A New Synchronous Lower Bound for Set Agreement

Maurice Herlihy[1], Sergio Rajsbaum[2], and Mark Tuttle[2]

[1] Brown University, Computer Science Department, Providence, RI 02912; mph@cs.brown.edu.

[2] Compaq, Cambridge Research Lab, One Cambridge Center, Cambridge, MA 02142; Sergio Rajsbaum is on leave from U.N.A.M., Mexico City; sergio.rajsbaum@compaq.com, mark.tuttle@compaq.com.

Abstract. We have a new proof of the lower bound that k-set agreement requires $\lfloor f/k \rfloor + 1$ rounds in a synchronous, message-passing model with f crash failures. The proof involves constructing the set of reachable states, proving that these states are highly connected, and then appealing to a well-known topological result that high connectivity implies that set agreement is impossible. We construct the set of reachable states in an iterative fashion using a round operator that we define, and our proof of connectivity is an inductive proof based on this iterative construction and using simple properties of the round operator. This is the shortest and simplest proof of this lower bound we have seen.

1 Introduction

The consensus problem [20] has received a great deal of attention. In this problem, $n+1$ processors begin with input values, and all must agree on one of these values as their output value. Fischer, Lynch, and Paterson [11] surprised the world by showing that solving consensus is impossible in an asynchronous system if one processor is allowed to fail. This leads one to wonder if there is any way to weaken consensus to obtain a problem that can be solved in the presence of $k-1$ failures but not in the presence of k failures. Chaudhuri [6] defined the k-set agreement problem and conjectured that this was one such problem, and a trio of papers [4, 16, 21] proved that she was right. The k-set agreement problem is like consensus, but we relax the requirement that processors agree: the set of output values chosen by the processors may contain as many as k distinct values, and not just 1. Consensus and set agreement are just as interesting in synchronous models as they are in asynchronous models. In synchronous models, it is well-known that consensus requires $f+1$ rounds of communication if f processors can crash [10, 8, 9], and that k-set agreement requires $\lfloor f/k \rfloor + 1$ rounds [7]. These lower bounds agree when $k = 1$ since consensus is just 1-set agreement. In this paper, we give a new proof of the $\lfloor f/k \rfloor + 1$ lower bound for set agreement in the synchronous message-passing model with crash failures.

All known proofs for the set agreement lower bound depend — either explicitly or implicitly — on a deep connection between computation and topology.

These proofs essentially consider the simplicial complex representing all possible reachable states of a set agreement protocol, and then argue about the connectivity of this complex. These lower bounds for set agreement follow from the observation that set agreement cannot be solved if the complex of reachable states is sufficiently highly-connected. This connection between connectivity and set agreement has been established both in a generic way [14] and in ways specialized to particular models of computation [2, 4, 7, 13–15, 21]. Once the connection has been established, however, the problem reduces to reasoning about the connectivity of a protocol's reachable complex.

Most of the prior work employing topological arguments has focused on the *asynchronous* model of computation, in which processors can run at arbitrary speeds, and fail undetectably. Reasoning about connectivity in the asynchronous model is simplified by the fact that the connectivity of the reachable complex remains unchanged over time. Moreover, the extreme flexibility of the processor failure model facilitates the use of invariance arguments to prove connectivity. In the *synchronous* model that we consider here, analyzing connectivity is significantly more complicated. The difficulty arises because the connectivity of the reachable complex changes from round to round, so the relatively simple invariance arguments used in the asynchronous model cannot possibly work here.

The primary contribution of this work is a new, substantially simpler proof of how the connectivity of the synchronous complex evolves over time. Our proof depends on two key insights:

1. The notion of a *round operator* that maps a global state to the set of global states reachable from this state by one round of computation, an operator satisfying a few simple algebraic properties.
2. The notion of an *absorbing poset* organizing the set of global states into a partial order, from which the connectivity proof follows easily using the round operator's algebraic properties.

We believe this new proof has several novel and elegant features. First, we are able to isolate a small set of elementary combinatorial properties of the round operator that suffice to establish the connection with classical topology. Second, these properties require only local reasoning about how the computation evolves from one round to the next. Finally, most connectivity arguments can be difficult to follow because they mix semantic, combinatorial, and topological arguments, but those arguments are cleanly separated here: The definition of the round operator captures the semantics of the synchronous model, the reasoning about the round operator is purely combinatorial, and the lower bound is completed with a "black box" application of well-known topological results without any need to make additional topological arguments.

In the next section, we give an overview of our proof strategy and discuss its relationship to other proofs appearing in the literature. In the main body of the paper, we sketch the proof itself. The full proof in the full paper fills just over a dozen pages, making it the shortest self-contained proof of this lower bound that we have seen.

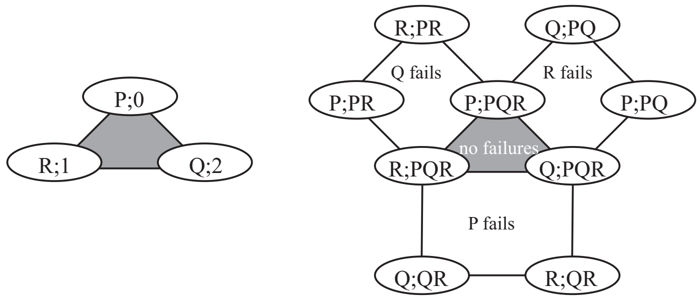

Fig. 1. A global state S and the set $\mathcal{R}_1(S)$ of global states after one round from S.

2 Overview

We assume a standard synchronous message-passing model with crash failures [3, 17]. The system has $n+1$ processors, and at most f of them can crash in any given execution. Each processor begins in an initial state consisting of its input value, and computation proceeds in a sequence of rounds. In each round, each processor sends messages to other processors, receives messages sent to it by the other processors in that round, performs some internal computation, and changes state. We assume that processors are following a full-information protocol, which means that each processor sends its entire local state to every processor in every round. This is a standard assumption to make when proving lower bounds. A processor can fail by crashing in the middle of a round, in which case it sends its state only to a subset of the processors in that round. Once a processor crashes, it never sends another message after that.

We represent the local state of a processor with a vertex labeled with that processor's id and its local state. We represent a global state as a set of labeled vertexes, labeled with distinct processors, representing the local state of each processor in that global state. In topology, a *simplex* is a set of vertexes, and a *complex* is a set of simplexes that is closed under containment. Applications of topology to distributed computing often assume that these vertexes are points in space and that the simplex is the convex hull of these points in order to be able to use standard topology results. As you read this paper, you might find it helpful to think of simplexes in this way, but in the purely combinatorial work done in this paper, a simplex is just a set of vertexes.

As an example, consider the simplex and complex illustrated in Figure 1. On the left side, we see a simplex representing an initial global state in which processor P, Q, and R start with input values 0, 2, and 1. Each vertex is labeled with a processor's id and its local state (which is just its input value in this case). On the right we see a complex representing the set of states that arise after one round of computation from this initial state if one processor is allowed to crash. The labeling of the vertexes is represented schematically by a processor id such as P and a string of processor ids such as PQ. The string PQ is intended to represent the fact that P heard from processors P and Q during the round but not from R, since R failed that round. (We are omitting input values on the right for notational simplicity.) The simplexes that represent states after one round are the 2-dimensional triangle in the center and the 1-dimensional edges that radiate from the triangle (including the edges of the triangle itself). The central triangle represents the state after a round in which no processor fails. Each edge represents a state after one processor failed. For example, the edge with vertexes labeled $P; PQR$ and $Q; PQ$ represent the global state after a round in which R fails by sending a message to P and not sending to Q: P heard from all three processors, but Q did not hear from R.

What we do in this paper is define round operators like the round operator $\mathcal{R}_1$ that maps the simplex S on the left of Figure 1 to the complex $\mathcal{R}_1(S)$ on the right, and then argue about the connectivity of $\mathcal{R}_1(S)$. Informally, connectivity in dimension 0 is just ordinary graph connectivity, and connectivity in higher dimensions means that there are no "holes" of that dimension in the complex. When we reason about connectivity, we often talk about the connectivity of a simplex S when we really mean the connectivity of the induced complex consisting of S and all of its faces. For example, both of the complexes in Figure 1 are 0-connected since they are connected in the graph theoretic sense. In fact, the complex on the left is also 1-connected, but the complex on the right is not since there are "holes" formed by the three cycles of 1-dimensional edges. The fundamental connection between k-set agreement and connectivity is that k-set agreement cannot be solved after r rounds of computation if the complex of states reachable after r rounds of computation is $(k-1)$-connected. In the remainder of this overview, we sketch how we define a round operator, and how we reason about the connectivity of the complex of reachable states.

2.1 Round operators

In the synchronous model, we can represent a round of computation with a round operator $\mathcal{R}_\ell$ that maps the state S at the start of a round to the set $\mathcal{R}_\ell(S)$ of all possible states at the end of a round in which at most ℓ processors fail. Suppose F is the set of processors that fail in a round, and consider the local state of a processor p at the end of that round. The full-information protocol has each processor send its local state to p, so p receives the local state of each processor, with the possible exception of some processors in F that fail before sending to p. Since each processor q sending to p sends its local state, and since this local state labels q's vertex in S, we can view p's local state at the end of the round as the face of S containing the local states p received from processors like q. If we define S/F to be the face of S obtained by deleting the vertexes of S labeled with processors in F, then p receives at least the local states labeling S/F, so p's local state after the round of computation can be represented by some face of S containing S/F.

This intuition leads us to define the round operator $\mathcal{R}_\ell$ as follows. For each set F of at most ℓ processors labeling a state S, define $\mathcal{R}_F(S)$ to be the set of simplexes obtained by labeling each vertex of S/F with some face of S containing S/F. This is the set of possible states after a round of computation from S in which F is the set of processors that fail, since the processors labeling S/F are the processors that are still alive at the end of the round, and since they each hear from some set of processors that contains the processors labeling S/F. The round operator $\mathcal{R}_\ell(S)$ is defined to be the union of all $\mathcal{R}_F(S)$ such that F is a set of at most ℓ processors labeling S.

To illustrate this informal definition, consider the complex $\mathcal{R}_1(S)$ of global states on the right of Figure 1. This complex is the union of four rather degenerate pseudospheres, which are complexes defined in Section 5.1 that are topologically similar to a sphere. The first pseudosphere is the central triangle. This is the pseudosphere $\mathcal{R}_\emptyset(S)$, where each processor hears from all other processors, so each processor's local state at the end of the round is the complete face $\{P, Q, R\}$ of S. The other three pseudospheres are the cycles hanging off the central triangle. These are the pseudospheres of the form $\mathcal{R}_{\{P\}}(S)$, where each processor hears from all processors with the possible exception of P, so each processor's local state at the end of the round is either the whole simplex $S = \{P, Q, R\}$ or the face $S/\{P\} = \{Q, R\}$, depending on whether the processor did or did not hear from P.

If $\mathcal{R}_\ell(S)$ is the set of possible states after one round of computation, then $\mathcal{R}_\ell^r(S) = \mathcal{R}_\ell\mathcal{R}_\ell^{r-1}(S)$ is the set of possible states after r rounds of computation. The goal of this paper is to prove that $\mathcal{R}_\ell^r(S)$ is highly-connected.

2.2 Absorbing posets

To illustrate the challenge of proving that $\mathcal{R}_\ell^r(S)$ is connected, let us assume that $\mathcal{R}_\ell(S)$ is ℓ-connected for every S and ℓ, and let us prove that $\mathcal{R}_\ell\mathcal{R}_\ell(S)$ is ℓ-connected. If $\mathcal{R}_\ell(S) = \{S_1, \ldots, S_k\}$ is the set of states after one round, then

$$\begin{aligned}\mathcal{R}_\ell\mathcal{R}_\ell(S) &= \mathcal{R}_\ell(S_1 \cup S_2 \cup \cdots \cup S_k)\\ &= \mathcal{R}_\ell(S_1) \cup \mathcal{R}_\ell(S_2) \cup \cdots \cup \mathcal{R}_\ell(S_k)\end{aligned}$$

is the set of states after two rounds. We know that the $\mathcal{R}_\ell(S_i)$ are ℓ-connected by assumption, but we need to prove that their union is ℓ-connected.

Proving that a union of complexes is connected is made easier by the Mayer-Vietoris theorem, which says that $A \cup B$ is c-connected if A and B are c-connected and $A \cap B$ is $(c-1)$-connected. This suggests that we proceed by induction on i to prove that

$$\mathcal{R}_\ell(S_1) \cup \mathcal{R}_\ell(S_2) \cup \cdots \cup \mathcal{R}_\ell(S_i)$$

is connected for $i = 1, \ldots, k$. We know that

$$\mathcal{R}_\ell(S_1) \cup \mathcal{R}_\ell(S_2) \cup \cdots \cup \mathcal{R}_\ell(S_{i-i}) \text{ and } \mathcal{R}_\ell(S_i)$$

are both ℓ-connected by hypothesis and assumption, so all we need to do is prove that their intersection

$$\begin{aligned}&[\mathcal{R}_\ell(S_1) \cup \mathcal{R}_\ell(S_2) \cup \cdots \cup \mathcal{R}_\ell(S_{i-1})] \cap \mathcal{R}_\ell(S_i)\\ &\quad= [\mathcal{R}_\ell(S_1) \cap \mathcal{R}_\ell(S_i)] \cup [\mathcal{R}_\ell(S_2) \cap \mathcal{R}_\ell(S_i)] \cup \cdots \cup [\mathcal{R}_\ell(S_{i-1}) \cap \mathcal{R}_\ell(S_i)]\end{aligned}$$

is $(\ell-1)$-connected. This union suggests another Mayer-Vietoris argument, but what do we know about the connectivity of the $\mathcal{R}_\ell(S_j) \cap \mathcal{R}_\ell(S_i)$?

One of the elegant properties of the round operator is that

$$\mathcal{R}_\ell(S_j) \cap \mathcal{R}_\ell(S_i) = \mathcal{R}_{\ell-c}(S_j \cap S_i)$$

where c is the number of vertices in S_i or S_j that do not appear in $S_j \cap S_i$, whichever number is larger. We refer to this number as the codimension of S_i and S_j, and it is a measure of how much the two states have in common. The $\mathcal{R}_{\ell-c}(S_j \cap S_i)$ are $(\ell-c)$-connected by our assumption, but we need to prove that they are $(\ell-1)$-connected for our inductive argument to go through, and it is not generally true that the S_i and S_j have codimension $c = 1$.

One of the insights in this paper — and one of the reasons that the lower bound proof for set agreement is now so simple — is that we can organize the inductive argument so that we need only consider pairs of simplexes S_i and S_j in this union that have codimension $c = 1$. If we order the set $\mathcal{R}_\ell(S) = \{S_1, \ldots, S_k\}$ of one-round states correctly, then we can prove that every set $\mathcal{R}_\ell(S_j) \cap \mathcal{R}_\ell(S_i)$ in the union is contained in another set $\mathcal{R}_\ell(T_j) \cap \mathcal{R}_\ell(S_i)$ in the union such that T_j and S_i have codimension $c = 1$. The larger set "absorbs" the smaller set, and while the smaller set may not have the desired $(\ell-1)$-connectivity, the larger set does. Now we can write this union as the union of the absorbing sets, which is a union of $(\ell-1)$-connected sets, and apply Mayer-Vietoris to prove that the union itself is $(\ell-1)$-connected. In this paper, we show how to define a partial order on the set $\mathcal{R}_\ell(S) = \{S_1, \ldots, S_k\}$ of one-round states that guarantees this absorption property holds during the Mayer-Vietoris argument. We call this partial order an *absorbing poset*.

To illustrate the notion of an absorbing poset, consider once again the complex $\mathcal{R}_1(S)$ of global states on the right of Figure 1. Suppose we order the pseudospheres making up $\mathcal{R}_1(S)$ by ordering the central triangle first and then ordering the cycles surrounding this triangle in some order. Within each cycle, let us order the edges of the cycle by ordering the edge of the central triangle first, then the two edges intersecting this edge in some order, and finally outermost edge that does not intersect the central triangle. To see that this ordering has the properties of an absorbing poset, consider the central triangle T and the edge E consisting of the vertexes $P; PR$ and $R; PQR$. The simplexes T and E intersect in the single vertex $R; PQR$ and hence have codimension two. On the other hand, consider the edge F consisting of the vertexes $R; PQR$ and $P; PQR$. This edge F appears between T and E in the simplex ordering, the intersection of F and E is actually equal to the intersection of T and E, and the codimension of F and E is one. This property of an absorbing poset is key to the simplicity of the connectivity argument given in this paper.

2.3 Related work

We are aware of three other proofs of the k-set agreement lower bound.

Chaudhuri, Herlihy, Lynch, and Tuttle [7] gave the first proof. Their proof consisted of taking the standard similarity chain argument used to prove the consensus synchronous lower bound and running that argument in k dimensions at once to construct a subset of the reachable complex to which a standard topological tool called Sperner's Lemma can be applied to obtain the desired impossibility. While their intuition is geometrically compelling, it required quite a bit of technical machinery to nail down the details.

Herlihy, Rajsbaum, and Tuttle [15] gave a proof closer to our "round-by-round" approach. In fact, the round operator that we define here is exactly the round operator they defined. Their connectivity proof for the reachable complex was not easy, however, and the inductive nature of the proof did not reflect the iterative nature of how the reachable complex is constructed by repeatedly applying the round operator locally to a global state S. The notion of an absorbing poset used in this paper dramatically simplifies the connectivity proof.

Gafni [12] gave another proof in an entirely different style. His proof is based on simple reductions between models, showing that the asynchronous model can simulate the first few rounds of the synchronous model, and thus showing that the synchronous lower bound follows from the known asynchronous impossibility result for set agreement [4, 16, 21]. While his notion of reduction is elegant, his proof depends on the asynchronous impossibility result, and that result is not easy to prove. We are interested in a simple, self-contained proof that gives as much insight as possible into the topological behavior of the synchronous model of computation.

Round-by-round proofs that show how the 1-dimensional (graph) connectivity evolves in the synchronous model have been described by Aguilera and Toueg [1] and Moses and Rajsbaum [18] (the latter do it in a more general way that applies to various other asynchronous models as well) to prove consensus impossibility results. These show how to do an elegant FLP style of argument, as opposed to the more involved backward inductive argument of the standard proofs [10, 8, 9]. They present a (graph) connectivity proof of the successors of a global state. Thus, our proofs are similar to this strategy in the particular case of $k = 1$, but give additional insights because they show more general ways of organizing these connectivity arguments.

There are also various set agreement impossibility results for asynchronous systems that are related to our work.

Attiya and Rajsbaum [2] and Borowsky and Gafni [4] present two similar proofs for the set agreement impossibility. The relation to our work is that they are also combinatorial. However, their proofs are for an asynchronous, shared memory model. Also, they do not have a round-by-round structure; instead they work by proving that the set of global states at the end of the computation has some properties (somewhat weaker than connectivity) that are sufficient to apply Sperner's Lemma and obtain the desired impossibility result.

Borowsky and Gafni [5] defined an asynchronous shared-memory model where variables can be used only once, a model they showed to be equivalent to general asynchronous shared-memory models. They defined a round operator as we do, and they showed that one advantage of their model was that it had a very regular iterative structure that greatly simplified computing its connectivity. Unfortunately, their elegant techniques for the asynchronous model do not extend to the synchronous model. Reasoning about connectivity is harder in the synchronous model than the asynchronous model for two reasons. First, the connectivity never decreases in the asynchronous model, whereas it does in the synchronous model, so their techniques cannot extend to our model. Second, processors never actually fail in their construction since dead processors can be modeled as slow processors, but this is not an option in our model, and we are forced to admit simplexes of many dimensions as models of global states where processors have failed.

3 Topology

We now give formal definitions of the topological ideas sketched in the introduction.

A *simplex* is just a set of vertexes. Each vertex v is labeled with a processor id $id(v)$ and a value $val(v)$. We assume that the vertexes of a simplex are labeled with distinct processor ids, and we assume a total ordering $\leq_{\mathrm{id}}$ on processor ids, which induces an ordering on the vertexes of a simplex. A *face* of a simplex is a subset of the simplex's vertexes, and we write $F \subseteq S$ if F is a face of S. A simplex X is *between* two simplexes S_0 and S_1 if $S_0 \subseteq X \subseteq S_1$. A *complex* is a set of simplexes closed under containment (which means that if a simplex belongs to a complex, then so do its faces). If $\mathcal{A}$ is a set of simplexes, denote by $\|\mathcal{A}\|$ the smallest simplicial complex containing every simplex of $\mathcal{A}$. It is easy to show that

$$\|\mathcal{A} \cup \mathcal{B}\| = \|\mathcal{A}\| \cup \|\mathcal{B}\| \quad \text{and} \quad \|\mathcal{A} \cap \mathcal{B}\| \subseteq \|\mathcal{A}\| \cap \|\mathcal{B}\|.$$

The *codimension* of a set $\mathcal{S} = \{S_1, \ldots, S_m\}$ of simplexes is

$$codim(\mathcal{S}) = \max_i \{\dim(S_i) - \dim(\cap_j S_j)\} = \max_i \{|S_i - \cap_j S_j|\},$$

where $\dim(\emptyset) = -1$ is the dimension of the empty simplex. This definition satisfies several simple properties, such as:

1. If X is between two simplexes S and T, then
$$codim(S, T) = codim(S, X) + codim(X, T).$$
2. If $codim(S_0, S_i) \leq 1$ for $i = 1, \ldots, m$, then
$$codim(S_0, S_1, \ldots, S_m) \leq m.$$
3. If $S_1, \ldots, S_m$ is a set of simplexes with largest dimension N and codimension c, then their intersection $S_1 \cap \cdots \cap S_m$ is a simplex with dimension $N - c$.

The *connectivity* of a complex is a direct generalization of ordinary graph connectivity. A complex is 0-connected if it is connected in the graph-theoretic sense, and while the definition of k-connectivity is more involved, the precise definition does not matter here since our work depends only on two fundamental properties of connectivity:

Theorem 1.

1. *If S is a simplex of dimension k, then the induced complex $\|S\|$ is $(k-1)$-connected.*
2. *If the complexes $\mathcal{L}$ and $\mathcal{M}$ are k-connected and $\mathcal{L} \cap \mathcal{M}$ is $(k-1)$-connected, then $\mathcal{L} \cup \mathcal{M}$ is k-connected.*

By convention, a nonempty complex is (-1)-connected, and every complex is $(-k)$-connected for $k \geq 2$. The first property follows from the well-known fact that $\|S\|$ is k-connected when S is a nonempty simplex of dimension k ($k \geq 0$), but when S is empty ($k = -1$), the best we can say is that $\|S\|$ is (-2)-connected (since every complex is (-2)-connected). The second property above is a consequence of the well-known Mayer-Vietoris sequence which relates the topology of $\mathcal{L} \cup \mathcal{M}$ with that of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{L} \cap \mathcal{M}$ (for example, see Theorem 33.1 in [19]).

4 Computing connectivity

Computing the connectivity of $\mathcal{R}_\ell(\mathcal{A})$ for some complex $\mathcal{A}$ depends on properties of the round operator $\mathcal{R}_\ell$ and on how the Mayer-Vietoris argument used in the proof is organized. In this section, we define the notion of an f-operator and the notion of an absorbing poset that structures the Mayer-Vietoris argument by imposing a partial order on simplexes in the complex $\mathcal{A}$, and we prove that an f-operator applied to this partially-ordered complex $\mathcal{A}$ is connected.

A *simplicial operator* $\mathcal{Q}$ is a function with an associated *domain*. It maps every simplex S in its domain to a set $\mathcal{Q}(S)$ of simplexes, and it extends to sets of simplexes in its domain in the obvious way with $\mathcal{Q}(\mathcal{A}) = \cup_{S \in \mathcal{A}} \mathcal{Q}(S)$. Proving the connectivity of $\mathcal{Q}(\mathcal{A})$ is simplified if $\mathcal{Q}$ satisfies the following property:

Definition 1. *Let $\mathcal{Q}$ be an operator, and let f be a function that maps each set $\mathcal{A}$ of simplexes in the domain of $\mathcal{Q}$ to an integer $f(\mathcal{A})$. We say that $\mathcal{Q}$ is an f-*operator *if for every set $\mathcal{A}$ of simplexes in the domain of $\mathcal{Q}$*

$$\bigcap_{S \in \mathcal{A}} \|\mathcal{Q}(S)\| \text{ is } (f(\mathcal{A}) - c - 1)\text{-connected}$$

where $c = codim(\mathcal{A})$.

To illustrate this definition, consider a single simplex S of dimension k, and remember the fundamental fact of topology that $\|S\|$ is $(k-1)$-connected. Now consider two simplexes S and T of dimension k that differ in exactly one vertex and hence have codimension one. Their intersection $S \cap T$ has dimension $k-1$, so $\|S \cap T\|$ is $(k-1-1)$-connected. In fact, we can show that $\|S\| \cap \|T\|$ is $(k-1-1)$-connected and, in general, that $\|S\| \cap \|T\|$ is $(k-c-1)$-connected if c is the codimension of S and T. In other words, the connectivity of their intersection is reduced by their codimension. In the definition above, if we interpret $f(\mathcal{A})$ as the maximum connectivity of the complexes $\|\mathcal{Q}(S)\|$ taken over all simplexes S in $\mathcal{A}$, then this definition says that taking the intersection of the $\|\mathcal{Q}(S)\|$ reduces the connectivity by the codimension of the S. As a simple corollary, if we take the identity operator $I(S) = \{S\}$ and define $f(\mathcal{A}) = \max_{S\in\mathcal{A}} \dim(S)$ to be the maximum dimension of any simplex in $\mathcal{A}$, then we can prove that the identity operator is an f-operator.

Proving the connectivity of $\mathcal{Q}(\mathcal{A})$ is simplified if there is a partial order on the simplexes in $\mathcal{A}$ that satisfies the following absorption property:

Definition 2. *Given a simplicial operator $\mathcal{Q}$ and a nonempty partially-ordered set $(\mathcal{S}, \preceq)$ of simplexes in the domain of $\mathcal{Q}$, we say that $(\mathcal{S}, \preceq)$ is an* absorbing poset for $\mathcal{Q}$ *if for every two simplexes S and T in $\mathcal{S}$ with $T \not\preceq S$ there is $T_S \in \mathcal{S}$ with $T_S \preceq T$ such that*

$$\|\mathcal{Q}(S)\| \cap \|\mathcal{Q}(T)\| \subseteq \|\mathcal{Q}(T_S)\| \cap \|\mathcal{Q}(T)\| \tag{1}$$

$$codim(T_S, T) = 1. \tag{2}$$

For example, if $\mathcal{S}$ is totally ordered, then every intersection $\|\mathcal{Q}(S)\| \cap \|\mathcal{Q}(T)\|$ involving a simplex S preceding a simplex T is contained in another intersection $\|\mathcal{Q}(T_S)\| \cap \|\mathcal{Q}(T)\|$ involving another simplex T_S preceding T with the additional property that T_S and T have codimension 1.

To see why such an ordering is useful, consider the round operator $\mathcal{R}_\ell$, and remember the problem we faced in the overview of proving that

$$\mathcal{R}_\ell(S_1) \cup \mathcal{R}_\ell(S_2) \cup \cdots \cup \mathcal{R}_\ell(S_i)$$

is connected for $i = 1, \ldots, k$. We were concerned that $\mathcal{R}_\ell(S_j) \cap \mathcal{R}_\ell(S_i) = \mathcal{R}_{\ell-c}(S_j \cap S_i)$ was $(\ell-c)$-connected and in general might not be $(\ell-1)$-connected since the codimension c of S_j and S_i might be too high. If we can impose an ordering on the $S_1, \ldots, S_k$ and prove that the $S_1, \ldots, S_k$ form an absorbing poset for $\mathcal{R}_\ell$, then each $\mathcal{R}_\ell(S_j) \cap \mathcal{R}_\ell(S_i)$ with $j < i$ is contained in another $\mathcal{R}_\ell(S_{j'}) \cap \mathcal{R}_\ell(S_i)$ with $j' < i$ where $S_{j'}$ and S_i have codimension one. This means that when computing the connectivity of the union we can restrict our attention to the intersections $\mathcal{R}_\ell(S_{j'}) \cap \mathcal{R}_\ell(S_i)$ with codimension one, and the proof goes through.

In general, we can prove that applying an operator to an absorbing poset yields a connected complex:

Theorem 2. *If $\mathcal{Q}$ is an f-operator and $(\mathcal{A}, \preceq)$ is an absorbing poset for $\mathcal{Q}$, then*

$$\|\mathcal{Q}(\mathcal{A})\| = \bigcup_{S\in\mathcal{A}} \|\mathcal{Q}(S)\| \text{ is } (f-1)\text{-connected}$$

where $f = \min_{\mathcal{B}\subseteq\mathcal{A}} f(\mathcal{B})$.

In the special case of the identity operator, we say that $(\mathcal{A}, \preceq)$ is an absorbing poset if it is an absorbing poset for the identity operator. It is an easy corollary to show that if $(\mathcal{A}, \preceq)$ is an absorbing poset, then

$$\|\mathcal{A}\| = \bigcup_{S\in\mathcal{A}} \|S\| \text{ is } (N-1)\text{-connected},$$

where N is the minimum dimension of the simplexes in $\mathcal{A}$.

5 Synchronous connectivity

In this section, we show how to use the ideas of the previous section to prove that $\mathcal{R}_k^r(S)$ is $(k-1)$-connected, from which we conclude that k-set agreement is impossible to solve in r rounds.

5.1 Round operators

Given a simplex S representing the state at the beginning of a round, and given a set X of processors that fail during the round, let $F = S/X$ be the face of S obtained from S by deleting the vertexes labeled with processors in X. The set of all possible states at the end of a round of computation from S in which processors in X fail can be represented by the set of all possible simplexes obtained by labeling the vertexes of F with simplexes between F and S. This set of simplexes obtained in this way forms a set that we call a pseudosphere. For every simplex S, the *pseudosphere operator* $\mathcal{P}_S(F)$ maps a face F of S to the set of all labelings of F with simplexes between F and S. The set $\mathcal{P}_S(F)$ is called a *pseudosphere*, and the face F is called the *base simplex* of the pseudosphere. Given a simplex T contained in a pseudosphere $\mathcal{P}_S(F)$ we define $\mathit{base}(T)$ to be the base simplex F of the pseudosphere.

If ℓ processors fail during the round, there there are many ways to choose this set X of processors that fail, and hence many ways to choose the base simplexes $F = S/X$ for the pseudospheres whose simplexes represent the states at the end of the round. For every integer $\ell \geq 0$, the *ℓ-failure operator* $\mathcal{F}_\ell(S)$ maps a simplex S to the set of all faces F of S with $\mathit{codim}(F, S) \leq \ell$, which is the set of all faces obtained by deleting at most ℓ vertexes from S. The domain of the operator $\mathcal{F}_\ell(S)$ is the set of all simplexes S with $\dim(S) \geq \ell$.

Finally, for every integer $\ell \geq 0$, the *synchronous round operator* $\mathcal{R}_\ell(S)$ is defined by

$$\mathcal{R}_\ell(S) = \mathcal{P}_S(\mathcal{F}_\ell(S)).$$

The domain of this operator $\mathcal{R}_\ell(S)$ is the set of all simplexes S with $\dim(S) \geq \ell + k$. This round operator satisfies a number of basic properties such as:

Lemma 1.

1. *$\mathcal{R}_\ell(S) \subseteq \mathcal{R}_m(S)$ if $\ell \leq m$ and S is in the domain of $\mathcal{R}_m$.*
2. *$\mathcal{R}_\ell(S) \subseteq \mathcal{R}_{\ell+c}(T)$ if $S \subseteq T$ and $c = \mathit{codim}(S, T)$.*
3. *$\mathcal{R}_\ell(S_1) \cap \cdots \cap \mathcal{R}_\ell(S_m) = \mathcal{R}_{\ell-c}(S_1 \cap \cdots \cap S_m)$ if $c = \mathit{codim}(S_1, \ldots, S_m)$, $\ell \geq c$, and each S_i is in the domain of $\mathcal{R}_\ell$.*
4. *$\|\mathcal{R}_\ell(S_1)\| \cap \cdots \cap \|\mathcal{R}_\ell(S_m)\| = \|\mathcal{R}_\ell(S_1) \cap \cdots \cap \mathcal{R}_\ell(S_m)\|$.*

Proof. We sketch the proof of property 3.

For the $\supseteq$ containment, suppose $A \in \mathcal{R}_{\ell-c}(\cap_j S_j)$. This means that A is a labeling of a simplex F with simplexes between F and $\cap_j S_j$ for some face F of $\cap_j S_j$ satisfying $\mathit{codim}(F, \cap_j S_j) \leq \ell - c$. Since A is a labeling of F with simplexes between F and $\cap_j S_j$, it is obviously a labeling of F with simplexes between F and S_i. We have $A \in \mathcal{R}_\ell(S_i)$ since F is a face of $\cap_j S_j$ which is in turn a face of S_i, and hence

$$\begin{aligned} codim(F, S_i) &= codim(F, \cap_j S_j) + codim(\cap_j S_j, S_i) \\ &\leq codim(F, \cap_j S_j) + codim(S_1, \ldots, S_m) \\ &\leq (\ell - c) + c = \ell. \end{aligned}$$

For the $\subseteq$ containment, suppose $A \in \cap_j \mathcal{R}_\ell(S_j)$. For each i, we know that A is a labeling of F_i with simplexes between F_i and S_i for some face F_i of S_i satisfying $codim(F_i, S_i) \leq \ell$. Since A is a labeling of F_i for each i, it must be that the F_i are all equal, so let F be this common face of the S_i and hence of $\cap_j S_j$. Since A is a labeling of F with simplexes between F and S_i for each i, it must be that A is a labeling of F with simplexes between F and $\cap_j S_j$. Since F is a face of $\cap_j S_j$ which is in turn a face of each S_i, including any S_M satisfying $codim(\cap_j S_j, S_M) = codim(S_1, \ldots, S_M) = c$, we have $A \in \mathcal{R}_{\ell-c}(\cap_j S_j)$ since $codim(F, \cap_j S_j) = codim(F, S_M) - codim(\cap_j S_j, S_M) \leq \ell - c$. □

The round operator $\mathcal{R}_\ell$ models a single round of computation. We model multiple rounds of computation with the multi-round operator $\mathcal{R}_L^r \mathcal{R}_\ell$ defined inductively by $\mathcal{R}_L^0 \mathcal{R}_\ell(S) = \mathcal{R}_\ell(S)$ and $\mathcal{R}_L^r \mathcal{R}_\ell(S) = \mathcal{R}_L(\mathcal{R}_L^{r-1} \mathcal{R}_\ell(S))$ for $r > 0$. The domain of $\mathcal{R}_L^r \mathcal{R}_\ell$ is the set of all simplexes S with $\dim(S) \geq rL + \ell + k$. The properties of one-round operators given above generalize to multi-round operators where $\mathcal{R}_\ell$ is replaced by $\mathcal{R}_L^r \mathcal{R}_\ell$.

5.2 Absorbing posets

We now impose a partial order on $\mathcal{R}_\ell(S)$ and prove that it is an absorbing poset. First we order the pseudospheres $\mathcal{P}_S(F)$ making up $\mathcal{R}_\ell(S)$, and then we order the simplexes within each pseudosphere $\mathcal{P}_S(F)$.

Both of these orders depend on ordering the faces F of S, which we do lexicographically. First we order the faces F by decreasing dimension, so that large faces occur before small faces. Then we order faces of the same dimension with a rather arbitrary rule using on our total order on processor ids: we order F_0 before F_1 if the smallest processor id labeling vertexes in F_0 and not F_1 comes before the smallest processor id labeling F_1 and not F_0. Formally:

Definition 3. *Define the total order $\leq_f$ on the faces of a simplex S by $F_0 \leq_f F_1$ iff*

1. $\dim(F_0) > \dim(F_1)$ *or*
2. $\dim(F_0) = \dim(F_1)$ *and either*
   (a) $F_0 = F_1$ *or*
   (b) $F_0 \neq F_1$ *and* $p_0 <_{id} p_1$
   where $p_0 = \min\{ids(F_0) - ids(F_1)\}$ *and* $p_1 = \min\{ids(F_1) - ids(F_0)\}$.

This face ordering induces an ordering on pseudospheres: $\mathcal{P}_S(F_0)$ comes before $\mathcal{P}_S(F_1)$ if F_0 comes before F_1 in the face ordering. This face ordering also induces an ordering on the simplexes within a single pseudosphere $\mathcal{P}_S(F)$: S_0 comes before S_1 if for each vertex v of the base simplex F the face labeling v in S_0 comes before the face labeling v in S_1. This ordering of $\mathcal{P}_S(F)$ is defined formally as follows:

Definition 4. *Define the partial order $\preceq_p$ on $\mathcal{P}_S(F)$ by $S_0 \preceq_p S_1$ iff $S_{0,v} \leq_f S_{1,v}$ for each vertex v in F where $S_{0,v}$ and $S_{1,v}$ are the simplexes labeling the vertex v in S_0 and S_1.*

Pulling everything together, the partial order on $\mathcal{R}_\ell(S)$ is defined as follows:

Definition 5. *Define the partial order $\preceq_r$ on $\mathcal{R}_\ell(S)$ by $S_0 \preceq_r S_1$ iff*

1. different pseudospheres: $\mathit{base}(S_0) <_f \mathit{base}(S_1)$ *or*
2. same pseudosphere: $\mathit{base}(S_0) = \mathit{base}(S_1)$ *and* $S_0 \preceq_p S_1$

Now we can prove that $(\mathcal{R}_\ell(S), \preceq_r)$ is an absorbing poset, and that

Lemma 2. *$(\mathcal{R}_\ell(S), \preceq_r)$ is an absorbing poset for $\mathcal{R}^r_L$ for $\dim(S) \geq rL + \ell + k$.*

Proof. We prove only the base case that $(\mathcal{R}_\ell(S), \preceq_r)$ is an absorbing poset. Let A and B be two simplexes in $\mathcal{R}_\ell(S)$ satisfying $B \not\preceq_r A$.

Case 1: Suppose A and B are in the same pseudosphere $\mathcal{P}_S(F)$ for some face F of S. The simplexes A and B are labelings of F with simplexes between F and S, so let A_v and B_v denote the label of vertex v in A and B for every vertex $v \in F$. There must be some vertex v with $A_v <_f B_v$ since $B \not\preceq_r A$. Let B_A be B with the label of v changed from B_v to A_v. We have $B_A \prec_r B$ since the label of v in B_A is ordered before the label of v in B, and the labels of all other vertices are equal. We have $A \cap B \subseteq B_A$ since v is not in $A \cap B$ due to the conflicting labels for v, while all other vertexes of B and hence of $A \cap B$ are in B_A. Finally, we have $\mathit{codim}(B_A, B) = 1$ since B_A and B differ only in the label of v.

Case 2: Suppose A and B are in different pseudospheres $\mathcal{P}_S(F_A)$ and $\mathcal{P}_S(F_B)$ for distinct faces F_A and F_B of S. We can assume without loss of generality that every vertex of $B - A$ is labeled with S, and we can show that $F_A <_f F_B$.

Case 2a: Suppose $\dim(F_A) > \dim(F_B)$. Since $\dim(F_A) > \dim(F_B)$, the set $F_A - F_B$ must be nonempty, so choose any vertex $v \in F_A - F_B$. Since $A \in \mathcal{P}_S(F_A)$, the simplex A must be a labeling of F_A with simplexes between F_A and S. Since v is a vertex of F_A, this means that v appears in all simplexes labeling A, and hence in all simplexes labeling $A \cap B$. Since we have assumed that S is the label of every vertex in $B - A$, and since S certainly contains the vertex v, the vertex v appears in all labels of $B - A$. It follows that v appears in every simplex labeling B. Let B_A be the simplex consisting of B together with the vertex v labeled with S, and notice that B_A is a simplex in $\mathcal{R}_\ell(S)$. We have $B_A \prec_r B$ since $\dim(B_A) = \dim(B) + 1 > \dim(B)$. We have $A \cap B \subseteq B_A$ since $A \cap B \subseteq B \subseteq B_A$. We have $\mathit{codim}(B, B_A) = 1$ since B and B_A differ only in v.

Case 2b: Suppose $\dim(F_A) = \dim(F_B)$, in which case we have $p_A \prec_p p_B$ where $p_A = \min\{\mathit{ids}(F_A) - \mathit{ids}(F_B)\}$ and $p_B = \min\{\mathit{ids}(F_B) - \mathit{ids}(F_A)\}$. Let v_A and v_B be the vertexes for processors p_A and p_B in the faces F_A and F_B of S. Let F_C be the face of S obtained from F_B by replacing v_B with v_A, and let C be the labeling of F_C obtained by labeling v_A with S and every other vertex with its label in B. Since A is a labeling of F_A with simplexes between F_A and S, and since v_A is a vertex of F_A, the vertex v_A appears in every simplex labeling A and hence $A \cap B$; and since we are assuming that every vertex of $B - A$ is labeled with S which certainly contains v_A, it follows that every vertex of $B - A$ is labeled with a simplex containing v_A; and hence it follows that every label in B contains F_C. It follows that $C \in \mathcal{R}_\ell(S)$ since C is a labeling of a face F_C of S with simplexes between F_C and S. We have $C \prec_r B$ since

$$\min \{ids(F_C) - ids(F_B)\} = p_A \prec_p p_B = \min \{ids(F_B) - ids(F_C)\}.$$

We have $A \cap B \subseteq C$ and $codim(B, C) = 1$. Taking $B_A = C$, we are done. □

5.3 Connectivity

All that remains is to prove that $\mathcal{R}_k^r$ is a k-operator. This follows from the following pair of statements proven by mutual induction:

Theorem 3. *For all $r \geq 0$,*

1. *$\|\mathcal{R}_L^r \mathcal{R}_\ell(S)\|$ is $(k-1)$-connected for all $\ell \geq 0$, all $L \geq k$, and all S in the domain of $\mathcal{R}_L^r \mathcal{R}_\ell$.*
2. *$\mathcal{R}_L^r \mathcal{R}_\ell$ is a k-operator for all $L, \ell \geq k$.*

Since $\mathcal{R}_k^r$ is a k-operator, and since $(\mathcal{R}_k(S), \preceq_r)$ is an absorbing poset for $\mathcal{R}_k^r$, the connectivity follows by Theorem 2:

Corollary 1. *$\|\mathcal{R}_k^r(S)\|$ is $(k-1)$-connected if $\dim(S) \geq (r+1)k$.*

6 Conclusion

As we have said, the impossibility of k-set agreement now follows directly from the connectivity of $\|\mathcal{R}_k^r(S)\|$ using standard arguments based on variants of Sperner's Lemma that have appeared in several places now. We hope that the notions of a round operator and an absorbing poset will yield simple proofs of other results, and will show the way toward simple proofs in other models of computation such as the asynchronous message-passing model.

References

1. Marcos Kawazoe Aguilera and Sam Toueg. A simple bivalency proof that t-resilient consensus requires $t+1$ rounds. *Information Processing Letters*, 71(3-4):155–158, August 1999.
2. Hagit Attiya and Sergio Rajsbaum. The combinatorial structure of wait-free solvable tasks. In *Proceedings of the 10th International Workshop on Distributed Algorithms*, volume 1151 of *Lecture Notes in Computer Science*, pages 322–343. Springer-Verlag, Berlin, October 1996.

3. Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics*. McGraw–Hill, 1998.
4. Elizabeth Borowsky and Eli Gafni. Generalized FLP impossibility result for t-resilient asynchronous computations. In *Proceedings of the 25th ACM Symposium on Theory of Computing*, pages 91–100, May 1993.
5. Elizabeth Borowsky and Eli Gafni. A simple algorithmically reasoned characterization of wait-free computations. In *Proceedings of the 16th Annual ACM Symposium on Principles of Distributed Computing*, pages 189–198, August 1997.
6. Soma Chaudhuri. Agreement is harder than consensus: set consensus problems in totally asynchronous systems. In *Proceedings of the 9th Annual ACM Symposium on Principles of Distributed Computing*, pages 311–234, August 1990.
7. Soma Chaudhuri, Maurice Herlihy, Nancy Lynch, and Mark R. Tuttle. Tight bounds for k-set agreement. *Journal of the ACM*, 47(5):912–943, September 2000.
8. Danny Dolev and Ray Strong. Polynomial algorithms for multiple processor agreement. In *Proceedings of the 14th ACM Symposium on Theory of Computing*, pages 401–407, 1982.
9. Cynthia Dwork and Yoram Moses. Knowledge and common knowledge in a Byzantine environment: Crash failures. *Information and Computation*, 88(2):156–186, October 1990.
10. Michael J. Fischer and Nancy A. Lynch. A lower bound for the time to assure interactive consistency. *Information Processing Letters*, 14(4):183–186, June 1982.
11. Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty processor. *Journal of the ACM*, 32(2):374–382, April 1985.
12. Eli Gafni. A round-by-round failure detector — unifying synchrony and asynchrony. In *Proceedings of the 17th Annual ACM Symposium on Principles of Distributed Computing*. ACM, June 1998.
13. Maurice Herlihy and Sergio Rajsbaum. Set consensus using arbitrary objects. In *Proceedings of the 13th Annual ACM Symposium on Principles of Distributed Computing*, pages 324–333, August 1994.
14. Maurice Herlihy and Sergio Rajsbaum. Algebraic spans. *Mathematical Structures in Computer Science*, 10(4):549–573, August 2000. Special Issue: Geometry and Concurrency.
15. Maurice Herlihy, Sergio Rajsbaum, and Mark R. Tuttle. Unifying synchronous and asynchronous message-passing models. In *Proceedings of the 17th Annual ACM Symposium on Principles of Distributed Computing*, pages 133–142. ACM, June 1998.
16. Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *Journal of the ACM*, 46(6):858–923, November 1999.
17. Nancy Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
18. Yoram Moses and Sergio Rajsbaum. The unified structure of consensus: a layered analysis approach. In *Proceedings of the 17th Annual ACM Symposium on Principles of Distributed Computing*, pages 123–132. ACM, June 1998. To apper in *SIAM Journal of Computing.*
19. James R. Munkres. *Elements Of Algebraic Topology.* Addison Wesley, Reading MA, 1984.
20. Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2):228–234, 1980.
21. Michael Saks and Fotios Zaharoglou. Wait-free k-set agreement is impossible: The topology of public knowledge. *SIAM Journal on Computing*, 29(5):1449–1483, 2000.