in \mathcal{F}(\text{View}^n)$ there exists k between 1 and c such that $\Phi(U) = W_k$. We now want to show that starting with View^n , collapsing first W_1 , then W_2 , and so on, until W_c , will yield a collapsing sequence from View^n to the void complex.

For every $0 \leq k \leq c$ let V_k be the subcomplex of View^n consisting of all simplices τ , such that $\tau \not\supseteq W_i$, for all $i = 1, \dots, k$. In particular, $V_0 = \text{View}^n$. We shall show by induction on k , $1 \leq k \leq c$, that

$$\mathcal{F}(V_{k-1})_{\geq W_k} = I(W_k). \tag{4.3}$$

We start with $k = 1$. If $U \supseteq W_1$, then Proposition 4.3(4) implies that $\Phi(U) \supseteq \Phi(W_1) = W_1$. But we know that $\Phi(U) = W_l$, for some l ; hence we get $l = 1$, and $\Phi(U) = W_1$. This means that $U \in I(W_1)$. Altogether this implies that $\mathcal{F}(\text{View}^n)_{\geq W_1} = I(W_1)$.

For an induction step, let $2 \leq k \leq c$. Assume $U \supseteq W_k$, then as before $\Phi(U) \supseteq \Phi(W_k) = W_k$. Pick l such that $\Phi(U) = W_l$. If $k = l$, then $\Phi(U) = W_k$; hence $U \in I(W_k)$. Otherwise, we must have $W_l \supset W_k$, in particular $|W_l| > |W_k|$. This implies $k > l$, and so $U \notin V_{k-1}$. This shows that $\mathcal{F}(V_{k-1})_{\geq W_k} \subseteq I(W_k)$. The other direction follows from the fact that the intervals are disjoint, together with the induction hypothesis. Namely, we have $\mathcal{F}(V_{k-1}) = \mathcal{F}(\text{View}^n) \setminus \cup_{i=1}^{k-1} I(W_i)$, which implies $I(W_k) \subseteq \mathcal{F}(V_{k-1})$, and hence $I(W_k) \subseteq \mathcal{F}(V_{k-1})_{\geq W_k}$. Summarizing, we conclude that (4.3) holds for this k .

The equality (4.3) means that for every $1 \leq k \leq c$, the simplex W_k is free in V_{k-1} , and that the corresponding collapse results in V_k . Proposition 4.3(3) says that $\Phi(W) \neq \Psi(W)$, in particular $I(W) \simeq \mathcal{B}_t$, for $t \geq 1$, so the condition of Definition 2.2 is satisfied.

Since we already saw that $\mathcal{F}(\text{View}^n)$ is a disjoint union of the intervals $I(W_1), \dots, I(W_c)$, we conclude that View^n is collapsible.

We shall now adjust the collapsing sequence above to first lead to $\chi(\Delta^n)$, and then collapse $\chi(\Delta^n)$. First, note that the condition $|W_1| \geq \dots \geq |W_c|$ was strictly speaking unnecessarily strong to be able to conclude that we have a collapsing sequence. All we needed was the implication that if $W_l \supset W_k$, then $l < k$, so any linear extension of the set $\{W_1, \dots, W_c\}$, equipped with the reverse inclusion order, would do. On the other hand, note that Proposition 4.3(5) implies that for all W either $I(W) \subseteq \mathcal{F}(\chi(\Delta^n))$, or $I(W)$ and $\mathcal{F}(\chi(\Delta^n))$ are disjoint.

Since $\chi(\Delta^n)$ is a simplicial subcomplex of View^n , its face poset is a lower ideal of $\mathcal{F}(\text{View}^n)$. In particular, the linear extension of the set $\{W_1, \dots, W_c\}$ can be chosen in a special way: first take any linear extension of the subset $\{W_1, \dots, W_c\} \cap (\mathcal{F}(\text{View}^n) \setminus \mathcal{F}(\chi(\Delta^n)))$, and then concatenate it with any linear extension of $\{W_1, \dots, W_c\} \cap \mathcal{F}(\chi(\Delta^n))$. By what was said above, this concatenation is a linear extension by itself. This linear extension now yields a collapsing sequence from View^n to $\chi(\Delta^n)$, and then from $\chi(\Delta^n)$ to the void complex.

To finish the proof of our theorem, we need to modify the collapsing sequence once more, in order to incorporate the $\mathcal{S}_{[n]}$ -action. Until now we did everything for an arbitrary set $\{W_1, \dots, W_c\}$ such that the intervals $I(W_1), \dots, I(W_c)$ cover $\mathcal{F}(\text{View}^n)$. We shall now specify this set.

Set $L := \{\Phi(W) \mid W \in \text{View}^n, W \notin \chi(\Delta^n)\}$, and pick W_1 such that $|W_1| = \max_{W \in L} |W|$. By what is said above, W_1 is free in View^n . Let $\{W_1, \dots, W_p\}$ be any

set of representatives of the orbit $\mathcal{S}_{[n]}(W_1)$. Specifically, this means that the views W_1, \dots, W_p are all distinct, and for each $\pi \in \mathcal{S}_{[n]}$ there exists i , $1 \leq i \leq p$, such that $\pi(W_1) = W_i$. By Proposition 4.3(6), we know that the intervals $I(W_1), \dots, I(W_p)$ are disjoint, and for each $\pi \in \mathcal{S}_{[n]}$ there exists i , $1 \leq i \leq p$, such that $\pi(I(W_1)) = I(W_i)$. In other words, $I(W_1), \dots, I(W_p)$ is a set of representatives of the orbit $\mathcal{S}_{[n]}(I(W_1))$.

We can now repeat the entire procedure with $\mathcal{F}(\text{View}^n) \setminus \cup_{i=1}^p I(W_i)$, and then proceed iterating until the entire difference $\mathcal{F}(\text{View}^n) \setminus \mathcal{F}(\chi(\Delta^n))$ is covered with chosen intervals. After that we proceed to do the same for $\chi(\Delta^n)$.

As a final result, we will obtain a sequence of collapses such that

- the collapses come in $\mathcal{S}_{[n]}$ -equivariant batches; in each such batch, all collapses can be performed simultaneously, resulting in an $\mathcal{S}_{[n]}$ -collapse;
- the collapses first exhaust the difference $\mathcal{F}(\text{View}^n) \setminus \mathcal{F}(\chi(\Delta^n))$, and then proceed to collapse the simplicial complex $\chi(\Delta^n)$.

The arguments above adapt easily to this specific sequence of collapses, so we are finished with the proof of the theorem. \square

For $n = 2$ we illustrate the collapsing procedure from Theorem 4.4 in Figure 4.3.




Figure 4.3: The collapsing procedure for View^2 . The first step collapses the 6 extraordinary triangles.

References

- [AR02] H. Attiya and S. Rajsbaum, The combinatorial structure of wait-free solvable tasks, *SIAM J. Comput.* **31**(4) (2002), 1286–1313.
- [AW04] H. Attiya and J. Welch, Distributed computing: fundamentals, simulations, and advanced topics, Wiley Series on parallel and distributed computing, 2nd ed, Wiley-Interscience, 2004, 432 pp.
- [BG93] E. Borowsky and E. Gafni, Generalized FLP impossibility result for t -resilient asynchronous computations, STOC 1993, 91–100.

- [Co73] M. Cohen, A course in simple-homotopy theory, Graduate Texts in Mathematics, Vol. 10, Springer-Verlag, New York-Berlin, 1973.
- [Ha04] J. Havlicek, A Note on the homotopy type of wait-free atomic snapshot protocol complexes, *SIAM J. Computing* **33** Issue 5, (2004), 1215–1222.
- [HKR14] M. Herlihy, D.N. Kozlov, and S. Rajsbaum, Distributed computing through combinatorial topology, Elsevier, 2014, 336 pp.
- [HS99] M. Herlihy and N. Shavit, The topological structure of asynchronous computability, *J. ACM* **46** (1999), no. 6, 858–923.
- [Ko07] D.N. Kozlov, Combinatorial algebraic topology, algorithms and computation in mathematics **21**, Springer-Verlag Berlin Heidelberg, 2008, XX, 390 p. 115 illus.
- [Ko12] D.N. Kozlov, Chromatic subdivision of a simplicial complex, *Homology, Homotopy and Applications* **14**(2) (2012), 197–209.
- [Ko14a] D.N. Kozlov, Witness structures and immediate snapshot complexes, preprint, 26 pages, [arXiv:1404.4250](https://arxiv.org/abs/1404.4250) [cs.DC].
- [Ko14b] D.N. Kozlov, Topology of immediate snapshot complexes, *Topology Appl.* **178**(1) (2014), 160–184.
- [SZ00] M.E. Saks and F. Zaharoglou, Wait-free k-set agreement is impossible: the topology of public knowledge, *SIAM J. Comput.* **29**(5) (2000), 1449–1483.
- [St66] R.E. Stong, Finite topological spaces, *Trans. AMS* **123** (1966), no. 2, 325–340.

Dmitry N. Kozlov dfk@math.uni-bremen.de

Department of Mathematics, University of Bremen, 28334 Bremen, Federal Republic of Germany