New Combinatorial Topology Upper and Lower Bounds for Renaming

[Extended Abstract] *

Armando Castañeda and Sergio Rajsbaum
Instituto de Matemáticas, Universidad Nacional Autónoma de México
acastanedar@uxmcc2.iimas.unam.mx, rajsbaum@matem.unam.mx

ABSTRACT

In the *renaming* task $n+1$ processes start with unique input names from a large space and must choose unique output names taken from a smaller name space, namely $0, 1, \ldots, K$. To rule out trivial solutions, a protocol must be *anonymous*: the value chosen by a process can depend on its input name and on the execution, but not on the specific process id.

Attiya et al. showed in 1990 that renaming has a wait-free solution when $K \geq 2n$. Several proofs of a lower bound stating that no such protocol exists when $K < 2n$ have been published. In this paper we prove that, for certain values of n, this lower bound is incorrect, exhibiting a wait-free renaming protocol for $K = 2n - 1$. For the other values of n, we present the first completely combinatorial lower bound proof stating that no such protocol exists when $K < 2n$.

More precisely, our main theorem states that there exists a wait-free renaming protocol for $K < 2n$ if and only if the set of integers $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$ are relatively prime. Thus, such protocol exists for six processes, and not for less. The proof of the theorem uses combinatorial topology techniques, both for the lower bound and to derive the renaming protocol.

Categories and Subject Descriptors: F.0 [Theory of computation]: General

General Terms: Algorithms, theory.

Keywords: Combinatorial topology, distributed systems, renaming, shared memory systems, wait-free computation.

1. INTRODUCTION

The 2004 Gödel Prize for outstanding journal articles in theoretical computer science was shared between Herlihy and Shavit [15], and Saks and Zaharoglou [18]. These papers, together with Borowsky and Gafni [6], discovered the topological nature of distributed computing and provided a new perspective on the area. The papers showed that the runs of any wait-free protocol in a read/write shared memory model can be represented by a simplicial complex, that "has no holes" (as in Figure 1) and used this powerful topological invariant together with Sperner's lemma to prove that the *k-set agreement* task (deciding on at most k different input values) cannot be wait-free solved in an asynchronous system of $n+1$ processes, if $k \leq n$, where *wait-free* means that any process must produce an output value in a fixed number of steps, regardless of delays or failures by other processes. Furthermore, [15] presented the Asynchronous Computability Theorem that characterizes the tasks that can be wait-free solved in such a system.

The paper [15] included one more application of the topology perspective: a renaming lower bound. In the *renaming* task $n+1$ processes start with unique input names from a large space and must choose unique output names taken from a smaller name space, namely $0, 1, \ldots, K$. To rule out trivial solutions, a protocol must be *anonymous*: the value chosen by a process can depend on its input name and on the execution, but not on the specific process id. Attiya et al. [3] presented a renaming wait-free protocol[1] for $K \geq 2n$, and showed that there is no wait-free solution if $K \leq n+1$. It was not until 1993 that the gap was closed, in [14], the conference version of [15], with a lower bound stating that no wait-free renaming protocol exists when $K < 2n$. In this paper we show that this lower bound does not hold for an infinite number of values of n; for the other values of n it holds, and we present a new, combinatorial lower bound proof.

The proof of [14] was the first of four lower bound proofs [1, 13, 14, 15], all closely related and based on algebraic topology, stating that no wait-free renaming protocol exists when $K < 2n$. The second proof appeared in [13] where a chain map [17] methodology is developed that was used to obtain lower bounds for other tasks. The lower bound proof published in [15] is the third one, but is based on the proof in [13]. The last proof we know of, appeared in [1]. A goal there was to provide a combinatorial version of the lower bound for renaming, and also for set agreement, that could be accessible to a reader unfamiliar with algebraic topology. They succeeded in doing this only for set agreement; for the renaming lower bound, in the crucial step of the proof[2],

*For a full version see [8].

[1]The protocol was presented in the message passing model, but can be extended to the shared read/write memory model [2].

[2]Actually, in the conference version of [1] this step was also proved combinatorially, but an error was discovered and in the journal version replaced by the lemma of [13].

PODC'08, August 18–21, 2008, Toronto, Ontario, Canada.

they relied on Lemma 6.1 of [13]. Thus, the question of a fully combinatorial proof for renaming was left open, and all four proofs are based on the same result, stated as Lemma 6.1[13].

This lemma was expressed combinatorially as Theorem 6.2[1]: roughly, that any binary coloring of a subdivided simplex that is symmetric on the boundary must produce a non-zero number of monochromatic simplexes, counted by orientation. In the example in Figure 1, $n = 2$, there is a subdivided 2-simplex (2 dimensional triangle), that has three 1-monochromatic 2-simplexes (the color 0 is represented by a white circle and the color 1 by a black circle), but counted by orientation it has only $+1$ or -1 (to orient it assign to each 2-simplex $+1$ or -1 in a way that if two 2-simplexes share an edge they have opposite sign; the two monochromatic simplexes inside the bold region must have opposite sign). Furthermore, the reader can verify that any coloring of the interior vertexes with 0's and 1's will produce monochromatic simplexes. Also, this will be true for *any* chromatic and symmetric subdivision of the 2-simplex.

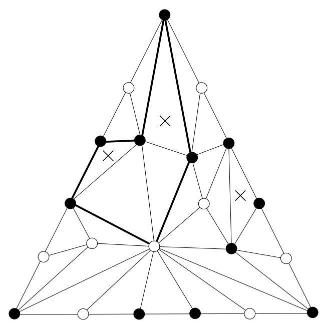

Figure 1: A subdivided simplex.

Our first result in this paper is to show that, for certain *exceptional* values of n, this is not true. We present subdivided simplexes K^n in Section 4.1, that are symmetric on the boundary and have zero monochromatic simplexes counted by orientation. The smallest such example exists for $n = 5$. That is, we discovered that while Theorem 6.2[1] is true in dimension 1 (a line), 2 (as in Figure 1), dimension 3 (subdivided tetrahedron), and dimension 4, it fails in dimension 5. These counterexamples K^n for exceptional n, imply that previous renaming lower bound proofs for $K < 2n$ are flawed.

Our second result, in Section 3, is a renaming lower bound proof for $K < 2n$, that holds for the other, non-exceptional values of n. For this, we prove a theorem that fully characterizes the number of monochromatic simplexes in a symmetric subdivided simplex. This is the first, fully combinatorial renaming lower bound proof, closing the open question left in [1].

Our third result is that the lower bound statement itself is incorrect, for the exceptional values of n. That is, we derive, for such values of n, a wait-free renaming protocol for $K = 2n - 1$. Technically, our contribution in Section 4.2 is an algorithm to eliminate all monochromatic simplexes from a counterexample K^n, without modifying its boundary. For example, in Figure 15 it is shown how to eliminate the two monochromatic simplexes in the bold region of the previous figure. Then, as explained in Section 5, the *Anonymous Computability Theorem* [15], or the Simplex Convergence algorithmic version [5], implies that the renaming protocol exists, using the equivalence of renaming and *weak-symmetry breaking* [11].

More precisely, combining the lower bound with the protocol, our main theorem states that there exists a wait-free renaming protocol for $K < 2n$ if and only if n is exceptional, where *exceptional* means that the integers in the set $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$ are relatively prime. For example, such protocol exists for $n = 5, 9, 11, 13, 14$ (exceptional), and does not exist for the other values smaller than 14 (as they are not exceptional). This is interesting in light of the result of [11] that states that renaming is weaker than set agreement when n is even. Our result implies that this claim is true for exceptional values of n.

We use the combinatorial topology framework of [1], both for the lower bound and to derive the renaming protocol. The lower bound is based on Fan's formula [10], a generalization of both Sperner's and Tucker's lemma (a combinatorial version of the Borsuk-Ulam Theorem). A particular case of the formula is called Index Lemma in [12]. The monochromatic simplex elimination algorithm of the upper bound is reminiscent of the equivariant Hopf theorem[3] [16].

For lack of space details and proofs have been omitted from this extended abstract. See [8] for a full version.

2. BASIC CONCEPTS

In this section we review basic concepts on distributed computing and combinatorial topology. For more details see [4] and [12].

2.1 Distributed Computing

Let $ID^n = \{0, \ldots, n\}$, and denote the set $ID^n - I$ as ID^n_I, where $I \subset ID^n$. We consider a set of $n + 1$ asynchronous processes with ids ID^n communicating through a read/write shared memory, and executing a *wait-free* protocol: any process that continues to run will halt with an output value in a fixed number of steps, regardless of delays or failures (only crash failures are considered) by other processes. A protocol is *anonymous* if it does not depend on process ids. In the K-*renaming* task [3] processes start with unique input names from a large name space, and must choose unique output names from the space $0, \ldots, K$. It can be formulated in other equivalent ways. For this paper it is convenient to use the *Weak Symmetry-Breaking* (WSB) formulation. It was shown in [11] that the WSB task is equivalent to K-renaming, $K = 2n - 1$. In the WSB task processes have no input values and the output values are 0 or 1. It is required that in every execution in which all $n+1$ processes participate, at least one process decides 1 and at least one process decides 0. Like renaming, any protocol that implements WSB must be anonymous.

We represent tasks and protocol executions following the combinatorial topology notation of [1]. An initial or final state of a process is modeled as a *vertex* $v = (P, w)$, a pair consisting of a process id P and a value w which represents its state, possibly including its output value. We say the vertex is *colored* with the process *id*, but we will also use other

[3]This well-known topology theorem states that two maps $f_1, f_2 : X \to S$ from a compact manifold X to a sphere S of the same dimension are homotopic if and only if they have the same Brouwer degree.

colorings for vertexes, especially colors 0 and 1 (this binary coloring will be represented in all the pictures with white and black circles). A set of $n+1$ mutually compatible initial or final states is modeled as an *n-simplex* (a nonempty set with $n+1$ elements) $\sigma^n = (a_0, \ldots, a_n)$, of dimension $dim(\sigma^n) = n$. A nonempty subset of a simplex is called a *face*. Notice that the number of i-faces (the faces of dimension i) of an n-simplex is $\binom{n+1}{i+1}$. Also, the coloring with *ids* of a simplex is *proper*, as it gives different values to different vertexes of the same simplex. We denote the *id* color of a simplex σ by $id(\tau)$. If a coloring of a simplex gives the same value b to every vertex of the simplex then the coloring is *b-monochromatic*. A *complex* K is a set of simplexes, closed under containment[4]. The *star complex* of a vertex v in K^n, denoted by $st(v, K^n)$, is the complex consisting of those n-simplexes of K^n that contain the vertex v, together with all their faces. Also, for simplex σ^n, we denote by $M(\sigma^n)$ the complex consisting of σ^n together with all its faces. A function from the vertexes of a complex K to the vertexes of a complex L is *simplicial* if the image of every simplex of K is a simplex of L.

The *j-graph* of a complex has a vertex for every j-simplex and an edge between two vertexes if they share a $(j-1)$-face. The complex is *j-connected* if its j-graph is connected, or if it consists of a single vertex when $j = 0$. A *j-path* P is a path in the j-graph. The *size* of P, $|P|$, is the number of its j-simplexes. Unless otherwise specified, we always consider simple paths.

A *task* is a problem where each process starts with a private input value, communicates with the others, and halts with a private output value. It is given by a colored *input complex* $\mathcal{I}$, a colored *output complex* $\mathcal{O}$, and a recursive relation Δ carrying each m-simplex of $\mathcal{I}$ to a set of m-simplexes of $\mathcal{O}$, for each $0 \le m \le n$. Δ has the following interpretation: if the $(m+1)$ processes named in σ^m start with the designated input values, and the remaining $n-m$ processes fail without taking any steps, then each simplex in $\Delta(\sigma^m)$ corresponds to a legal final state of the non-faulty processes.

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. The protocol complex corresponding to executions starting from an input simplex σ^m is denoted $\mathcal{P}(\sigma^m)$. A protocol *solves* a task if there exists a color-preserving simplicial *decision map* $\delta : \mathcal{P} \to \mathcal{O}$ such that for each simplex $\tau^m \in \mathcal{P}(\sigma^m)$, $\delta(\tau^m) \in \Delta(\sigma^m)$.

A fundamental result of [6, 15, 18] is that for any wait-free protocol there is a protocol complex $\mathcal{P}(\sigma^m)$ which is a chromatic subdivision of σ^m, for every σ^m in $\mathcal{I}$. Thus, the union of these complexes, over all $\sigma^m \in \mathcal{I}$, is a complex $K^n = \mathcal{P}(\mathcal{I})$, that is a chromatic subdivision of $\mathcal{I}$. If the protocol solves the task $< \mathcal{I}, \mathcal{O}, \Delta >$ then each vertex of $\mathcal{P}$ is colored with an output value, defining a decision function δ, which is a color-preserving (on ids) *simplicial map* from K^n to $\mathcal{O}$ that respects Δ. Also, if the protocol is anonymous, the subdivision of σ^m will be symmetric as explained later on. This is one direction of the (Anonymous) Asynchronous Computability Theorem [15]. The theorem says that this is an if and only if characterization of wait-free read/write computability. An algorithmic proof appears in [5].

[4]Sometimes it is convenient permits that a complex contains the empty simplex.

2.2 Combinatorial Topology

The combinatorial properties of the complex $K^n = \mathcal{P}(\mathcal{I})$ that are important for renaming, were identified in [1]: K^n is a symmetric, chromatic, orientable and connected divided image of $\mathcal{I}$. We briefly review what this means.

Pseudomanifolds and orientability.

We say that a complex K^n is *complete* for dimension i, $i \le n$, if every *j-simplex*, $j \le i$, of K^n is a face of at least one *i-simplex* of K^n. If the complex K^n is complete for dimension n then a simplex σ^{n-1} of K^n is *external* if it is a face of exactly one *n-simplex* of K^n, otherwise it is *internal*. The *boundary* of a complex K^n, $bd(K^n)$, is the subcomplex with all the faces of the external simplexes of K^n. A vertex of K^n is *external* if it is contained in a simplex of $bd(K^n)$, otherwise, it is *internal*. A complex K^n is an *n-pseudomanifold* if it is complete for dimension n and every $(n-1)$-simplex of K^n is contained in either one or two n-simplexes of K^n. A pseudomanifold K^n is *chromatic* if it has a coloring *id* with ID^n such that every n-simplex of K^n is properly colored under *id*.

Let σ^n be a simplex. An *orientation* of σ^n, $n > 0$, is a set consisting of a sequence of its vertexes and all even transpositions of this sequence. Hence, there are exactly two possible orientations for σ^n. For example, the two possible orientations of a 2-simplex are the clockwise and counter-clockwise directions, or the two possible orientations of a 1-simplex are the one from one of its vertexes to the other, and the other opposite direction. If σ^n has a properly coloring *id* with colors ID^n then we denote by $d = +1$ the orientation that contains the sequence $\langle 0, 1 \ldots n \rangle$, and denote by $d = -1$ the other orientation of σ^n which contains the sequence $\langle 1, 0 \ldots n \rangle$. For $n = 0$, there is only one sequence of the vertexes of a simplex σ^0, and then it has just one orientation, however we can associate $+1$ or -1 to this orientation. Hence, a 0-simplex also has two orientations. An orientation d of σ^n, $n > 0$, *induces an orientation* to all of its $(n-1)$-faces: σ_i^{n-1} has the orientation $(-1)^i d$, where σ_i^{n-1} is the $(n-1)$-face of σ^n without the vertex with *id i*. A pseudomanifold K^n is *orientable* if there is an orientation for each of its n-simplexes such that if $\sigma_1^n, \sigma_2^n \in K^n$ share a $(n-1)$-face τ then τ gets opposite induced orientations from σ_1^n and σ_2^n. An orientation of K^n like this is a *coherent orientation*.

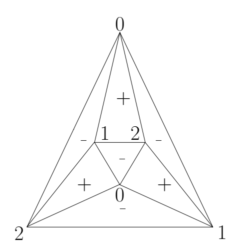

Figure 2: An oriented and chromatic pseudomanifold.

LEMMA 2.1 ([1, LEMMA 5.12]). *A chromatic pseudomanifold K^n is orientable if and only if its n-simplexes can be partitioned into two disjoint classes, such that if two n-simplexes share an $(n-1)$-face then they belong to different classes.*

We denote by $+1$ and -1 the two classes in the previous lemma. Moreover, the n-simplexes of class d have orientation d. Figure 2 contains a coherently oriented and chromatic pseudomanifold of dimension 2. Observe that all the 2-simplexes follow the counterclockwise direction. It is easy to verify that in a coherent orientation of a chromatic and 2-connected pseudomanifold of dimension 2, all the 2-simplexes follow either the clockwise or counterclockwise direction.

Index Lemma.

Let K^n be a coherently oriented pseudomanifold, with an induced orientation on its boundary. For the next definition we count a properly colored n-simplex by orientation with respect to a coloring c. This means that a properly colored n-simplex counts as $+1$ if the order of its vertexes induced by the sequence $\langle 0, 1 \ldots n \rangle$ belongs to its orientation, and as -1 otherwise. In a similar way we count the properly colored $(n-1)$-simplexes on the boundary.

Definition 2.2 (Index and Content). *Let c be a coloring, not necessarily proper, of K^n with ID^n. The* content *of K^n, $C(K^n)$, with respect to c is the number of the properly colored n-simplexes of K^n counted by orientation. The* index *of K^n, $I_i(K^n)$, with respect to c is the number of the properly colored $(n-1)$-simplexes of $bd(K^n)$ with ID_i^n counted by orientation.*

If there is no ambiguity we simply write C^n or I_i^n. The next lemma is the restatement of Corollary 2 in [10] using our notation (see [8] for a simple proof of it).

Lemma 2.3 (Generalized Index Lemma). *Let K^n be a coherently oriented pseudomanifold colored with ID^n. Then $C^n = (-1)^i I_i^n$.*

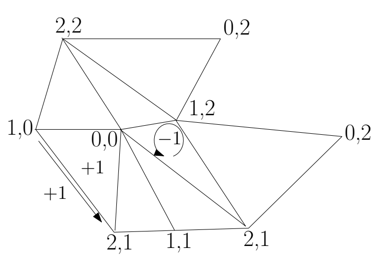

Figure 3: The Index Lemma.

Figure 3 contains a pseudomanifold K^2 in which every vertex has two colorings id and c, in this order. Observe that K^2 is chromatic with respect to id. We assume that the 2-simplexes are counterclockwise oriented and then the unique properly colored 2-simplex τ_1 by c has orientation -1 and the 2-simplex τ_2 that contains the unique properly 1-simplex ρ with 0 and 1 by c on the boundary, has orientation $+1$. Notice that the content C^2 with respect to c is equal to 1 because the c colors 0, 1 and 2 of τ_1 read in the counterclockwise direction, as denoted by the circular arrow. Moreover, observe that ρ has the induced orientation $+1$ by τ_2. That is, the direction from the vertex of ρ with id 1 to the vertex with id 2. We have that $I_2^2 = +1$ because the c colors 0 and 1 of ρ agree with the orientation of ρ, as denoted by the straight arrow. The reader can also check that $(-1)^2 I_2^2 = (-1)^1 I_1^2 = (-1)^0 I_0^2$.

The coloring c is a simplicial map from $bd(K^n)$ to the boundary of a properly colored n-simplex σ^n with ID^n. Thus, we can think of the content of K^n as the number of times that $bd(K^n)$ is wrapped around $bd(\sigma^n)$ (i.e., a combinatorial version of the notion of *degree* in topology).

Divided Images.

Divided images are defined and studied in [1]. These together with the Generalized Index Lemma are the principal tools in our combinatorial characterization of the number of monochromatic simplexes of a pseudomanifold. Let K^n, L^n be complexes and ψ be a function that maps every simplex of L^n to a finite subcomplex of K^n. Figure 4 presents a divided image of dimension 2 which maps the upper 2-simplex of L^2 to the upper subdivision of a 2-simplex of K^2 and the lower 2-simplex to the lower subdivision of a 2-simplex. The complex K^n is a *divided image* of L^n under ψ if: (1) $\psi(\emptyset) = \emptyset$ (2) for every $\tau \in K^n$ exists $\sigma \in L^n$ such that $\tau \in \psi(\sigma)$ (3) for every $\sigma^0 \in L^n$, $\psi(\sigma^0)$ is a vertex, (4) for every $\sigma_1, \sigma_2 \in L^n$, $\psi(\sigma_1 \cap \sigma_2) = \psi(\sigma_1) \cap \psi(\sigma_2)$ and (5) for every $\sigma \in L^n$, $\psi(\sigma)$ is a $dim(\sigma)$-pseudomanifold with $bd(\psi(\sigma)) = \psi(bd(\sigma))$.

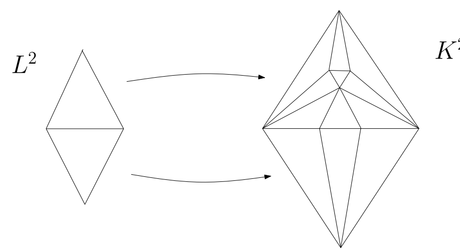

Figure 4: An example of a divided image.

Let K^n be a divided image of σ^n under ψ. We say K^n is *connected* if for every m-face σ of σ^n, if $m \geq 1$ then $\psi(\sigma)$ is m-connected, and if $m \geq 2$ then $bd(\psi(\sigma))$ is $(m-1)$-connected. Similarly, K^n is *orientable* if for every face σ of σ^n, $\psi(\sigma)$ is orientable. Also, we say K^n is *coherently oriented* if the pseudomanifold $\psi(\sigma^n)$ is coherently oriented.

The *carrier* of a simplex $\tau \in K^n$, denoted $carr(\tau)$, is the face σ of σ^n of smallest dimension such that $\tau \in \psi(\sigma)$. We say that a vertex $v \in K^n$ has an *i dimensional carrier* if $dim(carr(\{v\})) = i$.

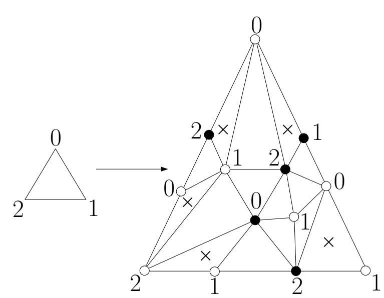

Figure 5: The 2-corners of a chromatic divided image of a 2-simplex with a symmetric binary coloring.

Now, assume that σ^n is properly colored with ID^n. We denote the colors of a face σ of σ^n by $ID(\sigma)$. The divided image K^n has a *Sperner coloring id* if for every vertex $v \in$

K^n, $id(v) \in ID(carr(v))$. Also, K^n is a *chromatic* divided image of σ^n if it has a Sperner coloring and every simplex $\tau \in K^n$ is properly colored with colors of $carr(\tau)$. Notice that for every face σ of σ^n, $\psi(\sigma)$ is a chromatic pseudomanifold with $ID(\sigma)$.

Intuitively, a divided image has *structural symmetry* if any pair of divided images of faces of the same dimension have the same subdivision. Also, K^n has a *symmetric binary coloring* if K^n has structural symmetry and a binary coloring b that preserves it. Figure 5 presents a chromatic divided image with symmetric binary coloring which uses white and black circles to represent the binary colors 0 and 1. Notice that the corners have the same binary color, and also the subdivision of all the edges have the same binary pattern.

Let K^n be a divided image of σ^n under ψ. A *cross edge* of K^n is a 1-simplex $\{u, v\} \in bd(K^n)$ such that $u \in \psi(\tau)$ and $v \in \psi(\rho)$, where τ, ρ are distinct m-faces of σ^n, $0 \leq m \leq n-2$. This implies that if a divided image K^n has no cross edges then the divided image of every proper face of σ^n has at least one internal vertex. Figure 6 contains a divided image of a 2-simplex and a 3-simplex, respectively. A cross edge of the divided image of dimension 2 is the bold edge connecting the bottom corners, and in the other case, the bold edges are cross edges. Assuming no cross edges facilitates our proofs.

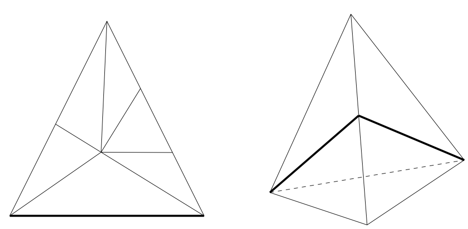

Figure 6: Divided images with cross edges.

The *n-corners* of a divided image K^n are the subset of its n-simplexes that have a face in the boundary of the divide image of some face σ^i of σ^n, for every i, $0 \leq i \leq n$. Figure 5 contains a divided image and its n-corners marked with small crosses.

LEMMA 2.4. *Let K^n be a chromatic, connected and coherently oriented divided image of σ^n under ψ. Then n-corners$(K^n) \neq \emptyset$ and every simplexes σ_1^n, σ_2^n of n-corners(K^n) have the same orientation.*

The following theorem summarizes the combinatorial properties of an anonymous wait-free protocol, whose outputs are either 0 or 1, assuming (1) the shared objects are immediate snapshots, (2) all processes execute at least one immediate snapshot operation (to guarantee that there are no cross edges) and (3) starts on a single input simplex σ^n. It follows directly from Theorem 5.14 in [1].

THEOREM 2.5 ([1]). *Let P^n be the complex of a protocol as described above. Then P^n is a chromatic, orientable and connected divided image of a simplex σ^n, with a symmetric binary coloring and no cross edges.*

In the particular case of the WSB task, using the Anonymous Asynchronous Computability Theorem [15] (see also [5]) we get the other direction also:

THEOREM 2.6. *There exists a WSB protocol if and only if there exists a chromatic, orientable and connected divided image K^n of a simplex σ^n, with a symmetric binary coloring and no cross edges with no monochromatic n-simplexes.*

3. A COMBINATORIAL CHARACTERIZATION OF THE NUMBER OF MONOCHROMATIC SIMPLEXES

In this section we present a characterization of the number of monochromatic n-simplexes of a chromatic, oriented and connected divided image K^n of σ^n under ψ, with a symmetric binary coloring b and no cross edges. Our characterization is obtained via an inductive process described in detail in [8]. The strategy is to start with a binary coloring equal to 0 on the boundary and modify it step by step until it is equal to the original binary coloring b, as explained below. Finally, we count the monochromatic simplexes using the Generalized Index Lemma. An alternate proof strategy suggested by Eli Gafni is described in [8]. It consists of inserting K^n inside the boundary of an n-simplex (by doing the join of two complexes; the definition of join is below), and computing the number of monochromatic n-simplexes in K^n by analyzing the number of monochromatic n-simplexes generated by this construction.

As in previous renaming lower bound proofs [13, 15], we add a third coloring c to K^n that will be used to count the monochromatic simplexes using the Generalized Index Lemma. This is done through the following function. Let K^n be a complex with a binary coloring b and a proper coloring id with colors ID^n. A function $f : ID^n \to ID^n$ is a *minimal permutation* of ID^n if f is a permutation (a one-to-one function) of ID^n and the restriction $f|_I$ to any proper subset $I \subset ID^n$ is not a permutation of I.

DEFINITION 3.1. *For a minimal permutation f of ID^n let*

$$c(v) = \begin{cases} id(v) & \text{if } b(v) = 0 \\ f(id(v)) & \text{if } b(v) = 1 \end{cases}$$

The next lemma shows the relation between b and c.

LEMMA 3.2. *An n-simplex of K^n is monochromatic under b if and only if it is properly colored under c.*

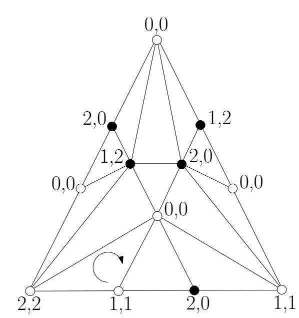

Figure 7: A pseudomanifold with the three colorings.

Figure 7 contains a 2-pseudomanifold with the three colorings, id, b and c, associated to each vertex for $ID^2 = \{0, 1, 2\}$. The binary coloring b is represented by white and black circles, and the *id* and *c* colorings, in this order, are the numbers near to the vertexes. Every 2-simplex is counterclockwise oriented and *c* uses the minimal permutation $f(x) = (x+1)\, mod\, 3$. The simplex with an arrow is the unique monochromatic simplex and it is properly colored under *c*. The direction of the arrow fallows the order of the vertexes with respect to *c*. As this order is opposite to the orientation of the simplex, It is counted -1 by C^2 with respect to *c*.

The next lemma implies that a 0-monochromatic n-simplex with orientation d is counted as d, but a 1-monochromatic n-simplex with orientation d is counted as d if n is even and as $-d$ if n is odd.

LEMMA 3.3. *Let τ be a b-monochromatic n-simplex of K^n with orientation d. The simplex τ is counted as $(-1)^{b*n}d$ by C^n.*

The inductive process over K^n starts with a binary coloring on the boundary equal to 0.

LEMMA 3.4. *If for every vertex v of $bd(K^n)$, $b(v) = 0$ then $I_i^n = (-1)^i$, assuming that the n-corners of K^n have the orientation +1.*

Then, we *process* groups of vertexes (change their binary color to 1) of $bd(K^n)$ with carriers of the same dimension, until we get the original given binary coloring. This action, i.e. processing of a group of vertexes with carrier of dimension ℓ, is called *ℓ-step,* and it may be done more than once in each dimension ℓ. A step guarantees that after executing it the binary coloring b of K^n remains symmetric. Steps are done by dimension: a vertex with carrier of dimension $\ell+1$ is processed if and only if every vertex with carrier of dimension ℓ has its correct binary color. For example, for dimension 3, we first process (if it is necessary) the corners, then the vertexes inside the subdivision of the edges, and finally the vertexes inside the subdivision of the triangles. Notice that the vertexes inside the subdivision of the tetrahedron are not modified and actually it does not matter their coloring. A crucial argument in the proof is to analyze how the steps affect the index of K^n. We can prove that all the changes in a step affect the index in the same way. In the next lemma, the content of a set of vertexes $v_1, \ldots, v_q$ of $\psi(\sigma_i^{n-1})$ is the number of properly colored $(n-1)$-simplexes with ID_i^n of $st(v_1, \psi(\sigma_i^{n-1})) \cup \ldots \cup st(v_q, \psi(\sigma_i^{n-1}))$, counted as in the definition of content, Definition 2.2.

LEMMA 3.5. *Let I_i^n be the index of K^n before every ℓ-step in the process is done and $\hat{I}_i^n$ be the index after all these ℓ-steps are done. For the i-th ℓ-step pick a vertex v_i of $\psi(\sigma_i^{n-1})$ processed in this step. Then $\hat{I}_i^n = I_i^n - \binom{n+1}{\ell+1}k_\ell$, where k_ℓ is the content with ID_i^n of the v_i vertexes before every ℓ-step is done.*

Figure 8 presents an example of the inductive process for $n = 2$. The vertexes have the colorings, *id*, *b* and *c*, associated to each vertex for $ID^2 = \{0, 1, 2\}$. The binary coloring *b* is represented by white and black circles, and the *id* and *c* colorings, in this order, are the numbers near to the vertexes. Also, *c* is with respect to $f(x) = (x+1)\ mod\ 3$. The -1 or $+1$ inside of a 2-simplexes τ is its orientation, the arrow is the induced orientation by τ to the 1-simplex ρ on the boundary, and the -1 or $+1$ outside of τ is the way in which I_0^2 counts ρ. The process begins with a binary coloring equal to 0 on the boundary, Figure 8(a). The index at the beginning of the process is equal to 1. The process has a 0-step presented in Figure 8 (b). Observe that this step adds a multiple of three to the index because a 2-simplex has three 0-faces. Figure 8 (c) contains one 1-step which adds a multiple of three to the index because a 2-simplex has three 1-faces. Something similar happens with the 1-step contained in Figure 8 (d).

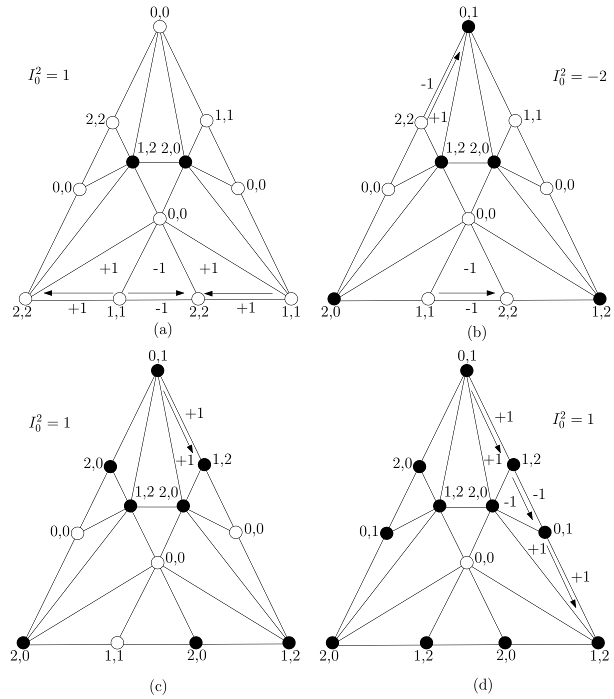

Figure 8: An example of the inductive process.

The main result of the section is the following theorem which is a characterization of the number of monochromatic n-simplexes of K^n counted as in Lemma 3.3 (see Lemma 3.2).[5]

THEOREM 3.6. *Let K^n be a chromatic, connected and coherently oriented divided image of σ^n, with a symmetric binary coloring and without cross edges. Then $C(K^n) = 1 + \sum_{i=0}^{n-1} \binom{n+1}{i+1} k_i$ for $k_i \in \mathbb{Z}$, $k_0 \in \{0, -1\}$, and assuming that the n-corners of K^n have the orientation +1.*

4. CONSTRUCTING DIVIDED IMAGES

Consider integers $k_0, k_1 \ldots k_{n-1}$, with $k_0 \in \{0, -1\}$, and let $x = 1 + \sum_{i=0}^{n-1} \binom{n+1}{i+1} k_i$. In Section 4.1 we show how to construct a divided image K^n of σ^n under ψ, as assumed in Theorem 3.6, with x monochromatic n-simplexes counted by orientation. In fact, the content of this divided image will be $(-1)^n x$ due to Lemma 3.3. Also, in Section 4.2 we present an algorithm that modifies the divided image, to make sure it has exactly x monochromatic n-simplexes all oriented $sign(x)$. The construction and the algorithm are generalized in [8] to any n and x. We explain in Section 5

[5] K^n does not actually need to be connected, but for simplicity we leave the assumption in this paper.

why the construction together whit the algorithm imply a renaming protocol for $K = 2n - 1$. First, some basic tools are presented.

Let σ and τ be two simplexes. The *join* of σ and τ, $\sigma * \tau$, is the simplex $\sigma \cup \tau$. If σ and τ are properly colored under id then we say that σ and τ are *compatible* if $id(\sigma) \cap id(\tau) = \emptyset$. Now, if K and L are properly colored complexes under id then the *join* of K and L, $K * L$, is the complex $\{\sigma * \tau | \sigma \in K, \tau \in L$ and σ and τ are compatible$\}$, assuming that every complex contains the empty simplex (the general definition of join complex does not require compatibility). Notice that due to the previous assumption, $K, L \subset K * L$.

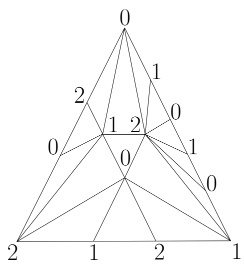

Figure 9: The cone construction.

Assume we are given $\psi(bd(\sigma^n))$, a chromatic, connected and orientable divided image of $bd(\sigma^n)$, and a properly colored n-simplex τ with ID^n. The *cone over* $\psi(bd(\sigma^n))$ *for* τ is obtained by putting the n-simplex τ at the center of $\psi(bd(\sigma^n))$ and joining every i-simplex of $\psi(\sigma^i)$, $0 \leq i \leq n-1$, with the simplexes of $M(\tau)$ with compatible colors, formally, it is the complex consisting of the simplexes in the set $\{\sigma * \rho \mid \sigma \in \psi(bd(\sigma^n)), \rho \in M(\tau), \sigma$ and ρ are compatible and $dim(\sigma) = dim(carr(\sigma))\}$, and all their faces. Figure 9 presents an example of the cone construction of dimension 2. This construction can be used to build inductively a divided image.

Lemma 4.1. *The cone $\psi(\sigma^n)$ is a chromatic, orientable and connected divided image of σ^n.*

The next subdivision is a particular case of the cone construction. The *basic chromatic subdivision* of two properly colored n-simplexes σ and τ with ID^n, is the cone over $bd(M(\sigma))$ for τ. The basic chromatic subdivision of a properly colored 2-simplex with colors ID^2 is illustrated in Figure 2.

Corollary 4.2. *A basic chromatic subdivision of an n-simplex σ is a chromatic, orientable and connected pseudomanifold.*

4.1 Constructing Divided Images Counting by Orientation

We describe here how to construct a divided image K^5 with $x = 0$ (0 monochromatic n-simplexes counted by orientation). This gives a counterexample to Theorem 6.2[1]. The construction is generalized in [8] to any n and x and also it implies the next theorem:

Theorem 4.3. *Let $k_0, k_1 \ldots k_{n-1}$ be integers such that $k_0 \in \{0, -1\}$. There exists a chromatic, orientable and connected divided image K^n of σ^n, with symmetric binary coloring, such that it has $1 + \sum_{i=0}^{n-1} \binom{n+1}{i+1} k_i$ monochromatic n-simplexes counted by orientation.*

For our example with $x = 0$, let $k_0 = -1, k_1 = -1, k_2 = 1, k_3 = 0, k_4 = 0$, as $1 - \binom{6}{1}k_0 - \binom{6}{2}k_1 + \binom{6}{3}k_2 + \binom{6}{4}k_3 + \binom{6}{5}k_4 = 1 + 6k_0 + 15k_1 + 20k_2 + 15k_3 + 6k_4 = 0$. The idea is to construct K^5 such that it has 42 monochromatic 5-simplexes, partitioned in two sets A^- and A^+: 21 negatively oriented and 21 positively oriented. We use the cone construction. Recall that in this construction, for every proper face σ^i of σ^5, every i-simplex of $\psi(\sigma^i)$ generates a 5-simplex in K^5. At the end of the construction, K^5 has a 1-monochromatic 5-simplex τ^5 at the center, and we orient it positively. The simplex τ^5 is 1-monochromatic, and every other vertex of K^5 is colored 0 unless explicitly stated otherwise below.

We construct K^5, the divided image of σ^5 under ψ, inductively by dimension. Recall that we denote the colors of a face σ of σ^5 by $ID(\sigma)$. First, for each face σ^0 of σ^5, $\psi(\sigma^0)$ is a vertex v with $id(v) = ID(\sigma^0)$ and $b(v) = 1$ (because $k_0 = -1$). Each such v is going to generate a 1-monochromatic 5-simplex τ_v in K^5. There are $\binom{6}{1}$ such 1-monochromatic simplexes, and they all have the same orientation in a coherent orientation of K^5, because they all are adjacent to τ^5 ($\tau_v \cap \tau^5$ is a 4-simplex), by Lemma 2.1. Namely, they all are negatively oriented, as we assume τ^5 is positively oriented. Thus, we have so far six 5-simplex in A^- and one in A^+.

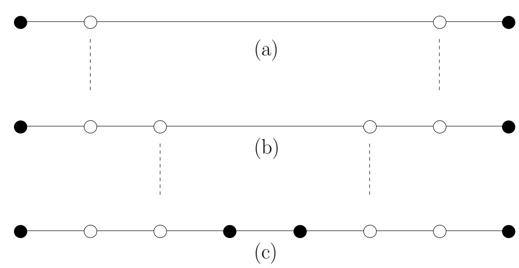

Figure 10: The construction for k_1.

The general idea, described in more detail below for this example, is to assume we have constructed $\psi(bd(\sigma^i))$, for every $\sigma^i \subset \sigma^5$. Then, take a face σ^i of σ^n, and construct the cone over $\psi(bd(\sigma^i))$ with some new simplex, τ^i. Once this is done, we chromatically subdivide $\psi(\sigma^i)$ iteratively using the basic chromatic subdivision, until $|k_i|$ internal i-simplexes are produced with the orientation required by $sign(k_i)$. We repeat the same construction for each such face, preserving id ranking, to make sure $\psi(\sigma^i)$ is symmetric. By Lemma 4.1 and Corollary 4.2, $\psi(\sigma^i)$ is chromatic, connected and orientable pseudomanifold. After repeating the procedure for $i = 1, 2, 3, 4, 5$, we obtain K^5, a chromatic, connected and orientable divided image of σ^5. Also, by construction, it has a symmetric binary coloring and does not have cross edges. To check that is has zero 1-monochromatic 5-simplexes by orientation, we will use the fact that every 5-simplex of K^5 contains at least one vertex of τ^5. Thus, K^5 does not contain 0-monochromatic 5-simplexes.

The case of $i = 1$ is illustrated in Figure 10(a) (recall that the binary coloring is represented with white and black circles). The goal is to create exactly one 1-monochromatic 1-simplex in the interior of $\psi(\sigma^1)$ (because $|k_1| = 1$) with a specific orientation, and we want that the 1-monochromatic 5-simplex generated by this 1-simplex in the cone construction is negatively oriented (because $sign(k_1) = -$). We cannot color τ^1 with 1 because we would create two other 1-monochromatic simplexes (the 1-corners), and also we cannot color with 1 just one of the vertexes of τ because the 1-monochromatic 1-simplex created would not have the correct orientation. Thus we subdivide chromatically $\psi(\sigma^1)$ once, as in Figure 10(b), but here we still cannot obtain exactly one 1-monochromatic 1-simplex with the correct orientation, so we subdivide it once more, as in Figure 10(c), and we color 1 the 1-simplex in the center, call it τ^1. It is easy to verify that the induced 5-simplex from τ^1 by the cone construction is negatively oriented. We have 15 more simplexes in A^-, because the same construction (preserving $id's$ ranking) is done for each face of dimension 1 of σ^5, and hence there will be $\binom{6}{2}$ induced 1-monochromatic 5-simplexes, and with the same orientation in a coherent orientation of K^5 (they all are at the same distance from τ^5, Lemma 2.1). And we get a a chromatic, orientable and connected divided image $\psi(sk^1(\sigma^5))$, with symmetric binary coloring ($skel^i$ is the subcomplex of all faces of dimension at most i).

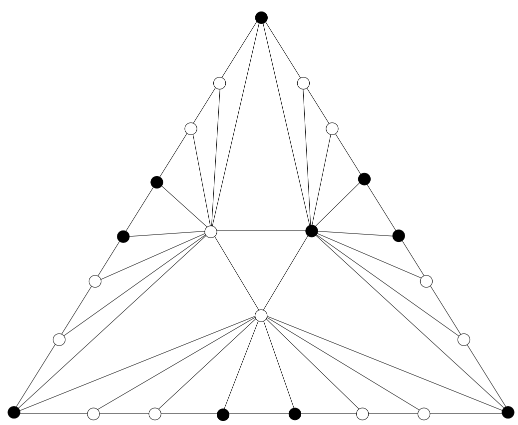

Figure 11: The construction for k_2.

We repeat this procedure to construct $\psi(sk^2(\sigma^5))$. For a 2-face σ^2 of σ^5, do the cone over $\psi(\sigma^2)$ (we do the same for each such face, preserving id ranking). Figure 11 presents an example of $\psi(\sigma^2)$. Now we need exactly one 1-monochromatic 2-simplex in the interior of $\psi(\sigma^2)$ (because $|k_2| = 1$), and we want that the 1-monochromatic 5-simplex generated by τ^2 by the cone construction is positively oriented (because $sign(k_2) = +$). There are $\binom{6}{3}$ such 1-monochromatic simplexes, and they all have the same orientation in a coherent orientation of K^5, because they all are at the same distance from τ^5. In this case we just have to color with 1 exactly one vertex of the 2-simplex at the center to obtain one 1-monochromatic 2-simplex τ^2 with the correct orientation. It is easy to verify that the induced 5-simplex from τ^2 by the cone construction is positively oriented, so we have 20 more simplexes in A^+. We get a chromatic, orientable and connected divided image $\psi(sk^2(\sigma^5))$, with symmetric binary coloring.

Finally, since $k_3 = k_4 = 0$, for $3 \le i \le 5$, $\psi(\sigma^i)$ is cone constructed with τ^i, without further subdivisions and no internal vertexes colored 1. Thus, $|A^-| = |A^+| = 21$.

Consider the content of K^5, C^5, with respect to a coloring c as presented in Definition 3.1. By Lemma 3.3, C^n counts a 1-monochromatic n-simplex with orientation d as d if n is even, and as $-d$ if n is odd. Since K^5 has only 1-monochromatic 5-simplexes and we have counted these simplexes by orientation, then $C^5 = -x$.

4.2 Eliminating monochromatic simplexes

In this section we present an algorithm which takes as input an orientable and chromatic pseudomanifold K^n with x monochromatic n-simplexes counted by orientation and modifies K^n, using basic chromatic subdivisions and without touching its boundary, so that it has exactly $|x|$ monochromatic n-simplexes, all oriented $sign(x)$ (see [8] for a detailed description). We can assume, as explained below, that K^n has no 0-monochromatic n-simplexes. This algorithm together with Theorem 4.3 implies:

THEOREM 4.4. *Let $k_0, k_1 \dots k_{n-1}$ be integers such that $k_0 \in \{0, -1\}$, and $x = 1 + \sum_{i=0}^{n-1} \binom{n+1}{i+1} k_i$. There exists a chromatic, orientable and connected divided image K^n of σ^n, with symmetric binary coloring, such that it has exactly $|x|$ monochromatic n-simplexes all oriented $sign(x)$.*

```
eliminate(K^n)
while ∃ 1-monochromatic n-simplexes with opposite orientation
    let P be a path in standard form
    eliminatePath(P)
```

Figure 12: Algorithm eliminate.

First we present some definitions. Let P be an n-path of size $|P| = 2(q+1)$, $P : S_0 - S_1 \cdots S_{2q+1}$, $q \ge 0$, such that the n-simplexes at its ends S_0 and S_{2q+1} are 1-monochromatic and it has no other monochromatic n-simplex. We say that P is in *standard* form. Notice that S_0 and S_{2q+1} have opposite orientation and then the content of P, $C(P)$, with respect to a coloring c as presented in Definition 3.1, is equal to zero and hence the index $I(P)$ is equal to zero by the Generalized Index Lemma 2.3. We denote by $S_{a,a+1}$ the $(n-1)$-face shared by S_a and S_{a+1}, by $x(S)$ the face of S with all the vertexes of S with binary color x and by $\#x(S) = |x(S)|$. A *good chromatic subdivision* of P is a chromatic subdivision $sub(P)$ that contains two disjoint (in the sense that they do not share n-simplexes) paths P_1 and P_2 in standard form, has no other monochromatic n-simplex and $bd(P) = bd(sub(P))$. And a *complete chromatic subdivision* of P is a chromatic subdivision $sub(P)$ that has no monochromatic n-simplex and $bd(P) = bd(sub(P))$. For all the figures in this section the id colors are the numbers near to the vertexes.

```
eliminatePath(P : S_0 − S_1 ··· S_{2q+1})
if |P| = 2 then
    subdivideComp(P)
else
    let m ← 0
    while true do
        if #1(S_{m+1,m+2}) ≥ n + 1 − m then
            P1, P2 ← subdivideGood(P, m)
            eliminatePath(P1)
            eliminatePath(P2)
            break       %end of while loop%
        m + +
```

Figure 13: Algorithm eliminatePath.

Algorithm **eliminate** (Figure 12) uses Algorithm **eliminatePath** (Figure 13) to produce a complete chromatic subdivision of a path in standard form. Notice that at the end of the execution of **eliminate**, K^n has exactly $|x|$ monochromatic n-simplexes, all oriented $sign(x)$ (recall that K^n has x monochromatic n-simplexes counted by orientation).

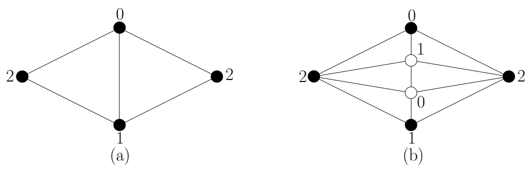

Figure 14: A path of size two.

We now describe how Algorithm **eliminatePath** produces a complete chromatic subdivision of a path P in standard form. If the input has size two, that is, P consists only of two 1-monochromatic n-simplexes, then it just subdivides their shared $(n-1)$-face with a 0-monochromatic n-simplex to obtain a complete chromatic subdivision of P. The function **subdivideComp** (line 2) does this subdivision. Figure 14 (a) presents en example of dimension 2 of this situation and Figure 14 (b) presents the result of **subdivideComp**.

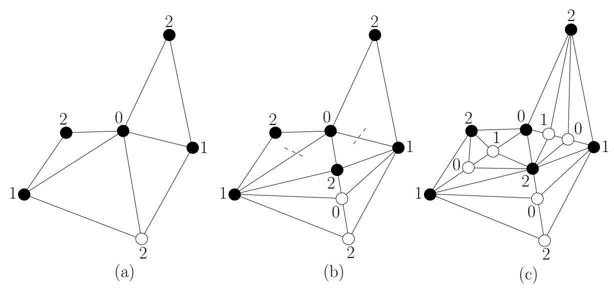

Figure 15: A path of size four.

Now, if P has size greater than two then **eliminatePath** inspects one by one the shared $(n-1)$-faces of P starting with $S_{1,2}$, to find the *cutting point* in which the algorithm will subdivide a shared $(n-1)$-face. The cutting point is identified when line 6 is true and the face that will be subdivided is $S_{m-1,m}$ or $S_{m,m+1}$. Also, **eliminatePath** subdivides this face using function **subdivideGood**. which creates a good chromatic subdivision of P such that P_1 and P_2 have sizes smaller or equal than $|P|$. The idea behind this subdivision is that since the boundary has to be the same, and so $I(sub(P)) = 0$, we can create exactly two 1-monochromatic n-simplexes with opposite orientation into P (adding zero to the index) to obtain two new paths in standard form. Finally, **eliminatePath** recursively calls itself on the two paths P_1 and P_2. Now, if $|P_i| < |P|$, $i \in \{1, 2\}$ then we say that the cutting point was *progressive*. We prove in [8], if $|P_i| = |P|$, for some $i \in \{1, 2\}$, then P_i always has a progressive cutting point. Observe that this condition guarantees progress of the algorithm.

We now present some examples of the subdivisions obtained by **eliminatePath**. Figure 15 (a) presents a path of size four which has the cutting point on 1 (i.e., $\#1(S_{2,3}) \geq n = 2$). **eliminatePath** subdivides the face $S_{1,2}$, Figure 15 (b), to add two 1-monochromatic n-simplexes and thus create two disjoint paths in standard form of size two. Then, the two resulting paths will be subdivided as in the case presented in Figure 14 on the next recursive call of **eliminatePath** on them, Figure 15 (c).

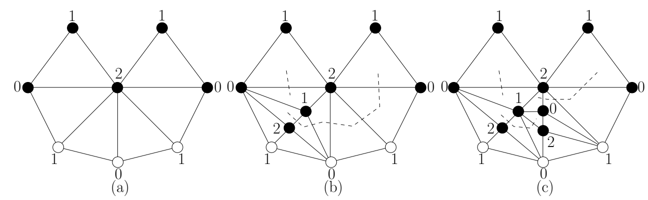

Figure 16: A path of size six.

As mentioned above, there exists cases which **eliminatePath** cannot obtain two paths of size smaller than the original path. Consider the path of size 6 in Figure 16 (a). Observe that it has the cutting point on 2 (i.e., $\#1(S_{3,4}) \geq n - 1 = 1$). In this case **eliminatePath** subdivides the face $S_{1,2}$ to obtain one path of size smaller than 6 and another path P' of size 6, Figure 16 (b). However, on the next recursive call on P', **eliminatePath** always can obtain paths of size 4, Figure 16 (c).

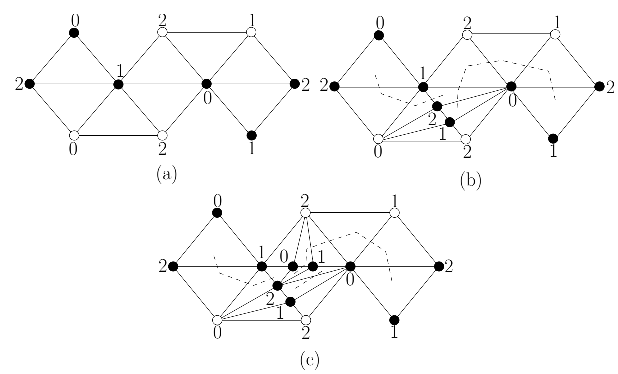

Figure 17: A path of size eight.

Sometimes **eliminatePath** does not obtain disjoint paths in the subdivision of a face in the cutting point, and then it has to do an extra subdivision. Figure 17 (a) presents a path of size 8 with the cutting point on 2. The subdivision of the face $S_{2,3}$, Figure 17 (b), generates two paths of size smaller than 8 but they share one 2-simplex. Therefore, **elminatePath** has to do an extra subdivision, Figure 17 (c), to create two paths of size 6 and one path of size 2 which is eliminated using **subdivideComp**.

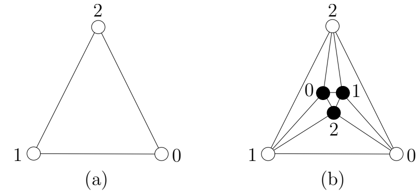

Figure 18: Changing the color.

Finally, we have been focusing on 1-monochromatic n-simplexes because a 0-monochromatic n-simplex can be easily transformed to a 1-monochromatic n-simplex using a basic chromatic subdivision, Figure 18.

5. CONCLUSIONS

One direction of Theorem 2.6 states that if there is a WSB protocol then there exists a chromatic, orientable and connected divided image K^n of a simplex σ^n, with a symmetric binary coloring and no cross edges with no monochromatic n-simplexes. That is, with $C(K^n) = 0$, by Lemma 3.2. However, by Theorem 3.6

$$C(K^n) = 1 + \sum_{i=0}^{n-1} \binom{n+1}{i+1} k_i$$

That is, if there is a WSB protocol then the linear Diophantine equation

$$\binom{n+1}{1} k_0 + \binom{n+1}{2} k_2 + \ldots + \binom{n+1}{n} k_{n-1} = 1 \quad (1)$$

has a solution. It is well known that there exist integers $k_i \in \mathbb{Z}$, $0 \leq i \leq n-1$, which satisfy (1) [9] if and only if $\binom{n+1}{1}, \binom{n+1}{2} \ldots \binom{n+1}{n}$ are relatively prime. Since $\binom{n+1}{i+1} = \binom{n+1}{n-i}$, we focus on $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$. We have the following result which is a special case of Theorem 6.2 in [13] and Theorem 6.3 in [15], recalling the WSB is equivalent to K-renaming, $K = 2n - 1$.

Corollary 5.1. *If $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$ are not relatively prime then there does not exist an anonymous wait-free protocol that solves the K-renaming, $K < 2n$.*

For example, it is easy to check that if $n+1$ is prime then $C^n \equiv 1 \; mod \, (n+1)$. Also, we can easily verify $C^3 \equiv 1 \; mod \, 4$. Therefore, if $n+1$ is prime or $n = 3$ then there does not exists an anonymous wait-free protocol that solves the K-renaming, $K < 2n$.

Now, by Theorem 4.4, if $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$ are relatively prime then there exists a chromatic, orientable and connected divided image K^n, with a symmetric binary coloring and without monochromatic n-simplexes. The other direction of Theorem 2.6 implies that there is a WSB protocol. As the $(2n-1)$-renaming and the weak-symmetry breaking are equivalent [11] we have

Corollary 5.2. *If $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$ are relatively prime then there exists an anonymous wait-free protocol that solves K-renaming, $K = 2n - 1$.*

Consider a prime $p \geq 3$. If $n+1 = 2p$ then $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$ are relatively prime. Just notice that $\binom{2p}{1} = 2p$ and 2 is not factor of $\binom{2p}{2} = p(2p-1)$. Also, it is well known that $\binom{2p}{p} = \frac{2^p(1*3*5\ldots(2p-1))}{p!} = \frac{2^p(1*3*5\ldots(p-2)*(p+2)\ldots(2p-1))}{(p-1)!}$ and hence p is not factor of $\binom{2p}{p}$. Thus, we have an infinite number of cases for which WSB, or a K-renaming protocol exists, $K = 2n - 1$.

Theorem 5.3 (Main). *There exists an anonymous wait-free protocol that solves K-renaming for $K < 2n$ if and only if $\{\binom{n+1}{i+1} : 0 \leq i \leq \lfloor \frac{n-1}{2} \rfloor\}$ are relatively prime.*

Acknowledgments

We would like to thank Eli Gafni for enlightening discussions, his involvement at the early stages of this work and his suggestion of an alternate proof strategy for Theorem 3.6 described in [8].

6. REFERENCES

[1] H. Attiya & S. Rajsbaum. *The Combinatorial Structure of Wait-Free Solvable Tasks.* SIAM Journal on Computing 31(4): 1286–1313 (2002).

[2] H. Attiya, A. Bar-Noy & D. Dolev. *Sharing Memory Robustly in Message-Passage Systems* Journal of the ACM 42(1): 124–142 (1995).

[3] H. Attiya, A. Bar-Noy, D. Dolev, D. Peleg & R. Reischuck. *Renaming in Asynchronous Environment.* Journal of the ACM 37(3): 524–548 (1990).

[4] H. Attiya & J. Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics.* McGraw-Hill 1998.

[5] E. Borowsky & E. Gafni. *A Simple Algorithmically Reasoned Characterization of Wait-Free Computations.* In: Proceedings of the 16th. Annual ACM Symposium on Principles on Distributed Computing, pp. 189–198 (1997).

[6] E. Borowsky & E. Gafni. *Generalized FLP Impossibility Result for t-Resilient Asynchronous Computations.* In: Proceedings of the 25th. Annual ACM Symposium on Theory of Computing, pp. 91–100 (1993).

[7] E. Borowsky & E. Gafni. *Immediate Atomic Snapshots and Fast Renaming.* In: Proceedings of the 12th. Annual ACM Symposium on Principles on Distributed Computing, pp. 41–51 (1993).

[8] A. Castañeda & S. Rajsbaum. *New Combinatorial Topology Upper and Lower Bounds for Renaming.* In: Publicación Preliminar del Instituto de Matemáticas de la Universidad Nacional Autónoma de México (June 9, 2008). http://www.matem.unam.mx/~rajsbaum/podc08.html

[9] J. B. Dence & T. P. Dence. *Elements of the Theory of Numbers.* Academic Press 1999.

[10] K. Fan. *Simplicial Maps from an Orientable n-Pseudomanifold into S^m with the Octahedral Triangulation.* Journal of Combinatorial Theory 2: 588-602 (1967).

[11] E. Gafni, S. Rajsbaum & M. Herlihy. *Subconsensus tasks: renaming is weaker than set agreement* In: 20th Int. Symp. Distributed Computing (DISC), Lecture Notes in Computer Science 4167, Springer, pp.329–338, (2006).

[12] M. Henle. *A Combinatorial Introduction to Topology.* Dover 1994.

[13] M. Herlihy & S. Rajsbaum. *Algebraic Spans.* Mathematical Structures in Computer Science 10(4): 549–573 (2000).

[14] M. Herlihy & N. Shavit. *The asynchronous computability theorem for t-resilient tasks.* In: Proc. of the 25th ACM Symposium on the Theory of Computing, 111–120 (1993).

[15] M. Herlihy & N. Shavit. *The Topological Structure of Asynchronous Computability.* Journal of the ACM 46(6): 858–923 (1999).

[16] H. Hopf. *Abbildungklassen n-dimensionaler Mannigfaltigkeiten.* Math. Ann. 96: pp. 209–224 (1927).

[17] J. R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley 1993.

[18] M. Saks & F. Zaharoglou, *Wait-Free k-Set Agreement Is Impossible: The Topology of Public Knowledge.* SIAM J. on Computing 29(5), 1449–1483 (2000).