

A Topological Treatment of Early-deciding Set-agreement

Rachid Guerraoui^{1,2}, Maurice Herlihy³, and Bastian Pochon¹

¹ School of Computer and Communication Sciences, EPFL

² Laboratory of Computer Science and Artificial Intelligence, MIT

³ Computer Science Department, Brown University

Abstract. This paper considers the k -set-agreement problem in a synchronous message passing distributed system where up to t processes can fail by crashing. We determine the number of communication rounds needed for all correct processes to reach a decision in a given run, as a function of k , the degree of coordination, and $f \leq t$ the number of processes that actually fail in the run. We prove a lower bound of $\min(\lfloor f/k \rfloor + 2, \lfloor t/k \rfloor + 1)$ rounds. Our proof uses simple topological tools to reason about runs of a full information set-agreement protocol. In particular, we introduce a topological operator, which we call the *early deciding* operator, to capture rounds where k processes fail but correct processes see only $k - 1$ failures.

Keywords: Set-agreement, topology, time complexity, lower bound, early global decision.

1 Introduction

This paper studies the inherent trade-off between the degree of coordination that can be obtained in a synchronous message passing distributed system, the time complexity needed to reach this degree of coordination in a given run of the system, and the actual number of processes that crash in that run. We do so by considering the time complexity of the k -set-agreement [3] (or simply set-agreement) problem. The problem consists for the processes of the system, each starting with its own value, possibly different from all other values, to agree on less than k among all initial values, despite the crash of some of the processes. The problem is a natural generalization of consensus [9], which correspond to the case where $k = 1$.

Most studies of the time complexity of k -set-agreement focused on *worst-case global decision* bounds. Chaudhuri et al. in [4], Herlihy et al. in [14], and Gafni in [10], have shown that, for any k -set agreement protocol tolerating at most t process crashes, there exists a run in which $\lfloor t/k \rfloor + 1$ communication rounds are needed for all correct (non-crashed) processes to decide. This (worst-case global decision) bound is tight and there are indeed protocols that match it, e.g., [4].

This paper studies the complexity of *early global decisions* [5]. Assuming a known maximum number of t processes that may crash, early-deciding protocols are those that takes advantage of the effective number $f \leq t$ of failures in any run. In particular, for runs where f is significantly smaller than t , such protocols are appealing for it is often claimed that it is good practice to optimize for the best and plan for the worst.

More specifically, assuming a maximum number t of failures in a system of $n + 1$ processes, we address in this paper the question of how many communication rounds are needed for all correct (non-crashed) processes to decide (i.e., to reach a *global decision*) in any run of the system where f processes fail. Interestingly, there is a protocol through which all correct processes decide within $\min(\lfloor f/k \rfloor + 2, \lfloor t/k \rfloor + 1)$ rounds in every run in which at most f processes crash [11].

We prove in this paper a lower bound of $\min(\lfloor f/k \rfloor + 2, \lfloor t/k \rfloor + 1)$ on the round complexity needed to reach a global decision in any run in which at most f processes crash. The bound is thus tight. Our result generalizes, on the one hand, results on worst-case global decisions for set agreement [4, 14], and on the other hand, results on early global decisions for consensus [16, 2]. As we discuss in the related work section, our bound is also complementary to a recent result on early *local* decisions for set-agreement [11] with an unbounded number of processes.

To prove our lower bound result, we use the topological notions of *connectivity* and *pseudo-sphere*, as used in [14], and we combine them with a mathematical object which we introduce and which we call the *early-deciding* operator. This combination provides a convenient way to describe the topological structure of a bounded number of rounds of an early-deciding full information synchronous message-passing set-agreement protocol.

We prove our result by contradiction. Roughly speaking, we construct the *complex* (set of points in an Euclidean space) representing a bounded number of rounds of the protocol, where k processes crash in each round, followed by a single round in which k processes crash but no process sees more than $k - 1$ crashes. In a sense, we focus on all runs where processes see a maximum of k failures in each round, except in the last round where they only see a maximum of $k - 1$ failures. Interestingly, even if all failures are different, all correct processes need to decide in this round (to comply with the assumption, by contradiction, of $(\lfloor f/k \rfloor + 1)$). We prove nevertheless that the *connectivity* of the resulting complex is high enough, and this leads directly to show that not all correct processes can decide in that complex, without violating the safety properties of k -set-agreement.

Roadmap. The rest of the paper is organized as follows. Section 2 discusses the related work. Section 3 gives an overview of our lower bound proof. Section 4 presents our model of computation. Section 5 presents some topological preliminaries, used in our lower bound proof. Section 6 presents the actual proof. Section 7 concludes the paper with an open problem.

2 Related work

The set-agreement problem was introduced in 1990 by Chaudhuri in [3]. Chaudhuri presented solutions to the problem in the asynchronous system model where $k - 1$ processes may crash, and gave an impossibility proof for the case where at least k processes might crash, assuming a restricted class of distributed protocols called *stable vector protocols*.

In 1993, three independent teams of researchers, namely Herlihy and Shavit [15], Borowsky and Gafni [1], and Saks and Zaharoglou [18], proved, concurrently, that k -set-agreement is impossible in an asynchronous system when k processes may crash. All used topological arguments for showing the results. (Herlihy and Shavit later introduced in [15] a complete topological characterization of asynchronous shared-memory runs, using the concept of algebraic spans [13] for showing the sufficiency of the characterization.)

Chaudhuri et al. in [4], and Herlihy et al. in [14], then investigated the k -set-agreement problem in the synchronous message-passing system, and established that, any k -set-agreement protocol tolerating at most t process crashes, has at least one run in which $\lfloor t/k \rfloor + 1$ rounds are needed for all processes to decide. This is a worst-case complexity bound for synchronous set-agreement.

Dolev, Reischuk and Strong were the first to consider early-stopping protocols (best-case complexity). In particular they studied in [5] the Byzantine agreement problem, for which they gave the first early-stopping protocol. Keidar and Rajsbaum in [16], and Charron-Bost and Schiper in [2],

considered early-deciding consensus and proved that $f + 2$ rounds are needed in the synchronous message-passing system for all processes to decide, in runs with at most f process crashes.

Early-deciding k -set-agreement was first studied by Gafni et al. in [11]. An early-deciding k -set-agreement protocol was proposed, together with a matching lower bound. As we discuss now, the bound we prove in this paper and that of [11] are in a precise sense incomparable. On the one hand, the bound was given in [11] for the case where the number n of processes is unbounded. It is in this sense a *weaker* result than the one we prove here. Indeed, that lower bound does not generalize the results on consensus where $n + 1$ (the total number of processes), and t (the number of failures that may occur in any run) are fixed, nor on the (worst-case) complexity of k -set-agreement. In the present paper, we assume that n and t are fixed and known, and we present a *global* decision lower bound result that thus generalizes the results on the time complexity of early-deciding consensus and the worst-case time complexity of k -set-agreement [4, 14, 16, 2]. All considered global decision with a fixed number of processes.

On the other hand, the bound of [11] states that *no* single process may decide within $\lfloor f/k \rfloor + 1$ rounds. In this sense, the result of [11] characterizes a *local decision* [7] bound and is in this sense *stronger* than the bound of this paper. Coming up with a bound on local decisions and a bounded number of processes is an open question that is out of the scope of this paper.

3 Overview of the Proof

Our lower bound proof relies on some notions of algebraic topology applied to distributed computing, following in particular the work of [15]. In short, an impossibility of solving set-agreement comes down to showing that the runs, or a subset of the runs, produced by a full-information protocol (a generic protocol where processes exchange their complete local state in any round), gathered within a *protocol complex*, have a sufficiently high *connectivity*. Connectivity is an abstract notion of algebraic topology which, when used in the context of set-agreement, captures the fact that the processes are sufficiently *confused* so that they would violate set-agreement if they were to decide some value; i.e., they would decide on more than k values in at least one of the runs. Basically, 0-connectivity corresponds to the traditional graph connectivity, whereas $(k - 1)$ -connectivity means the absence of "holes" of dimension k .

Our proof proceeds by contradiction. We assume that all processes decide by the end of round $\lfloor f/k \rfloor + 1$ in any run with at most f failures, and we derive a contradiction in two steps. The first step concerns rounds 1 to $\lfloor f/k \rfloor$, whereas the second part concerns round $\lfloor f/k \rfloor + 1$. The second step builds on the result of the first part. In both steps, we show that a full information protocol $\mathcal{P}$, remains highly connected, thus preventing processes from achieving k -set-agreement.

In both steps, we only focus on a subset of all possible runs. In the first step, we gather all the runs in which at most k processes crash in any round, starting from the set of all system states where $n + 1$ processes propose different values from a value range V . The protocol complex corresponding to this subset of runs is $(k - 1)$ -connected, at the end of any round r [14]. Roughly speaking, the $(k - 1)$ -connectivity of the protocol complex at the end of round $\lfloor f/k \rfloor$ is made by those runs in which $k + 1$ processes have $k + 1$ distinct *estimate* values (potential decisions), and would thus decide on $k + 1$ distinct values if these processes had to decide at the end of round $\lfloor f/k \rfloor$.

Then, in the second step, we focus on round $\lfloor f/k \rfloor + 1$, and we extend the protocol complex obtained at round $\lfloor f/k \rfloor$ with a round in which, as before, at most k processes crash, but now every

process observes at most $k - 1$ crashes. In other words, in this additional round $\lfloor f/k \rfloor + 1$, every process that reaches the end of the round receives a message from at least one process that crashes in round $r + 1$. The intuition behind this round is to force processes to decide at the end of round $\lfloor f/k \rfloor + 1$, and then obtain the desired contradiction with the computation of the connectivity. Indeed, any process p_i that receives, in round $\lfloor f/k \rfloor + 1$, at least one message from one of the k processes that crash in round $\lfloor f/k \rfloor + 1$, decides at the end of round $\lfloor f/k \rfloor + 1$.

This is because the subset of runs that we consider is indistinguishable for any process at the end of round $\lfloor f/k \rfloor + 1$, from a run that has at most k crashes in the first $\lfloor f/k \rfloor$ rounds, and at most $k - 1$ crashes in round $\lfloor f/k \rfloor + 1$: a total of $k \lfloor f/k \rfloor + (k - 1)$ crashes. In this case, processes must decide at the end of round $\lfloor f/k \rfloor + 1$.

We finally obtain our contradiction by showing that extending the protocol complex obtained at the end of round $\lfloor f/k \rfloor$, with the round $\lfloor f/k \rfloor + 1$ described in the previous paragraph, i.e., where at most k processes crash but any process observes at most $k - 1$ crashes, preserves the $(k - 1)$ -connectivity of the protocol complex, at the end of round $\lfloor f/k \rfloor + 1$. By applying the result relating high connectivity and the impossibility of set-agreement, formalized in Theorem 3, we derive the fact that not all processes may decide at the end of round $\lfloor f/k \rfloor + 1$.

The main technical difficulty is to prove that the connectivity of the complex obtained at the end of round $\lfloor f/k \rfloor + 1$ is high-enough. The approach here is similar to that of [14] in the sense that we compute connectivity by induction, using the topological notions of *pseudosphere* and union of pseudospheres. Basically, the protocol complexes of which we compute the connectivity can be viewed as a union of n -dimensional pseudospheres which makes it possible to apply (a corollary of) the Mayer-Vietoris theorem [17]. We also use here a theorem from [12], which itself generalizes Theorem 9 and Theorem 11 of [14].

The main originality in our work is the introduction of our *early-deciding* operator, which is key to showing that the connectivity is preserved from round $\lfloor f/k \rfloor$ to round $\lfloor f/k \rfloor + 1$, i.e., even if processes see less than k failures in the last round.

4 Model

Processes. We consider a distributed system made of a set Π of $n + 1$ processes, $p_0, \dots, p_n$. Each process is an infinite state-machine. The processes communicate via message passing through reliable channels, in synchronous rounds. Every round r proceeds in three phases: (1) first any process p_i sends a message to all processes in Π ; (2) then process p_i receives all the messages that have been sent to it in round r ; (3) at last p_i performs some local run, changes its state, and starts round $r + 1$.

Failures. The processes may fail by crashing. When a process crashes, it stops executing any step from its assigned protocol. If any process p_i crashes in the course of sending its message to all the processes, a subset only of the messages that p_i sends are received. We assume that at most t out of the $n + 1$ processes may crash in any run. The identity of the processes that crash vary from one run to another and is not known in advance. We denote by $f \leq t$ the effective number of crashes that occur in any run.

Problem. In this paper, we consider the k -set-agreement problem. In this problem, any process p_i is supposed to propose a value $v_i \in V$, such that $|V| > k$ (otherwise, the problem is trivially solved), and eventually decide on a value v'_i , such that the following three conditions are satisfied:

(**Validity**) Any decided value v'_i is a value v_j proposed by some process p_j .

(**Termination**) Eventually, every correct process decides.

(**k -set-agreement**) There are at most k distinct decided values.

5 Topological Background

This section recalls some general notions and results from basic algebraic topology from [17], together with some specific ones from [14] used to prove our result.

5.1 Simplexes and complexes

It is convenient to model a global state of a system of $n + 1$ processes as an n -dimensional simplex $S^n = (s_0, \dots, s_n)$, where $s_i = \langle p_i, v_i \rangle$ defines local state v_i of process p_i [15]. We say that the vertexes $s_0, \dots, s_n$ span the simplex S^n . We say that a simplex T is a *face* of a simplex S if all vertexes of T are vertexes of S . A set of global states is modeled as a set of simplexes, closed under containment, called a *complex*.

5.2 Protocols

A *protocol* $\mathcal{P}$ is a subset of runs of our model. For any initial state represented as an n -simplex S , a *protocol complex* $\mathcal{P}(S)$ defines the set of final states reachable from them through the runs in $\mathcal{P}$. In other words, a set of vertexes $\langle p_{i_0}, v_{i_0} \rangle, \dots, \langle p_{i_n}, v_{i_n} \rangle$ span a simplex in $\mathcal{P}(S)$ if and only if (1) S defines the initial state of $p_{i_0}, \dots, p_{i_n}$, and (2) there is a run in $\mathcal{P}$ in which $p_{i_0}, \dots, p_{i_n}$ finish the protocol with states $v_{i_0}, \dots, v_{i_n}$. For a set $\{S_i\}$ of possible initial states, $\mathcal{P}(\cup_i S_i)$ is defined as $\cup_i \mathcal{P}(S_i)$. If S^m is a face of S^n , then we define $\mathcal{P}(S^m)$ to be a subcomplex of $\mathcal{P}(S^n)$ corresponding to the runs in $\mathcal{P}$ in which only processes of S^m take steps and processes of $S^n \setminus S^m$ do not take steps. For $m < n - t$, $\mathcal{P}(S^m) = \emptyset$, since in our model, there is no run in which more than t processes may fail.

For any two complexes $\mathcal{K}$ and $\mathcal{L}$, $\mathcal{P}(\mathcal{K} \cap \mathcal{L}) = \mathcal{P}(\mathcal{K}) \cap \mathcal{P}(\mathcal{L})$: any state of $\mathcal{P}(\mathcal{K} \cap \mathcal{L})$ belongs to both $\mathcal{P}(\mathcal{K})$ and $\mathcal{P}(\mathcal{L})$, any state from $\mathcal{P}(\mathcal{K}) \cap \mathcal{P}(\mathcal{L})$ defines the final states of processes originated from $\mathcal{K} \cap \mathcal{L}$ and, thus, belongs to $\mathcal{P}(\mathcal{K} \cap \mathcal{L})$.

We denote by $\mathcal{I}$ a complex corresponding to a set of possible initial configurations. Informally, a protocol $\mathcal{P}$ solves k -set-agreement for $\mathcal{I}$ if there exists a map δ that carries each vertex of $\mathcal{P}(\mathcal{I})$ to a decision value in such a way that, for any $S^m = (\langle p_{i_0}, v_{i_0} \rangle, \dots, \langle p_{i_m}, v_{i_m} \rangle) \in \mathcal{I}$ ($m \geq n - f$), we have $\delta(\mathcal{P}(S^m)) \subseteq \{v_{i_0}, \dots, v_{i_m}\}$ and $|\delta(\mathcal{P}(S^m))| \leq k$. (The formal definition of a *solvable task* is given in [15].)

Thus, in order to show that k -set-agreement is not solvable in r rounds, it is sufficient to find an r -round protocol $\mathcal{P}$ that cannot solve the problem for some $\mathcal{I}$. Such a protocol can be interpreted as a set of worst-case runs in which no decision can be taken.

5.3 Pseudospheres

To prove our lower bound, we use the notion of *pseudosphere* introduced in [14] as a convenient abstraction to describe the topological structure of a bounded number of rounds of distributed protocol in our model. To make the paper self-contained, we recall the definition of [14] here:

Definition 1. Let $S^m = (s_0, \dots, s_m)$ be a simplex and $U_0, \dots, U_m$ be a sequence of finite sets. The pseudosphere $\psi(S^m; U_0, \dots, U_m)$ is a complex defined as follows. Each vertex of $\psi(S^m; U_0, \dots, U_m)$ is a pair $\langle s_i, u_i \rangle$, where s_i is a vertex of S^m and $u_i \in U_i$. Vertexes $\langle s_{i_0}, u_{i_0} \rangle, \dots, \langle s_{i_l}, u_{i_l} \rangle$ define a simplex of $\psi(S^m; U_0, \dots, U_m)$ if and only if all s_{i_j} ($0 \leq j \leq l$) are distinct. If for all $0 \leq i \leq m$, $U_i = U$, the pseudosphere is written $\psi(S^m; U)$.

The following properties of pseudospheres follow from their definition:

1. If $U_0, \dots, U_m$ are singleton sets, then $\psi(S^m; U_0, \dots, U_m) \cong S^m$.
2. $\psi(S^m; U_0, \dots, U_m) \cap \psi(S^m; V_0, \dots, V_m) \cong \psi(S^m; U_0 \cap V_0, \dots, U_m \cap V_m)$.
3. If $U_i = \emptyset$, then $\psi(S^m; U_0, \dots, U_m) \cong \psi(S^{m-1}; U_0, \dots, \hat{U}_i, \dots, U_m)$, where circumflex means that U_i is omitted in the sequence $U_0, \dots, U_m$.

5.4 Connectivity

Computing the connectivity of a given protocol complex plays a key role in characterizing whether the corresponding protocol may solve k -set-agreement. Informally speaking, a complex is said to be k -connected if it has no holes in dimension k or less. Theorem 3 below states that a protocol complex that is $(k-1)$ -connected cannot solve k -set-agreement.

Before giving a formal definition of connectivity, we briefly recall the standard topological notions of a *disc* and of a *sphere*. We say that a complex $\mathcal{C}$ is an m -disc if $|\mathcal{C}|$ (the convex hull occupied by $\mathcal{C}$) is homeomorphic to $\{x \in \mathbb{R}^m | d(x, 0) \leq 1\}$ whereas it is an $(m-1)$ -sphere if $|\mathcal{C}|$ is homeomorphic to $\{x \in \mathbb{R}^m | d(x, 0) = 1\}$. For instance, the 2-disc is the traditional two-dimensional disc, whereas the 2-sphere is the traditional three-dimensional sphere.

We adopt the following definition of connectivity, given in [15]:

Definition 2. For $k > 0$, a complex $\mathcal{K}$ is k -connected if, for every $m \leq k$, any continuous map of the m -sphere to $\mathcal{K}$ can be extended to a continuous map of the $(m+1)$ -disk. By convention, a complex is (-1) -connected if it is non-empty, and every complex is k -connected for $k < -1$.

The following corollary to the Mayer-Vietoris theorem [17] helps define the connectivity of the result of $\mathcal{P}$ applied to a union of complexes:

Theorem 1. If $\mathcal{K}$ and $\mathcal{L}$ are k -connected complexes, and $\mathcal{K} \cap \mathcal{L}$ is $(k-1)$ -connected, then $\mathcal{K} \cup \mathcal{L}$ is k -connected.

The following theorem from [12] generalizes Theorem 9 and Theorem 11 of [14], and helps define the connectivity of a union of pseudospheres. The proof basically reuses the arguments from [14]. Later in the paper, we use Theorem 2 to compute the connectivity of a complex to which we apply our early-deciding operator.

Theorem 2. Let $\mathcal{P}$ be a protocol, S^m a simplex, and c a constant integer. Let for every face S^l of S^m , the protocol complex $\mathcal{P}(S^l)$ be $(l-c-1)$ -connected. Then for every sequence of finite sets $\{A_{0_j}\}_{j=0}^m, \dots, \{A_{l_j}\}_{j=0}^m$, such that for any $j \in [0, m]$, $\bigcap_{i=0}^l A_{i_j} \neq \emptyset$, the protocol complex

$$\mathcal{P} \left(\bigcup_{i=0}^l \psi(S^m; A_{i_0}, \dots, A_{i_m}) \right) \text{ is } (m-c-1)\text{-connected.} \quad (\text{Eq. 1})$$

Proof. Since for any sequence $V_0, \dots, V_l$ of singleton sets, $\psi(S^l; V_0, \dots, V_l) \cong S^l$, we notice that $\mathcal{P}(\psi(S^l; V_0, \dots, V_l)) \cong \mathcal{P}(S^l)$ is $(l - c - 1)$ -connected.

- (i) First, we prove that, for any m and any non-empty sets $U_0, \dots, U_m$, the protocol complex $\mathcal{P}(\psi(S^m; U_0, \dots, U_m))$ is $(m - c - 1)$ -connected. We introduce here the partial order on the sequences $U_0, \dots, U_m$: $(V_0, \dots, V_m) \prec (U_0, \dots, U_m)$ if and only if each $V_i \subseteq U_i$ and for some j , $V_j \subset U_j$. We proceed by induction on m . For $m = c$ and any sequence $U_0, \dots, U_m$, the protocol complex $\mathcal{P}(\psi(S^m; U_0, \dots, U_m))$ is non-empty and, by definition, (-1) -connected.

Now assume that the claim holds for all simplexes of dimension less than m ($m > c$). We proceed by induction on the partially-ordered sequences of sets $U_0, \dots, U_m$. For the case where $(U_0, \dots, U_m)$ are singletons, the claim follows from the theorem condition. Assume that the claim holds for all sequences smaller than $U_0, \dots, U_m$ and there is an index i , such that $U_i = v \cup V_i$, where V_i is non-empty ($v \notin V_i$). $\mathcal{P}(\psi(S^m; U_0, \dots, U_m))$ is the union of $\mathcal{K} = \mathcal{P}(\psi(S^m; U_0, \dots, V_i, \dots, U_m))$ and $\mathcal{L} = \mathcal{P}(\psi(S^m; U_0, \dots, \{v\}, \dots, U_m))$ which are both $(m - c - 1)$ -connected by the induction hypothesis. The intersection is:

$$\begin{aligned} \mathcal{K} \cap \mathcal{L} &= \mathcal{P}(\psi(S^m; U_0, \dots, V_i \cap \{v\}, \dots, U_m)) = \\ &= \mathcal{P}(\psi(S^m; U_0, \dots, \emptyset, \dots, U_m)) \cong \\ &\cong \mathcal{P}(\psi(S^{m-1}; U_0, \dots, \widehat{\emptyset}, \dots, U_m)). \end{aligned}$$

The argument of $\mathcal{P}$ in the last expression represents an $(m - 1)$ -dimensional pseudosphere which is $(m - c - 2)$ -connected by the induction hypothesis. By Theorem 1, $\mathcal{K} \cup \mathcal{L} = \mathcal{P}(\psi(S^m; U_0, \dots, U_m))$ is $(m - c - 1)$ -connected.

- (ii) Now we prove our theorem by induction on l . We show that for any $l \geq 0$ and any sequence of sets $\{A_{i_j}\}$ satisfying the condition of the theorem, Equation 1 is guaranteed. The case $l = 0$ follows directly from (i). Now assume that, for some $l > 0$,

$$\mathcal{K} = \mathcal{P} \left(\bigcup_{i=0}^{l-1} \psi(S^m; A_{i_0}, \dots, A_{i_m}) \right) \text{ is } (m - c - 1)\text{-connected.} \quad (\text{Eq. 2})$$

By (i), $\mathcal{L} = \mathcal{P}(\psi(S^m; A_{l_0}, \dots, A_{l_m}))$ is $(m - c - 1)$ -connected. The intersection is

$$\begin{aligned} \mathcal{K} \cap \mathcal{L} &= \mathcal{P} \left(\left(\bigcup_{i=0}^{l-1} \psi(S^m; A_{i_0}, \dots, A_{i_m}) \right) \cap \psi(S^m; A_{l_0}, \dots, A_{l_m}) \right) = \\ &= \mathcal{P} \left(\bigcup_{i=0}^{l-1} \psi(S^m; A_{i_0} \cap A_{l_0}, \dots, A_{i_m} \cap A_{l_m}) \right). \end{aligned}$$

By the initial assumption (Equation 2), for any j , $\bigcap_{i=0}^{l-1} (A_{i_j} \cap A_{l_j}) = \bigcap_{i=0}^l A_{i_j} \neq \emptyset$. Thus by the induction hypothesis,

$$\mathcal{K} \cap \mathcal{L} = \mathcal{P} \left(\bigcup_{i=0}^{l-1} \psi(S^m; A_{i_0} \cap A_{l_0}, \dots, A_{i_m} \cap A_{l_m}) \right) \text{ is } (m - c - 1)\text{-connected.}$$

By Theorem 1, $\mathcal{K} \cup \mathcal{L}$ is $(m - c - 1)$ -connected.

5.5 Impossibility and connectivity

The following theorem, borrowed from [14], is based on Sperner's lemma [17]: it relates the connectivity of a protocol complex derived from a pseudosphere, with the impossibility of k -set-agreement:

Theorem 3. *Let $\mathcal{P}$ be a protocol. If for every n -dimensional pseudosphere $\psi(p_0, \dots, p_n; V)$, where V is non-empty, $\mathcal{P}(\psi(p_0, \dots, p_n; V))$ is $(k-1)$ -connected, and there are more than k possible input values, then $\mathcal{P}$ cannot solve k -set agreement.*

6 The Lower bound

As we pointed out in Section 3, our lower bound proof proceeds by contradiction. We assume that there is a full information protocol $\mathcal{P}$ using which all correct processes can decide by round $\lfloor f/k \rfloor + 1$. We construct a complex of $\mathcal{P}$ that satisfies the precondition of Theorem 3: namely, for any pseudosphere $\psi(p_0, \dots, p_n; V)$, where V is non-empty, $\mathcal{P}(\psi(p_0, \dots, p_n; V))$ is $(k-1)$ -connected. Basically, the $(k-1)$ -connectivity of the protocol complex at the end of round $\lfloor f/k \rfloor + 1$ is made by those runs in which $k+1$ processes have $k+1$ distinct estimate values, and would thus decide on $k+1$ distinct values if these processes had to decide at the end of round $\lfloor f/k \rfloor + 1$. The protocol complex corresponding to the subset of runs of $\mathcal{P}$ where, in every run, at most k processes are allowed to fail, is $(k-1)$ -connected, at the end of any round r , in particular $\lfloor f/k \rfloor$: this follows from the use of the topological operator $\S$, introduced in [14]. In round $\lfloor f/k \rfloor + 1$, we extend the protocol complex with a last round in which at most k process crash, but every process observes at most $k-1$ crashes. In other words, in this last round, every process that reaches the end of the round receives a message from at least one process that crashes in the round. We show that this extension still preserves the $(k-1)$ -connectivity of the protocol complex at the end of round $r+1$. We use here a notion topological operator $\mathcal{E}$. We conclude by applying the result of Theorem 3, and derive the fact that not all processes may decide at the end of round $r+1 = \lfloor f/k \rfloor + 1$.

6.1 Single round and Multiple Round Operators

In the proof, we use the topological round operator $\S$, which generates a set of runs in a synchronous message-passing model, in which at most k processes may crash in any round. Operator $\S$ was introduced in [14]. We recall some results about $\S$ that are necessary for presenting our lower bound proof.

The protocol complex $\S^1(S^l)$ corresponds to all single-round runs of our model, starting from an initial configuration S^l , in which up to k processes can fail by crashing. We consider the case where $k \leq l$, otherwise the protocol complex is trivial. $\S^1(S^l)$ is the union of the complexes $\S_K^1(S^n)$ of single-round runs starting from S^n in which *exactly* the processes in K fail. Given a set of processes, let $S^n \setminus K$ be the face of S^n labeled with the processes *not* in K . Lemmas 1, 2 and 3 below, are Lemmas 18, 21 and 22 from [14]. The first lemma says that $\S_K^1(S^n)$ is a pseudosphere, which means that $\S^1(S^n)$ is a union of pseudospheres.

Lemma 1. $\S_K^1(S^n) \cong \psi(S^n \setminus K; 2^K)$.

Lemma 2. *If $n \geq 2k$ and for all l , then $\S^1(S^l)$ is $(l - (n - k) - 1)$ -connected.*

The connectivity result over a single round is now used to compute the connectivity over runs spanning multiple rounds.

Lemma 3. *If $n \geq rk + k$, and ξ^r is an r -round, $(n + 1)$ -process protocol with degree k , then $\xi^r(S^l)$ is $(l - (n - k) - 1)$ -connected for any $0 \leq m \leq n$.*

6.2 Early-deciding Operator

So far, we have characterized runs in which at most k processes may crash in a round, without being interested in how many of these crashes other processes actually see. To derive our lower bound, we focus on runs where processes see less than k failures in the last round.

We introduce for that purpose a new round operator, $\mathcal{E}^1(S^n)$, which generates all single-round runs from the initial simplex S^n (obtained following the construction of the previous paragraph), in which at most k processes crash, and any process that does not crash misses at most $k - 1$ messages from crashed processes (in other words, any process that does not crash receives a message from at least one crashed process). $\mathcal{E}^1(S^n)$ is the complex of one-round runs of an $(n + 1)$ -process protocol with input simplex S^n in which at most k processes crash and every non-crashed process misses at most $k - 1$ messages. It is the union of complexes $\mathcal{E}_K^1(S^n)$ of one-round runs starting from S^n in which *exactly* the processes in K fail and any process that does not crash misses at most $k - 1$ messages.

We first show that $\mathcal{E}_K^1(S^n)$ is a pseudo-sphere, which means that $\mathcal{E}^1(S^n)$ is a union of pseudo-spheres. In the following lemma, 2_k^K denotes the set of all subsets of K of size at most $k - 1$.

Lemma 4. $\mathcal{E}_K^1(S^n) \cong \psi(S^n \setminus K; 2_k^K)$.

Proof. The processes that do not crash are those in $S^n \setminus K$. Each process that does not crash may be labeled with all messages from processes that do not crash (processes in $S^n \setminus K$), plus any combination of size at most $k - 1$ of the messages from processes that crash, represented by the subsets in 2_k^K . Hence, for any $i \in \text{ids}(S^n \setminus K)$, then $\text{label}(i)$ concatenates $S^n \setminus K$, plus a particular subset of K .

To compute the union of all pseudo-spheres, we characterize their intersection and apply Theorem 2. We order the sets K in the lexicographic order of process ids, starting from the empty set, singleton sets, 2-process sets, etc. Let $K_0, \dots, K_l$ be the ordered sequence of process ids less than or equal to K_l , listed in lexicographic order.

Lemma 5.

$$\bigcup_{i=0}^{l-1} \mathcal{E}_{K_i}^1(S^n) \cap \mathcal{E}_{K_l}^1(S^n) \cong \bigcup_{j \in K_l} \psi(S^n \setminus K_l; 2_k^{K_l - \{j\}}).$$

Proof. The proof proceeds in two parts, first for the $\subseteq$ inclusion, then for the $\supseteq$ inclusion.

For the $\subseteq$ inclusion, we show that any $\mathcal{E}_{K_i}^1(S^n) \cap \mathcal{E}_{K_l}^1(S^n)$ is included in $\psi(S^n \setminus K_l; 2_k^{K_l - \{j\}})$ for some j in K_l :

$$\mathcal{E}_{K_i}^1(S^n) \cap \mathcal{E}_{K_l}^1(S^n) \cong \psi(S^n \setminus K_i; 2_k^{K_i}) \cap \psi(S^n \setminus K_l; 2_k^{K_l}) \quad (1)$$

$$\cong \psi((S^n \setminus K_i) \cap (S^n \setminus K_l); (2_k^{K_i}) \cap (2_k^{K_l})) \quad (2)$$

$$\cong \psi(S^n \setminus (K_i \cup K_l); 2_k^{K_i \cap K_l}) \quad (3)$$

$$\subseteq \psi(S^n \setminus K_l; 2_k^{K_l - \{j\}}). \quad (4)$$

Equation 1 follows from the definition. Equations 2 and 3 follow from basic properties of pseudo-spheres. Equation 4 follows from the following observation: since K_i precedes K_l in the sequence and $K_i \neq K_k$, then there exists at least one process $p_j \in K_l$ and $p_j \notin K_i$. Thus we have (i) $S^n \setminus (K_i \cup K_l) \subseteq S^n \setminus K_l$ and (ii) $2_k^{K_j \cap K_l} \subseteq 2_k^{K_l - \{j\}}$.

For the $\supseteq$ inclusion, we observe that for any process p_j , each set $K_l - \{j\}$ precedes K_l in the sequence. Hence for any process p_j , we have:

$$\mathcal{E}_{K_l - \{j\}}^1(S^n) \cap \mathcal{E}_{K_l}^1(S^n) \cong \psi(S^n \setminus K_l - \{j\}; 2_k^{K_l - \{j\}}) \cap \psi(S^n \setminus K_l; 2_k^{K_l}) \quad (5)$$

$$\cong \psi((S^n \setminus K_l - \{j\}) \cap (S^n \setminus K_l); 2_k^{K_l - \{j\}} \cap 2_k^{K_l}) \quad (6)$$

$$\cong \psi(S^n \setminus K_l; 2_k^{K_l - \{j\}}). \quad (7)$$

Equation 5 follows from the definition of the early-deciding operator. Equation 6 follows from basic properties of pseudo-spheres, presented in Section 5.3. Equation 7 follows from the fact that $K_l - \{j\} \cap K_l = K_l - \{j\}$.

We denote $\mathcal{E}^1(S^n)$ the protocol complex for a one-round synchronous $(n+1)$ -process protocol in which no more than k processes crash, and every process that does not crash misses at most $k-1$ messages from processes that crash.

Lemma 6. *For $n \geq 2k$, then $\mathcal{E}^1(S^n)$ is $(k - (n - m) - 1)$ -connected.*

Proof. We have three cases: (i) $m = n$, (ii) $n - k \leq m < n$, and (iii) $m < n - k$.

For case (i), let $K_0, \dots, K_l$ be the sequence of sets of k processes that crash in the first round ordered lexicographically, that are less or equal to K_l . Let K_l be the maximal set of k processes, i.e., $K_l = \{p_{n-k+1}, \dots, p_n\}$. Then we have:

$$\mathcal{E}^1(S^n) = \bigcup_{i=0}^l \mathcal{E}_{K_i}^1(S^n).$$

We inductively show on l that $\mathcal{E}^1(S^n)$ is $(k-1)$ -connected. First, observe that for $l = 0$, then $\mathcal{E}_{K_0}^1(S^n) \cong \psi(S^n; \{\emptyset\}) \cong S^n$ which is $(n-1)$ -connected. As $n \geq 2k$, $n-1 \geq k-1$, and $\mathcal{E}_{K_0}^1(S^n)$ is $(k-1)$ -connected.

For the induction hypothesis, assume that:

$$\mathcal{K} = \bigcup_{i=0}^{l-1} \mathcal{E}_{K_i}^1(S^n)$$

is $(k-1)$ -connected. Let the complex $\mathcal{L}$ be:

$$\mathcal{L} = \mathcal{E}_{K_l}^1(S^n) = \psi(S^n \setminus K_l; 2_k^{K_l}).$$

As $\dim(S^n \setminus K_l) \geq n - k$, $\mathcal{L}$ is $(n - k - 1)$ -connected by Corollary 10 of [14]. As $n \geq 2k$, $\mathcal{L}$ is $(k-1)$ -connected.

We want to show that $\mathcal{K} \cup \mathcal{L}$ is $(k-1)$ -connected, and for that, we need to show that $\mathcal{K} \cap \mathcal{L}$ is at least $(k-2)$ -connected. We have:

$$\mathcal{K} \cap \mathcal{L} = \bigcup_{i=0}^{l-1} \mathcal{E}_{K_i}^1(S^n) \cap \mathcal{E}_{K_l}^1(S^n) \quad (8)$$

$$= \bigcup_{j \in K_l} \psi(S^n \setminus K_l; 2_k^{K_l - \{j\}}). \quad (9)$$

Equation 8 follows from the definition of $\mathcal{K}$ and $\mathcal{L}$. Equation 9 follows from Lemma 5.

Now let $A_i = 2_k^{K_l - \{i\}}$. We know that:

$$\bigcap_{i \in K_l} A_i = \{\emptyset\} \neq \emptyset.$$

and $S^n \setminus K_l$ has dimension at least $n - k$, so Corollary 12 of [14] implies that $\mathcal{K} \cap \mathcal{L}$ is $(n - k - 1)$ -connected. As $n \geq 2k$, $\mathcal{K} \cap \mathcal{L}$ is $(k - 1)$ -connected.

For case (ii), $n - k \leq m < n$. Recall that $\mathcal{E}^1(S^m)$ is the set of runs in which only processes in S^m take steps. As a result, the corresponding protocol complex is equivalent to the complex made of runs of $m + 1$ processes, out of which $k - n + m$ may be faulty. If we now substitute m for n , and $k - n + m$ for k , $\mathcal{E}^1(S^m)$ is $(k - (n - m) - 1)$ -connected.

For case (iii), $m < n - k$, $k - (n - m) - 1 < -1$ and thus, $\mathcal{E}^1(S^m)$ is empty.

Combining our one-round operator $\mathcal{E}$ and the round operator $\mathcal{S}$ corresponding to the set of runs in which at most k processes crash in a round, we obtain the following:

Lemma 7. *If $n \geq (r + 1)k + k$, $\mathcal{E}^1(\mathcal{S}^r(S^m))$ is an $(r + 1)$ -round, $(n + 1)$ -process protocol with degree k , then $\mathcal{E}^1(\mathcal{S}^r(S^m))$ is $(k - (n - m) - 1)$ -connected, for any $0 \leq m \leq n$.*

Proof. We prove the theorem by induction on r . For the base case $r = 0$, $n \geq 2k$ and thus in this case, Lemma 6 proves that $\mathcal{E}^1(S^m)$ is $(k - (n - m) - 1)$ -connected. For the induction hypothesis, assume the claim holds for $r - 1$.

We first consider the case where $m = n$. We denote by $K_0, \dots, K_l$ the sequence of all sets of processes less than or equal to K_l , listed in lexicographic order. The set of r -round runs in which *exactly* the processes in K_i fail in the first round can be written as $\mathcal{E}_i^{r-1}(\mathcal{E}_{K_i}^1(S^n))$, where $\mathcal{E}_i^{r-1}$ is the complex of for an $(r - 1)$ -round, $(t - |K_i|)$ -faulty, $(n + 1 - |K_i|)$ -process full-information protocol. The $\mathcal{E}_i^{r-1}$ are considered as different protocols because the $\mathcal{E}_{K_i}^1(S^n)$ have varying dimensions. We inductively show that if $|K_l| \leq k$, then:

$$\bigcup_{i=0}^l \mathcal{E}_i^{r-1}(\mathcal{E}_{K_i}^1(S^n)) \text{ is } (k - 1)\text{-connected.}$$

The claim then follows when K_l is the maximal set of size k .

For the base case, we have $l = 0$, $K_0 = \emptyset$, and thus $\mathcal{E}_0^1(S^n)$ is $\psi(S^n; 2^\emptyset) \cong S^n$, and $\mathcal{E}^1(\mathcal{E}_0^{r-1}(S^n))$ is $(k - 1)$ -connected by the induction hypothesis on r .

For the induction step on l , assume that:

$$\mathcal{K} = \bigcup_{i=0}^{l-1} \mathcal{E}_i^{r-1}(\mathcal{E}_{K_i}^1(S^n)) \text{ is } (k - 1)\text{-connected.}$$

By Lemma 1, we have:

$$\mathcal{L} = \mathcal{E}^1(\mathfrak{S}_l^{r-1}(\mathfrak{S}_{K_l}^1(S^n))) = \mathcal{E}^1(\mathfrak{S}_l^{r-1}(\psi(S^n \setminus K_l; 2^{K_l}))).$$

We recall that $\mathcal{E}^1(\mathfrak{S}_l^{r-1})$ is a rk -faulty, $(n+1-|K_l|)$ -process, r -round protocol, where $n+1-|K_l| \geq rk$, so by the induction hypothesis, for each simplex $S^d \in \mathfrak{S}_{K_l}^1(S^n) = \psi(S^n \setminus K_l; 2^{K_l})$, $\mathcal{E}^1(\mathfrak{S}_l^{r-1}(S^d))$ is $(k - (n - |K_l| - d) - 1)$ -connected. By Theorem 2, $\mathcal{E}^1(\mathfrak{S}_l^{r-1}(\psi(S^n \setminus K_l; 2^{K_l}))) = \mathcal{E}^1(\mathfrak{S}_l^{r-1}(\mathfrak{S}_{K_l}^1(S^n))) = \mathcal{L}$ is $(k - 1)$ -connected.

We claim the following property:

Claim.

$$\begin{aligned} \mathcal{K} \cap \mathcal{L} &= \bigcup_{i=0}^{l-1} \mathcal{E}^1(\mathfrak{S}_i^{r-1}(\psi(S^n \setminus K_i; 2^{K_i}))) \cap \mathcal{E}^1(\mathfrak{S}_l^{r-1}(\psi(S^n \setminus K_l; 2^{K_l}))) \\ &= \mathcal{E}^1(\tilde{\mathfrak{S}}_l^{r-1} \left(\bigcup_{i \in K_l} \psi(S^n \setminus K_l; 2^{K_l - \{i\}}) \right)), \end{aligned}$$

where $\tilde{\mathfrak{S}}_l^{r-1}$ is a protocol identical to $\mathfrak{S}_l^{r-1}$ except that $\tilde{\mathfrak{S}}_l^{r-1}$ fails at most $k - 1$ processes in its first round.

Proof. For the $\subseteq$ inclusion, in the exact same manner as we have seen in the proof of Lemma 5 and, for each i , there is some $j \in K_l$ such that:

$$\psi(S^n \setminus K_i \cap S^n \setminus K_l; 2^{K_i \cap K_l}) \subseteq \psi(S^n \setminus K_l; 2^{K_l - \{j\}}).$$

We still need to show how $\mathcal{E}^1(\mathfrak{S}_i^{r-1})$ and $\mathcal{E}^1(\mathfrak{S}_l^{r-1})$ intersect. Because p_j has already failed in $\mathcal{E}^1(\mathfrak{S}_l^{r-1})$, the only runs $\mathcal{E}^1(\mathfrak{S}_i^{r-1})$ that are also present in $\mathcal{E}^1(\mathfrak{S}_l^{r-1})$ are ones in which p_j fails without sending any messages to non-faulty processes. But then $\mathcal{E}^1(\mathfrak{S}_i^{r-1})$, and therefore $\mathcal{E}^1(\mathfrak{S}_l^{r-1})$, can fail at most $k - 1$ processes that do send messages to non-faulty processes. Any such run is also a run of $\mathcal{E}^1(\tilde{\mathfrak{S}}_l^{r-1})$.

For the reverse inclusion $\supseteq$, we have seen in Lemma 5 that for each $j \in K_l$:

$$\mathcal{E}_{K_l - \{j\}}^1(S^n) \cap \mathcal{E}_{K_l}^1(S^n) \cong \psi(S^n \setminus K_l; 2^{K_l - \{j\}}).$$

It turns out that the same argument also holds for the case:

$$\mathfrak{S}_{K_l - \{j\}}^1(S^n) \cap \mathfrak{S}_{K_l}^1(S^n) \cong \psi(S^n \setminus K_l; 2^{K_l - \{j\}}).$$

The set of runs in which the two protocols overlap are exactly those runs in which $\mathcal{E}^1(\mathfrak{S}_i^{r-1})$ immediately fails p_j , and in which $\mathcal{E}^1(\mathfrak{S}_l^{r-1})$ fails no more than $k - 1$ processes. These runs comprise $\mathcal{E}^1(\tilde{\mathfrak{S}}_l^{r-1})$.

While $\mathfrak{S}_l^{r-1}$ has degree k , $\tilde{\mathfrak{S}}_l^{r-1}$ has degree $k - 1$. By the induction hypothesis on r , for any simplex S^{n-k} , $\tilde{\mathfrak{S}}_l^{r-1}(S^{n-k})$ is $(k - 2)$ -connected. Let $A_i = 2^{K_l - \{i\}}$, for $i \in K_l$. As $\bigcap_{i \in K_l} A_i = \{\emptyset\} \neq \emptyset$, $\mathcal{K} \cap \mathcal{L}$ is $(k - 2)$ -connected by Claim 6.2 and Theorem 2. The claim now follows from Theorem 1.

If $n > m \geq n - k$, $\mathcal{E}^1(\mathfrak{S}^r(S^m))$ is equivalent to an m -process protocol in which at most $k - (n - m)$ processes fail in the first round, and k thereafter. This protocol has degree $k - (n - m)$, so $\mathcal{E}^1(\mathfrak{S}^r(S^m))$ is $(k - (n - m) - 1)$ -connected.

When $m < n - k$, $k - (n - m) - 1 < -1$ and $\mathcal{E}^1(\mathcal{S}^r(S^m))$ is empty, so the condition holds vacuously.

Theorem 4. *If $n \geq k \lfloor t/k \rfloor + k$, then in any solution to k -set-agreement, not all processes may decide earlier than within round $\lfloor f/k \rfloor + 2$ in any run with at most f failures, for $0 \leq \lfloor f/k \rfloor \leq \lfloor t/k \rfloor - 1$.*

Proof. Consider the protocol complex $\mathcal{E}^1(\mathcal{S}^{\lfloor f/k \rfloor}(S^m))$. We have $k(\lfloor f/k \rfloor + 1) + 1 \leq k \lfloor t/k \rfloor + k \leq n$, thus Lemma 7 applies. Hence $\mathcal{E}^1(\mathcal{S}^{\lfloor f/k \rfloor}(S^m))$ is $(k - (n - m) - 1)$ -connected for any f such that $\lfloor f/k \rfloor \leq \lfloor t/k \rfloor - 1$, and $0 \leq m \leq n$. The result now holds immediately from Theorem 3.

7 Concluding Remark

This paper establishes a lower bound on the time complexity of early-deciding set-agreement in a synchronous model of distributed computation. This lower bound also holds for synchronous runs of an eventually synchronous model [8] but we conjecture a larger lower bound for such runs. Determining such a bound, which would generalize the result of [6], is an intriguing open problem.

As we discussed in the related work section, although, at first glance, the local decision lower bound presented in [11] seems to imply a global decision on k -set-agreement, the model in which early-deciding k -set-agreement was investigated in [11] relies on the fact that the number of processes is not bounded. In fact, the proof technique we used here is fundamentally different from [11]: in [11], the proof is based on a pure algorithmic reduction whereas we use here a topological approach. Unifying these results would mean establishing a local decision lower bound assuming a bounded number of processes. This, we believe, is an open challenging question that might require different topological tools to reason about on-going runs.

References

1. E. Borowsky and E. Gafni. Generalized FLP impossibility result for t -resilient asynchronous computation. In *Proceedings of the 25th ACM Symposium on the Theory of Computing (STOC'93)*, pages 91–100. ACM Press, 1993.
2. B. Charron-Bost and A. Schiper. Uniform consensus is harder than consensus. *Journal of Algorithms*, 51(1):15–37, 2004.
3. S. Chaudhuri. More choices allow more faults: set consensus problems in totally asynchronous systems. *Information and Computation*, 105(1):132–158, July 1993.
4. S. Chaudhuri, M. Herlihy, N. Lynch, and M. Tuttle. Tight bounds for k -set agreement. *Journal of the ACM*, 47(5):912–943, 2000.
5. D. Dolev, R. Reischuk, and H. R. Strong. Early stopping in Byzantine agreement. *Journal of the ACM*, 37(4):720–741, 1990.
6. P. Dutta and R. Guerraoui. The inherent price of indulgence. *Distributed Computing*, 18(1):85–98, 2005.
7. P. Dutta, R. Guerraoui, and B. Pochon. Tight lower bounds on early local decisions in uniform consensus. In *Proceedings of the 17th International Symposium on Distributed Computing (DISC'03)*, Lecture Notes in Computer Science, pages 264–278. Springer-Verlag, October 2003.
8. C. Dwork, N. Lynch, and L. Stockmeyer. Consensus in the presence of partial synchrony. *Journal of ACM*, 35(2):288–323, 1988.
9. M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.
10. E. Gafni. Round-by-round fault detector — unifying synchrony and asynchrony. In *Proceedings of the 17th ACM Symposium on Principles of Distributed Computing (PODC'98)*, 1998.

11. E. Gafni, R. Guerraoui, and B. Pochon. From a static impossibility to an adaptive lower bound: the complexity of early deciding set agreement. In *Proceedings of the 37th ACM Symposium on Theory of Computing (STOC'05)*, May 2005.
12. R. Guerraoui, P. Kouznetsov, and B. Pochon. A note on set agreement with omission failures. *Electronic Notes in Theoretical Computing Science*, 81, 2003.
13. M. Herlihy and S. Rajsbaum. Algebraic spans. In *Proceedings of the 14th ACM Symposium on Principles of Distributed Computing (PODC'95)*, pages 90–99, New York, NY, USA, 1995. ACM Press.
14. M. Herlihy, S. Rajsbaum, and M. Tuttle. Unifying synchronous and asynchronous message-passing models. In *Proceedings of the 17th ACM Symposium on Principles of Distributed Computing (PODC'98)*, pages 133–142, 1998.
15. M. Herlihy and N. Shavit. The topological structure of asynchronous computability. *Journal of the ACM*, 46(6):858–923, 1999.
16. I. Keidar and S. Rajsbaum. On the cost of fault-tolerant consensus when there are no faults – a tutorial. *SIGACT News, Distributed Computing Column*, 32(2):45–63, 2001.
17. J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Reading MA, 1984.
18. M. Saks and F. Zaharoglou. Wait-free k -set agreement is impossible: The topology of public knowledge. *SIAM Journal on Computing*, 29(5):1449–1483, March 2000. A preliminary version appeared in the Proceedings of the 25th ACM Symposium on the Theory of Computing (STOC'93).