Lower Bounds for k-Set Agreement in Fault-Prone Networks

Pierre Fraigniaud
Inst. de Recherche en Informatique Fondamentale (IRIF), CNRS and Université Paris Cité, France

Minh Hang Nguyen
Inst. de Recherche en Informatique Fondamentale (IRIF), CNRS and Université Paris Cité, France

Ami Paz
Lab. Interdisciplinaire des Sciences du Numérique (LISN), CNRS & Université Paris-Saclay, France

Ulrich Schmid
TU Wien, Austria

Hugo Rincon-Galeana
TU Berlin, Germany

Abstract

We develop a new lower bound for k-set agreement in synchronous message-passing systems connected by an arbitrary directed communication network, where up to t processes may crash. Our result thus generalizes the $\lfloor t/k \rfloor + 1$ lower bound for complete networks in the t-resilient model by Chaudhuri, Herlihy, Lynch, and Tuttle [JACM 2000]. Moreover, it generalizes two lower bounds for oblivious algorithms in synchronous systems connected by an arbitrary undirected communication network known to the processes, namely, the domination number-based lower bound by Castañeda, Fraigniaud, Paz, Rajsbaum, Roy, and Travers [TCS 2021] for failure-free processes, and the radius-based lower bound in the t-resilient model by Fraigniaud, Nguyen, and Paz [STACS 2024].

Our topological proof non-trivially generalizes and extends the connectivity-based approach for the complete network, as presented in the book by Herlihy, Kozlov, and Rajsbaum (2013). It is based on a sequence of shellable carrier maps that, starting from a shellable input complex, determine the evolution of the protocol complex: During the first $\lfloor t/k \rfloor$ rounds, carrier maps that crash exactly k processes per round are used, which ensure high connectivity of their images. A Sperner's lemma style argument can thus be used to prove that k-set agreement is still impossible by that round. From round $\lfloor t/k \rfloor + 1$ up to our actual lower bound, a novel carrier map is employed, which maintains high connectivity. As a by-product, our proof also provides a strikingly simple lower-bound for k-set agreement in synchronous systems with an arbitrary communication network, where exactly $t \geq 0$ processes crash initially, i.e., before taking any step. We demonstrate that the resulting additional agreement overhead can be expressed via an appropriately defined radius of the communication graphs, and show that the usual input pseudosphere complex for k-set agreement can be replaced by an exponentially smaller input complex based on Kuhn triangulations, which we prove to be also shellable.

2012 ACM Subject Classification Theory of computation → Distributed computing models

Keywords and phrases Distributed computing, k-set agreement, time complexity, lower bounds, topology

Digital Object Identifier 10.4230/LIPIcs.DISC.2025.31

Related Version *Full Version*: `http://arxiv.org/abs/2508.15562` [11]

Funding *Pierre Fraigniaud*: Additional support from ANR projects DUCAT (ANR-20-CE48-0006), ENEDISC (ANR-24-CE48-7768-01), and PREDICTIONS (ANR-23-CE48-0010), and from the InIDEX Project METALG.
Minh Hang Nguyen: Additional support from ANR projects DUCAT (ANR-20-CE48-0006), TEMPORAL (ANR-22-CE48-0001), ENEDISC (ANR-24-CE48-7768-01), and the European Union's

Horizon 2020 program H2020-MSCA-COFUND-2019 Grant agreement n° 945332.
Hugo Rincon-Galeana: Supported by the German Research Foundation (DFG) SPP 2378 (project ReNO).

1 Introduction

In the k-set agreement task, introduced by Chaudhuri [7], each process starts with some input value belonging to an finite set of possible input values, and must irrevocably output a value usually referred to as its *decision* value. The output value decided by a process has to be the input of some process (strong validity condition), and, system-wide, no more that k different decision values may be decided (k-agreement condition). Whereas the case $k = 1$ (consensus) is well-understood, properly understanding k-set agreement for general $k > 1$ is notoriously difficult, even for simple computing models. Besides the inherent difficulty of handling a task that is less constrained than consensus, its analysis is considerably complicated by the fact that "classic" proof techniques are inadequate [1–3]. As a consequence, methods from combinatorial topology must usually be resorted to [12]. Such methods are very powerful, but often difficult to apply to concrete scenarios.

Unsurprisingly, these complications affect not only impossibility proofs for k-set agreement, but also termination-time lower bounds. In particular, in message-passing synchronous systems (where the processes communicate with each other in a sequence of synchronous, communication-closed *rounds* over some communication network), we are aware of only two substantially new results since the seminal tight $\lfloor t/k \rfloor + 1$ lower-bound established by Chaudhuri, Herlihy, Lynch, and Tuttle [8] for complete networks in the t-resilient model (where at most t processes may fail by crashing during any run).

The first one is the lower-bound by Castañeda, Fraigniaud, Paz, Rajsbaum, Roy, and Travers [5] (see also [10]) for *failure-free* processes connected by an arbitrary (connected) bidirectional communication network G that is commonly known to all nodes – this model is referred to as the KNOW-ALL model. It holds for *oblivious* algorithms only, that is, algorithms which exchange the sets of different input values seen so far using flooding-based communication, and only take decisions by these sets (and not, e.g., the time of a message arrival or the neighbor it arrived from). The lower bound essentially states that r rounds are necessary, where r is the smallest integer such that the graph $G_r = (V, E_r)$ obtained from $G = (V, E)$ by connecting by an edge every two nodes at distance at most r in G has domination number at most k.

The second one is the lower bound, established by Fraigniaud, Nguyen, and Paz [9], for *oblivious* algorithms in the t-resilient model with an arbitrary undirected communication network G. For $k = 1$, it essentially states that consensus requires r rounds, where $r = \text{radius}(G, t)$ is the *radius* of the network when up to t nodes may fail by crashing. Informally, the radius is defined as the minimum, taken over all nodes of the network, of the worst-case *finite* number of rounds required for broadcasting from a node over all possible failure patterns, hence can be defined via the *eccentricity* of certain nodes in G. The lower bound in [9] is tight for oblivious algorithms thanks to the algorithm in [6]. The consensus lower bound can be extended to k-set agreement using the same techniques as [10], but only if assuming a priori knowledge on the failure pattern.

In the current paper, we generalize the above results by developing a lower bound for the number of rounds for solving k-set agreement in the t-resilient model for arbitrary directed communication networks. It is fomulated via the "agreement overhead" caused by the presence of an arbitrary communication network over the mere case of the complete network.

▶ **Definition 1.1.** *Let G be a directed graph, and let $k \geq 1$ and $t \geq 0$ be integers. The agreement overhead* $\mathrm{ovh}(G, k, t)$ *is the smallest integer such that k-set agreement in G can be solved in $\lfloor t/k \rfloor + 1 + \mathrm{ovh}(G, k, t)$ rounds in the t-resilient model.*

The agreement overhead can hence be viewed as the penalty for not using the complete network but solely G. For the n-process complete network K_n, for every $k \geq 1$ and $t \geq 0$, $\mathrm{ovh}(K_n, k, t) = 0$, thanks to the lower bound established in [8].

1.1 Contributions

Our main lower bound result (Theorem 4.10) relies on two cornerstones:

(1) a proof that the $\lfloor t/k \rfloor + 1$ lower-bound for t-resilient systems over the complete communication network [8] also holds for every arbitrary network (which motivates the notion of agreement overhead), and

(2) a lower bound on the agreement overhead for an arbitrary communication graph $G = (V, E)$. For specifying the latter, recall that, for every $U \subseteq V$, $G[U]$ denotes the subgraph of G induced by the vertices in U. Given a set D of vertices, we denote by $\mathrm{ecc}(D, G)$ the eccentricity of D in G, i.e., the number of rounds D need to collectively broadcast to all the graph's nodes. We then define the (t, k)-radius of G as

$$\mathrm{rad}(G, t, k) = \min_{D \subseteq V, |D| = t+k} \max_{D' \subseteq D, |D'| = t} \mathrm{ecc}(D \setminus D', G[V \setminus D']) \tag{1}$$

and show that the agreement overhead satisfies $\mathrm{ovh}(G, k, t) \geq \mathrm{rad}(G, k, t) - 1$.

Consequently, any algorithm solving k-set agreement in G under the t-resilient model must perform at least $\lfloor \frac{t}{k} \rfloor + \mathrm{rad}(G, t, k)$ rounds (Theorem 4.10). For the special case of $t = 0$, our lower bound is equivalent to the one established in [5] for the KNOW-ALL model. For $t > 0$, our lower bound on the agreement overhead also gives a lower bound of $\mathrm{rad}(G, k, t)$ for solving k-set agreement in arbitrary networks with t initially dead processes (Theorem 4.11).

Our paper also advances the state of the art of topological modeling as follows:

(3) We introduce a novel carrier map that governs the evolution of a shellable protocol complex in systems connected by an arbitrary directed communication graph G (that may even vary from round to round) with t initially dead processes, and show that it maintains high connectivity during $\mathrm{ovh}(G, k, t)$ rounds. For $t = 0$, our carrier map allows a much simpler analysis of the setting studied in [5]. For $t > 0$, we also demonstrate how to generalize the scissors cut-based analysis in [5] for handling the case $t > 0$ as well, and show that the resulting lower bound is equivalent to the one $\mathrm{ovh}(G, k, t) + 1$ established by our approach.

(4) We non-trivially generalize, extend and also correct the topological proof technique for the $\lfloor t/k \rfloor + 1$ lower bound in complete networks sketched in [12, Ch. 13] to arbitrary directed communication graphs (that may also vary from round to round). Our approach starts out from a shellable input complex, and utilizes a sequence of shellable carrier maps that crash exactly k processes per round for modeling the evolution of the protocol complex. Since these carrier maps maintain high connectivity during the first $\lfloor t/k \rfloor$ rounds, a Sperner-lemma style argument can be used to prove that k-set agreement is still impossible. Our contribution not only adds details missing in [12, Ch. 13] (e.g., the strictness proof of the carrier maps, and the Sperner-style argument), but also fixes a non-trivial error by replacing the rigidity requirement for the carrier maps (which does not hold) by a novel, weaker condition.

(5) We prove that the Kuhn triangulation [8], which is exponentially smaller than the standard pseudosphere complex used as the input complex for k-set agreement in [12, Ch. 13], is shellable. We can hence seamlessly replace the pseudosphere input complex in our analysis by Kuhn triangulations.

Whereas the focus of our results are lower bounds, the question of tightness obviously arises. So far, we do not know whether and for which choices of G, k and t our lower bound in Theorem 4.10 is tight. We must hence leave this question to future research. We nevertheless include the following upper-bound result:

(6) We present an upper bound on the agreement overhead by generalizing the algorithm for the clique K_n in [8] to an arbitrary communication network $G = (V, E)$. For $S \subseteq V$, let $G[V \setminus S]$ denote the subgraph of G induced by the nodes in $V \setminus S$, and let $D(G, t) = \max_{S \subseteq V, |S| \leq t} \operatorname{diam}(G[V \setminus S])$, where diam denotes the diameter. By following the arguments in [8], we show that there exists an algorithm solving k-set agreement in G in $\lfloor \frac{t}{k} \rfloor + D(G, t)$ rounds. As a consequence, $\operatorname{ovh}(G, k, t) \leq D(G, t) - 1$.

Paper organization.[1] In Section 2, we introduce our system model, and the basics of the topological modeling. In Section 3, we revisit the topological round-by-round connectivity analysis of [12, Ch. 13]. Section 4 provides our lower bounds on the agreement overhead, and Section 5 provides our upper-bound result. Some conclusions in Section 6 complete our paper. All our proofs are provided in the appendix.

2 System Model

2.1 Computational Model

Our computational model is similar to the one used by Fraigniaud, Nguyen and Paz [9], albeit we consider full-history protocols, and general (i.e., non-necessarily oblivious) algorithms. We consider a finite set of n processes with names $\Pi = \{p_1, \ldots, p_n\}$, that are ordered according to their index set $[n] = 1, \ldots, n$. Processes communicate in lock-step synchronous rounds via point-to-point directed links, that is, any message that is sent during a round r will be received in the same round r, and we do not consider the possibility that messages arrive at later rounds. We consider that all processes start simultaneously at round 1.

We assume that processes are represented by deterministic state machines, and have a well-defined local state that also includes the complete history of received messages. Thus, we consider a protocol to be defined by state transitions as well as a communication function and a decision function. In this paper, we will not focus on the protocol specifications, since it is fairly simple to derive them from the particular protocols that we consider. Instead, for the sake of readability and succinctness, we will sketch the protocols by specifying the messages that a process is able to send at each round, the information captured by the local states, and whether or not a process is ready to decide on an output value.

In a given round, processes can communicate using a fixed network topology, represented by a directed *communication graph* $G = (V, E)$, where $V = \Pi$. That is, a process p can send a message directly to another process q if and only if $(p, q) \in E$. We assume that E contains all self-loops $\{(p, p) \mid p \in V\}$. The processes are aware of the communication graph G. A

[1] Lack of space does not allow us to include all our results, which can be found in the full version [11] of our paper.

process p is able to send a message to any other process in its set of out-neighbors in G, denoted by $\mathrm{Out}_p(G) := \{q \in V \mid (p,q) \in E\}$. Symmetrically, p can only receive a message from a process in its set of in-neighbors in G, denoted by $\mathrm{In}_p(G) := \{q \in V \mid (q,p) \in E(G)\}$.

We consider the t-resilient model, where up to t processes may permanently crash in every execution, in any round. Crashes may be unclean, thus a process p may still send a message to a non-empty subset of $\mathrm{Out}_p(G)$ before crashing. The set of faulty processes crashing in a given execution is denoted as F with $|F| \leq t$, and is arbitrary and unknown to the processes.

For the sake of completeness, we also provide a formal description of the k-set agreement problem, which constitutes the main focus of this paper. Every process p_i has a local *input value* x_i taken arbitrarily from a finite set $\mathcal{V}$ with $|\mathcal{V}| \geq k+1$, which is often assumed to be just $\mathcal{V} = [k+1] = \{1, \ldots, k+1\}$. Every correct process p_i must irrevocably assign some decision value to a local *output variable* y_i eventually, which is initialized to $y_i = \perp$ with $\perp \notin \mathcal{V}$. In essence, k-set agreement is a relaxed instance of consensus, in which the agreement condition is relaxed to accept at most k different decision values. More precisely, k-set agreement is defined by the following conditions:

- **Strong Validity:** If a process p_i decides output value y_i, then y_i is the input value x_j of a process p_j
- ***k*-Agreement:** In every execution, if $\mathcal{O}$ denotes the set of all decision values of the processes that decide in that execution, then $|\mathcal{O}| \leq k$.
- **Termination:** Every non-faulty process $p_i \notin F$ must eventually and irrevocably decide on some value $y_i \neq \perp$.

2.2 Basics of combinatorial topology

Our analysis of k-set agreement relies on combinatorial topology [12]. Most notably, we develop novel topological techniques that allow us to ensure high-order connectivity, which seamlessly translates to a lower bound for k-set agreement. We now provide some basic definitions on simplicial complexes, which will be used heavily in the paper.

Intuitively, simplicial complexes may be thought of as a "higher dimensional" instance of an undirected graph. Indeed, in addition to vertices and edges, a simplicial complex may have faces of higher dimension.

▶ **Definition 2.1** (Simplicial Complex). *A pair $\mathcal{K} = (V(\mathcal{K}), F(\mathcal{K}))$, where $V(\mathcal{K})$ is a set, and $F(\mathcal{K}) \subseteq 2^{V(\mathcal{K})} \setminus \varnothing$ is a collection of subsets of $V(\mathcal{K})$ is a simplicial complex if, for any $\sigma \in F(\mathcal{K})$, and any $\sigma' \subseteq \sigma$, $\sigma' \in F(\mathcal{K})$. $V(\mathcal{K})$ is called the vertex set, and $F(\mathcal{K})$ the set of faces called simplices (singular: simplex). For notational simplicity, we will occasionally refer to simplicial complexes as* complexes.

Note that, following the convention in [12], we will very rarely (cf. Definition 3.4) also consider the empty "simplex" $\varnothing$.

The maximal faces (by inclusion) of a simplicial complex are called *facets*. Since faces are downward closed, then the facets are sufficient for fully determining a simplicial complex. The dimension of a face σ is defined as $\dim(\sigma) = |\sigma| - 1$. The dimension of a simplicial complex $\mathcal{K} = (V(\mathcal{K}), F(\mathcal{K}))$ is defined as $\max_{\sigma \in F(\mathcal{K})} \dim(\sigma)$. A complex is *pure* if all of its facets have the same dimension, and a complex is *impure* if it is not pure. For a face σ of a pure complex with facet dimension d, we denote by $\mathrm{codim}(\sigma) = d - \dim(\sigma)$ the *co-dimension* of σ, and by $\mathrm{Face}_k\, \sigma = \{\rho \mid \rho \subseteq \sigma \text{ with } \dim(\rho) = k\}$ the set of all k-faces of σ.

For any two simplicial complexes $\mathcal{K}$ and $\mathcal{L}$, $\mathcal{L}$ is a *subcomplex* of $\mathcal{K}$, denoted by $\mathcal{L} \subseteq \mathcal{K}$, if $V(\mathcal{L}) \subseteq V(\mathcal{K})$ and $F(\mathcal{L}) \subseteq F(\mathcal{K})$. For any $d \geq 0$, the d-skeleton $\mathrm{skel}_d(\mathcal{K})$ is the subcomplex of $\mathcal{K}$ consisting of all simplices of dimension at most d.

The morphisms (i.e., structure-preserving maps) for simplicial complexes are called *simplicial* maps:

▶ **Definition 2.2** (Simplicial map). *Let $\mathcal{K}$ and $\mathcal{L}$ be simplicial complexes. A mapping $\mu : V(\mathcal{K}) \to V(\mathcal{L})$ is a simplicial map if, for every $\sigma \in F(\mathcal{K})$, $\mu(\sigma) \in F(\mathcal{L})$.*

For our analysis, we also need to consider other maps beyond simplicial maps. Since we are interested in the evolution of configurations of processes, which are represented via faces of a simplicial complex, we need to consider functions that map individual simplices to sets of simplices.

▶ **Definition 2.3** (Carrier maps). *Let $\mathcal{K}$ and $\mathcal{L}$ be simplicial complexes, and $\Psi : F(\mathcal{K}) \to 2^{F(\mathcal{L})}$ be a function that maps faces of $\mathcal{K}$ into sets of faces of $\mathcal{L}$ such that, for every simplex $\sigma \in \mathcal{K}$, $\Psi(\sigma)$ is a subcomplex of $\mathcal{L}$. Ψ is a* carrier map *if, for every two simplices σ and κ in $F(\mathcal{K})$, $\Psi(\sigma \cap \kappa) \subseteq \Psi(\sigma) \cap \Psi(\kappa)$. Moreover,*

- *Ψ is* strict *if $\Psi(\sigma \cap \kappa) = \Psi(\sigma) \cap \Psi(\kappa)$, and*
- *Ψ is* rigid *if, for every $\sigma \in F(\mathcal{K})$, $\Psi(\sigma)$ is pure and of dimension $\dim(\sigma)$.*

Note that the definition above also allows $\Psi(\sigma) = \emptyset$, the empty complex.

In addition to the vertices and faces, a simplicial complex $\mathcal{K}$ may be endowed with a vertex coloring $\chi : V(\mathcal{K}) \to \mathcal{C}$, where $\mathcal{C}$ is the *color set*. We say that a vertex coloring χ is proper on $\mathcal{K}$ if for any simplex $\sigma \in F(\mathcal{K})$, the restriction $\chi_{|\sigma}$ of χ on σ is injective. We say that a pair $\mathcal{K}_\chi := (\mathcal{K}, \chi)$ is a *chromatic simplicial complex* if $\mathcal{K}$ is a simplicial complex, and $\chi : V(\mathcal{K}) \to \mathcal{C}$ is a proper vertex coloring. (As we shall see in the next section, in the context of using complexes to model distributed computing, the color of a vertex is merely a process ID.) Let $\mathcal{K}_\chi := (\mathcal{K}, \chi)$ and $\mathcal{L}_{\chi'} := (\mathcal{L}, \chi')$ be chromatic simplicial complexes. A simplicial map $\mu : V(\mathcal{K}) \to V(\mathcal{L})$ is a *chromatic map* if, for every $v \in V(\mathcal{K})$, $\chi(v) = \chi'(\mu(v))$, i.e., μ is color-preserving. For notational simplicity, when it is clear from the context, we will omit mentioning the vertex coloring explicitly.

2.3 Topological modeling of distributed systems

Simplicial complexes are particularly useful for representing system configurations, both for inputs and outputs, and for describing mid-run states. Vertices are used for representing local states, while faces represent (partial) configurations.

The *input complex* $\mathcal{I}$ is used for representing all possible initial configurations. Its vertices (p_i, x_i) consist of a process name $p_i = \chi((p_i, x_i)) \in \Pi$ that is used as its color, and some input value $x_i \in \mathcal{V}$. A facet σ of the input complex consists of n vertices $v_1, \ldots, v_n$, with $v_i = (p_i, x_i)$ for $1 \leq i \leq n$ that represent some initial configuration.

The *output complex* $\mathcal{O}$ is used for representing all possible decision configurations. Its vertices (p_i, y_i) consist of a process name $p_i = \chi((p_i, y_i)) \in \Pi$ that is used as its color, and some output value $y_i \in \mathcal{V}$.

The *protocol complex* $\mathcal{P}^r$ at the end of round $r \geq 1$ is used for representing all possible system configurations after r rounds of execution. Its vertices (p_i, λ_i) consist of a process name $p_i = \chi((p_i, \lambda_i)) \in \Pi$ that is used as its color, and the local state λ_i of p_i at the end of round r. Since processes can crash in the t-resilient model, the protocol complex may not be pure. A facet σ of the protocol complex consists of $n' \geq n - t$ vertices $v_{\pi(1)}, \ldots, v_{\pi(n')}$, with $v_{\pi(i)} = (p_{\pi(i)}, \lambda_{\pi(i)})$ for $1 \leq i \leq n'$ that represent some possible system configuration after r rounds. We set $\mathcal{P}^0 = \mathcal{I}$.

For any of the simplicial complexes above, $\mathsf{names}(\sigma)$ denotes the set of process names corresponding to the vertices of a face σ, i.e., $\mathsf{names}(\sigma) = \chi(\sigma)$.

In topological modeling, problems like k-set agreement are specified as a *task* $\mathcal{T} = (\mathcal{I}, \mathcal{O}, \Delta)$, where $\Delta : \mathcal{I} \to \mathcal{O}$ is a carrier map that specifies the allowed decision configurations $\Delta(\sigma)$ for a given face $\sigma \in \mathcal{I}$.

▶ **Definition 2.4** (Task solvability). *A task $\mathcal{T}$ is solvable with respect to a protocol complex $\mathcal{P}$ if there exists a simplicial chromatic map $\delta : \mathcal{P} \to \mathcal{O}$ that agrees with Δ, that is, for every $\sigma \in \mathcal{I}$, and for every $\kappa \in \mathcal{P}_\sigma$, $\delta(\kappa) \in \Delta(\sigma)$, where $\mathcal{P}_\sigma$ is the collection of faces of $\mathcal{P}$ reachable when only processes in σ run, with inputs taken from σ.*

Note that, for each face $\sigma \in \mathcal{I}$, $\mathcal{P}_\sigma$ may or may not be empty depending on the executions allowed by the underlying model. For instance, for wait-free computing in the IIS model, $\mathcal{P}_\sigma \neq \varnothing$ for every σ because $|\sigma| \geq 1$ and up to all but one processes can crash. Instead, for synchronous *failure-free* shared-memory computing, $\mathcal{P}_\sigma \neq \varnothing$ if and only if $|\sigma| = n$, whereas for t-resilient synchronous message-passing, $\mathcal{P}_\sigma \neq \varnothing$ if and only if $|\sigma| \geq n - t$.

3 Connectivity-Based Topological Analysis of Synchronous Systems

In this section, we re-visit the round-by-round topological analysis of the lower-bound for deterministic k-set agreement algorithms in the synchronous t-resilient model for the complete graph in [12, Ch. 13]. In a nutshell, this analysis shows that too few rounds of communication lead to a protocol complex that is too highly connected for allowing the existence of a simplicial chromatic map to the output complex of k-set agreement. Since some parts of the original proof are not entirely correct or have been omitted, we revisit these parts in Section 3.2 after presenting some basic ingredients for the proof. In [11, Sec. 3.3], we generalize the analysis for the complete graph to arbitrary communication graphs, and prove formally that the lower bound $\lfloor t/k \rfloor + 1$ for complete graphs also applies to arbitrary communication graphs.

Note that [12, Ch. 13] assumes a system consisting of $n+1$ processes named $\{P_0, \ldots, P_n\}$, with index set $\{0, \ldots, n\}$, whereas our system model considers n processes named $\{p_1, \ldots, p_n\}$, with index set $\{1, \ldots, n\}$. For uniformity, we decided to stick to the latter notation, which makes it necessary to "translate" the original and revised definitions and lemmas of [12, Ch. 13]. In a nutshell, this primarily requires replacing n occurring in the dimension of a face by $n-1$.

3.1 Basic ingredients

In this subsection, we introduce the key ingredients needed for the topological analysis in [12], which tracks the connectivity properties of the sequence of protocol complexes over successive rounds. We start out with the definition of a pseudosphere, a particular type of simplicial complexes. As we shall see, the input complex $\mathcal{I}$ of k-set agreement is a pseudosphere.

▶ **Definition 3.1** (Pseudosphere [12, Def. 13.3.1]). *Let $\varnothing \neq I \subseteq [n]$ be a finite index set. For each $i \in I$, let p_i be a process name, indexed such that if $i \neq j$, then $p_i \neq p_j$, and let V_i be a non-empty set. The* pseudosphere *complex $\Psi(\{(p_i, V_i) \mid i \in I\})$ is defined as follows:*

- *Every pair (p_i, v) with $i \in I$ and $v \in V_i$ is a vertex., and*
- *for every index set $J \subseteq I$, any set $\{(p_j, v_j) \mid j \in J\}$ such that $v_j \in V_j$ for all $j \in J$ is a simplex.*

Note that, for a given simplex σ, and a given set of values $\mathcal{V}$, we sometimes use the shorthand $\Psi(\sigma, \mathcal{V})$ to denote $\Psi(\{(p, \mathcal{V}) \mid p \in \mathsf{names}(\sigma)\})$, where the processes in $\mathsf{names}(\sigma)$ define the corresponding index set I. An essential feature of the protocol complexes arising in the round-by-round connectivity analysis is that they are *shellable*. Intuitively, a pure d-dimensional complex is shellable if it can be built by gluing together its facets in some specific order, called *shelling order*, such that a newly added facet intersects the already built subcomplex in $(d-1)$-dimensional faces only.

▶ **Definition 3.2** (Shellable complex). *A simplicial complex $\mathcal{K}$ is* shellable *if it is pure, of dimension d for some $d \geq 0$, and its facets can be arranged in a linear order $\phi_0, \ldots, \phi_N$, called a* shelling order*, in such a way that, for every $k \in \{1, \ldots, N\}$, the subcomplex $\left(\bigcup_{i=0}^{k-1} \phi_i\right) \cap \phi_k$ is the union of $(d-1)$-dimensional faces of ϕ_k.*

The following alternative definition of shellability is easier to use in proofs.

▶ **Definition 3.3** (Shellability properties [12, Fact 13.1.3]). *An order $\phi_0, \ldots, \phi_N$ of the facets of a pure complex $\mathcal{K}$ is a shelling order if and only if, for any two facets ϕ_a and ϕ_b with $a < b$ in that order, there exists ϕ_c with $c < b$ such that (i) $\phi_a \cap \phi_b \subseteq \phi_c \cap \phi_b$, and (ii) $|\phi_b \setminus \phi_c| = 1$.*

Note that the face ϕ_c guaranteed by Definition 3.3 can depend on ϕ_a. Moreover, ϕ_c is usually not unique, i.e., there might be several choices all satisfying the above properties.

It was shown in [12, Lem. 13.2.2] (resp., [12, Thm. 13.3.6]) that the d-skeleton, for any dimension $d \geq 0$, of any simplex (resp., of any pseudosphere) is shellable. The proofs of these facts are based on the following orderings.

▶ **Definition 3.4** (Face order, adapted from [12, Def. 13.2.1]). *Let $\sigma = \{v_1, \ldots, v_n\}$ be an $(n-1)$-simplex, together with a total ordering on its vertices $v_1, \ldots, v_n$ given by index order. Each face τ of σ has an associated* signature*, denoted by* $\mathrm{sig}(\tau)$*, defined as the Boolean string* $(\mathrm{sig}(\tau)[1], \ldots, \mathrm{sig}(\tau)[n])$ *of length n whose i-th entry is*

$$\mathrm{sig}(\tau)[i] = \begin{cases} \bot & \text{if } v_i \in \tau, \\ \top & \text{if } v_i \notin \tau. \end{cases} \tag{2}$$

The face order $<_f$ of two faces τ_1, τ_2 of σ is defined as $\tau_1 <_f \tau_2$ if $\mathrm{sig}(\tau_1) <_{lex} \mathrm{sig}(\tau_2)$, i.e., $\mathrm{sig}(\tau_1)$ is lexicographically smaller than $\mathrm{sig}(\tau_2)$, where $\bot < \top$.

Note that the empty "simplex" $\varnothing$ is the smallest in the face order of Definition 3.4, and $\mathrm{sig}(\{v_1, \ldots, v_n\})$ is the largest. Definition 3.4 has been slightly adapted from [12, Sec. 13.2.1], by using the more precise notation $\mathrm{sig}(\tau)$ instead of just τ. Informally, $\mathrm{sig}(\tau)$ just encodes the set of processes involved in (the vertices of) a simplex τ. The face order in Definition 3.4 can be used to show that, for any dimension $d \geq 0$, the d-skeleton, of any simplex is shellable.

▶ **Definition 3.5** (Pseudosphere order [12, Def. 13.3.5]). *Let $\phi_a = \{(p_i, \lambda_i) \mid i \in [n]\}$ and $\phi_b = \{(p_i, \mu_i) \mid i \in [n]\}$ be two facets of a pseudosphere $\Psi(\{(p_i, V_i) \mid i \in [n]\})$ where each V_i is an ordered set. The order relation $<_p$ orders these facets lexicographically by value, i.e., $\phi_a <_p \phi_b$ if there exists $\ell \in [n]$ such that $\lambda_i = \mu_i$ for every $1 \leq i < \ell$, and $\lambda_\ell < \mu_\ell$.*

It is known that every skeleton of a shellable complex is shellable [4] (see [11, Thm. 3.6] for the detailed proof):

▶ **Theorem 3.6** (Shellability of skeletons of pseudospheres). *Let $\Psi = \Psi(\{(p_i, V_i) \mid i \in [n]\})$ be an $(n-1)$-dimensional pseudosphere. For every $d \in \{0, \ldots, n-1\}$, the d-skeleton of Ψ is shellable via the shelling order $<$ defined as $\phi_a < \phi_b \iff (\phi_a <_f \phi_b) \vee \big((\mathrm{sig}(\phi_a) = \mathrm{sig}(\phi_b)) \wedge (\phi_a <_p \phi_b)\big)$.*

As our last basic ingredient, we provide a proof of the well-known but often quite informally argued fact that k-set agreement is impossible if the protocol complex is too highly connected (see, e.g., [12, Thm. 10.3.1]). Informally, a complex $\mathcal{K}$ is k-connected, if it does not contain a "hole" of dimension k or lower. For $k = 0$, which captures the consensus impossibility, for example, $\mathcal{K}$ must not be (path-)connected, i.e., 0-connected. More generally, 1-connectivity refers to the ability to contract 1-dimensional loops, 2-connectivity refers to the ability to contract 2-dimensional spheres, etc.

Let $\mathcal{T} = (\mathcal{I}, \mathcal{O}, \Delta)$ be the k-set agreement task as specified in Section 2.3. In particular, $\mathcal{I} = \Psi(\{(p_i, [k+1]) \mid i \in [n]\})$ is a pseudosphere.

▶ **Definition 3.7.** *For every $J \subseteq [k+1]$, we define $\mathcal{P}[J]$ as the minimal complex including $\mathcal{P}_\sigma$ for all $\sigma \in \Psi(\{(p_i, J) \mid i \in [n]\})$.*

▶ **Theorem 3.8.** *If $\mathcal{P}[J]$ is $(\dim(J) - 1)$-connected for all $J \subseteq [k+1]$, then k-set agreement is not solvable with respect to $\mathcal{P}$.*

3.2 The round-by-round connectivity analysis of [12] revisited

During our attempts to generalize the round-by-round topological modeling and analysis of [12, Ch. 13] for the complete graph to arbitrary networks, we figured out that the original analysis in [12] is not entirely correct. More specifically, the analysis there assumes that the involved carrier maps are rigid, which cannot be guaranteed in the executions considered for synchronous k-set agreement where exactly k processes crash per round. We hence provide here a revised analysis for the case of the complete graph (i.e., the clique of our n processes), where rigidity is replaced by a weaker condition (Definition 3.9 below). Unfortunately, this change forces us to re-phrase and re-prove most of the lemmas of the original analysis. Moreover, we have to add a non-trivial strictness proof in Lemma 3.14 below, which was lacking in [12, Ch. 13]. The next definition relaxes Definition 13.4.1 in [12] by removing the rigidity condition.

▶ **Definition 3.9** (q-connected carrier map)**.** *Let $q \geq 0$ be an integer, and let $\mathcal{L}$ and $\mathcal{K}$ be simplicial complexes, where $\mathcal{K}$ is pure. A carrier map $f : \mathcal{K} \to 2^{\mathcal{L}}$ is q-connected if it is strict, and, for every $\sigma \in \mathcal{K}$, the simplicial complex $f(\sigma)$ is $(q - \operatorname{codim}(\sigma))$-connected.*

The following Lemma 3.10 is exactly the same as [12, Lem. 13.4.2]. Indeed, thanks to our new Definition 3.9 of a q-connected carrier map, the original proof holds literally as well, as faces with $\operatorname{codim}(\sigma) > q$ are not considered anyway.

▶ **Lemma 3.10** ([12, Lem. 13.4.2])**.** *For every integer $q \geq 0$, if $\mathcal{K}$ is a pure shellable simplicial complex, and $f : \mathcal{K} \to 2^{\mathcal{L}}$ is a q-connected carrier map, then the simplicial complex $f(\mathcal{K})$ is q-connected.*

In Definition 13.4.3 in [12], a shellable carrier map $f : \mathcal{K} \to 2^{\mathcal{L}}$ was defined as a *rigid* and strict carrier map such that, for each $\sigma \in \mathcal{K}$, the complex $f(\sigma)$ is shellable. Since the carrier maps we study later on are not rigid, we need to weaken this definition as follows.

▶ **Definition 3.11** (q-shellable carrier map)**.** *Let $q \geq 0$ be an integer. A carrier map $f : \mathcal{K} \to 2^{\mathcal{L}}$ is q-shellable if it is strict, and, for each $\sigma \in \mathcal{K}$ satisfying $\operatorname{codim}(\sigma) \leq q + 1$, the complex $f(\sigma)$ is shellable (and hence pure).*

Note carefully that a q-shellable carrier map only guarantees that the image $f(\sigma)$ of a given *single* face $\sigma \in \mathcal{K}$ is shellable. To set the stage for extending the analysis of this section for arbitrary communication graphs later on, we introduce some notation already

here: For a given simplex $\kappa \in \mathcal{K}_{i-1}$ (that will become clear from context), we abbreviate $\overline{\mathcal{K}_i} := f_{i-1}(\kappa) \subseteq \mathcal{K}_i$ and restrict our attention only to (facets of) this subcomplex when proving shellability. Indeed, this will become necessary, since one cannot infer that the entire complex $\mathcal{K}_i$ is also shellable, as we prove in [11, Sec. 3.4].

The following is an appropriately refined version of Lemma 13.4.4 in [12].

▶ **Lemma 3.12.** *Let $q \geq 0$ be an integer, and let us consider a sequence of pure complexes, and carrier maps $\mathcal{K}_0 \xrightarrow{f_0} \mathcal{K}_1 \xrightarrow{f_1} \mathcal{K}_2$, where f_0 is a q-shellable carrier map, f_1 is a q-connected carrier map, and, for every $\sigma \in \mathcal{K}_0$ with $\operatorname{codim}(\sigma) \leq q+1$, $\operatorname{codim}(\sigma) \geq \operatorname{codim}(f_0(\sigma))$. Then, $f_1 \circ f_0$ is a q-connected carrier map.*

Similarly, we need a refined version of Lemma 13.4.5 in [12].

▶ **Lemma 3.13.** *Let $q \geq 0$ and $\ell \geq 0$ be integers, and let us consider a sequence of pure complexes, and carrier maps $\mathcal{K}_0 \xrightarrow{f_0} \mathcal{K}_1 \xrightarrow{f_1} \ldots \xrightarrow{f_\ell} \mathcal{K}_{\ell+1}$, such that the carrier maps $f_0, \ldots, f_{\ell-1}$ are q-shellable, the carrier map f_ℓ is q-connected, and, for every $i \in \{0, \ldots, \ell-1\}$, and every $\sigma \in \mathcal{K}_i$ with $\operatorname{codim}(\sigma) \leq q+1$, $\operatorname{codim}(\sigma) \geq \operatorname{codim}(f_i(\sigma))$. Then, $f_\ell \circ \cdots \circ f_0$ is a q-connected carrier map.*

We want to prove a variant of [12, Thm. 13.5.7] adapted to our refined modeling. For some N to be determined later, consider a sequence

$$\mathcal{K}_0 \xrightarrow{f_0} \mathcal{K}_1 \xrightarrow{f_1} \ldots \xrightarrow{f_{N-1}} \mathcal{K}_N \xrightarrow{id} \mathcal{K}_N \tag{3}$$

where $\mathcal{K}_0$ is the (shellable) input complex for k-set agreement, each $\mathcal{K}_i$ is the image of $\mathcal{K}_{i-1}$ under f_{i-1} (i.e., f_{i-1} is surjective), and

$$f_i(\sigma) = \bigcup_{\tau \in \operatorname{Face}_{n-1-k(i+1)} \sigma} \Psi(\tau, [\tau, \sigma]), \tag{4}$$

where $\Psi(\tau, [\tau, \sigma])$ denotes the pseudosphere obtained by independently labeling the processes in $\mathsf{names}(\tau)$ with one of the simplices in $\{\rho \mid \tau \subseteq \rho \subseteq \sigma\}$. Intuitively, a vertex (p_j, ρ) represents the situation where p_j receives exactly the information in ρ in round $i+1$. As in [12, Sec. 13.5.2], the carrier map f_i applied to σ is the execution map representing round $i+1$ starting from some face $\sigma \in \mathcal{I}$, where the communication graph is a clique, and exactly k additional processes crash in round $i+1$ (i.e., a total of $i \cdot k$ processes have already failed during the i previous rounds).

▶ **Lemma 3.14.** *For every $i \in \{0, \ldots, N-1\}$, f_i is $(k-1)$-shellable, and, for every $\sigma \in \mathcal{K}_i$ with $\operatorname{codim}(\sigma) \leq k$, $\operatorname{codim}(f_i(\sigma)) \leq \operatorname{codim}(\sigma)$.*

Since the identity map id in Equation (3) is trivially a q-connected carrier map satisfying Definition 3.9, we can apply Lemma 3.13 for $N = \lfloor t/k \rfloor$ equal to the maximum number of rounds where k processes can crash to immediately get:

▶ **Lemma 3.15.** *For $N = \lfloor t/k \rfloor$, $f_{N-1} \circ \cdots \circ f_0 : \mathcal{K}_0 \to \mathcal{K}_N$ is a $(k-1)$-connected carrier map.*

▶ **Theorem 3.16.** *For integers $t \geq 0$ and $k \geq 1$, let $\mathcal{P}$ be the protocol complex of k-set agreement after $N = \lfloor \frac{t}{k} \rfloor$ rounds in the synchronous t-resilient model with the complete communication graph, where exactly k processes crash per round. For every $J \subseteq [k+1]$, $\mathcal{P}[J]$ is a $(\dim(J)-1)$-connected subcomplex.*

By combining Theorems 3.8 and 3.16, we finally get the well-known lower bound:

▶ **Corollary 3.17.** *The k-set agreement task cannot be solved in less than $\lfloor \frac{t}{k} \rfloor + 1$ rounds in the synchronous t-resilient model.*

4 A Lower Bound for the Agreement Overhead for Arbitrary Graphs

In this section, we will derive a lower bound for the agreement overhead (Definition 1.1), i.e., the number of additional rounds necessary for solving k-set agreement in t-resilient systems after the first $N = \lfloor t/k \rfloor$ crashing rounds. Interestingly, this can be done via two substantially different approaches, which will be presented below and in [11, Sec. 4.2]. Moreover, as a byproduct of our analyses, we will also establish a lower bound for k-set agreement in systems with t initially dead processes connected by arbitrary communication graphs. For simplicity, we henceforth assume that k evenly divides t, and no further crashes occur after round $N = t/k$. If this is not the case, the missing $t - k\lfloor t/k \rfloor$ crashes could only increase our lower bound, which would further complicate our analysis, and so we discard this option here.

In Definition 4.1 below, we will introduce a novel carrier map g that captures the agreement overhead caused by *arbitrary communication graphs*, beyond the mere case of the clique. This carrier map has been inspired by the *scissors cuts* introduced in [5], which were used to prove a lower bound for solving k-set agreement with oblivious algorithms in the KNOW-ALL model (which is *failure-free*). Our approach however differs from the original scissors cuts in several important ways. First, we admit directed graphs and general full-information algorithms, and replace the original pseudosphere input complex by the *source complex* $\mathcal{P}_N$, which is one of the following two cases:

- $\mathcal{P}_N$ is the (locally shellable) complex $\mathcal{K}_N = \mathcal{P}^{(N)}$ for arbitrary graphs (formally introduced in [11, Sec. 3.3]). This will allow us to paste together the agreement overhead lower bound determined in this section with the $\lfloor t/k \rfloor$ lower bound caused by process crashes.
- $\mathcal{P}_N$ is the (shellable) $(n-t-1)$-skeleton $\mathrm{skel}_{n-t-1}\big(\Psi(\{(p_i, [k+1]) \mid i \in [n]\})\big)$ of the pseudosphere given in Theorem 3.6, which we subsequently abbreviate as $\Psi(n, k+1)$ for conciseness, which contains all the faces of the full pseudosphere $\Psi(\{(p_i, [k+1]) \mid i \in [n]\})$ with at most $n-t$ vertices. Note that actually $N = 0$ in this case, albeit we will not make this explicit later on, but just stick to $\mathcal{P}_N$ to denote the source complex for uniformity. This will result in a lower bound (Theorem 4.11) for k-set agreement with t initially dead processes. Note that the total number of processes that appear in *all* the facets of $\Psi(n, k+1)$ together is n here.

The second main difference w.r.t. [5] is that g will be a proper carrier map, which also specifies the images of arbitrary faces of the source complex, and not only images of facets. It is particularly simple, however, since it resembles a (non-rigid) simplicial map in that $g(\sigma)$ returns the subcomplex corresponding to a *single* simplex ρ, or else $g(\sigma) = \varnothing$. In order not to unnecessarily clutter our notation, we will hence subsequently pretend as if $g(\sigma)$ only consisted of a single simplex ρ or $\varnothing$ only. One of the particularly appealing consequences of g's simple image is that it allows us to replace the very complex topological analysis in [5] by a strikingly simple connectivity argument.

Generally, we assume that $g : \mathcal{P}_N \to \mathcal{P}_M$ given in Definition 4.1 models the failure-free execution in rounds $N+1, N+2, \ldots, M$ for some $M > N$, where $M-N$ will finally determine our desired agreement overhead lower bound. Consider the execution starting in some facet $\phi \in \mathcal{P}_N$, which involves exactly $n-t$ processes with index set $I_\phi \subseteq [n]$ and consists of vertices of the form (p_i, λ_i), $i \in I_\phi$, each with a process name $p_i = \mathsf{names}((p_i, \lambda_i)) \in \Pi$ used as its color, and a label λ_i (which denotes p_i's local view at the end of round N), consisting of

- the vertices of the processes that managed to successfully send to p_i up to round N, if $N > 0$, or
- p_i's initial value x_i if $N = 0$, i.e., when the source complex is $\mathrm{skel}_{n-t-1}\big(\Psi(n, k+1)\big)$.

In either case, λ_i encodes the complete heard-of history of p_i up to round N, due to the fact that we are assuming full-information protocols.

We assume that $G_{N+1}, \ldots, G_M$ is the sequence of communication graphs governing rounds $N+1, \ldots, M$. These graphs may be different and known to the processes; clearly, assuming a static graph $G = G_{N+1} = \ldots G_M$ as in Section 1 can make our lower bound only stronger. We will abbreviate this sequence by G for brevity, and define G_ϕ to be the product $G_{N+1,\phi} \circ G_{N+2,\phi} \circ \cdots \circ G_{M,\phi}$, where $G_{N+1,\phi}, \ldots, G_{M,\phi}$ are the graphs induced by the nodes $\mathsf{names}(\phi)$ on the graphs in the sequence G.

It is worth mentioning that our carrier map g actually focuses on a subset of all the possible executions, as it is sufficient for a lower bound. Namely, g considers the case where processes that crashed in round N did crash *cleanly* only (they failed to send messages to all their neighbors). Note that this somehow resembles the situation of the carrier maps f_i used in Section 3, which also only covered a submodel of all possible executions, namely, the one where exactly k processes crash per round. The map g accomplishes this by "discarding" executions starting from facets in $\mathcal{P}_N$ that involve unclean crashes in round N, in the sense that it (non-rigidly) maps such an "unclean" facet ϕ' to some face in the image $g(\phi)$ of some "clean" facet $\phi \in \mathcal{P}_N$, where the processes that crashed uncleanly in ϕ' crashed cleanly in ϕ or not at all.

In more detail, g maps a facet ϕ' to the maximal face ρ contained in the "full", i.e., unconstrained, image of ϕ' (that would be used without the discarding of "dirty" source vertices), where *no* vertex hears from a "dirty" witness of an unclean crash in ϕ'. Note that any such $\rho \neq \varnothing$ is also present as part of $g(\phi)$ for some facet ϕ where the crashing processes are not participating at all or are correct, so no "new" face needs to be included for mapping $g(\phi')$ here. On the other hand, there is no a priori guarantee that such a face ρ exists, as $g(\phi') = \varnothing$ is also possible; we will show in our non-emptyness proofs below (see Lemma 4.6 and Lemma 4.9) that this cannot happen under the conditions of our impossibility proof, however. In fact, this very intuitive property of g is what enables g to completely replace the complicated analysis of [5] by a simple connectivity argument based on its non-emptyness.

To formally define g, we need the following notation.

- For a facet $\phi = \{(p_i, \lambda_i^N) \mid i \in I_\phi\}$ of $\mathcal{P}_N$, where $I_\phi \subseteq [n]$ and $|I_\phi| = n - t$, let

$$\breve{\phi} = \{(p_i, \lambda_i^N) \mid i \in I_\phi \wedge \mathsf{names}(\lambda_i^N) \setminus \mathsf{names}(\phi) \neq \varnothing\}$$

  be the set of "dirty vertices" in ϕ. That is, $\breve{\phi} \subseteq \phi$ is the set of vertices in ϕ where the corresponding processes received a message from a process that has crashed uncleanly in round N.
- For a vertex $(p_i, \lambda_i^M) \in \mathcal{P}_M$, let $\mathrm{hist}(\lambda_i^M)$ be the set of all vertices (p_j, λ_j^N) contained in the heard-of history λ_i^M, i.e., the ones have been received by p_i directly or indirectly in any of the rounds $N+1, \ldots, M$.

▶ **Definition 4.1** (Agreement overhead carrier map)**.** *We define the carrier map g as follows:*

- *For a facet $\phi \in \mathcal{P}_N$,*

$$g(\phi) = \big\{\{(p_i, \lambda_i^M) \mid i \in I_\phi \wedge \mathrm{hist}(\lambda_i^M) \cap \breve{\phi} = \varnothing\}\big\}. \tag{5}$$

  That is, $g(\phi)$ is the (possibly empty) face of $\mathcal{P}_M$ consisting of the vertices that do not have any vertex in $\breve{\phi}$ in their heard of history.
- *For a face $\sigma \in \mathcal{P}_N$,*

$$g(\sigma) = \{\rho\} \text{ with } \rho = \text{maximal simplex in } \textstyle\bigcap_{\phi \in T_\sigma} g(\phi) \text{ s.th. } \mathsf{names}(\rho) \subseteq \mathsf{names}(\sigma), \tag{6}$$

  where T_σ denotes the set of all facets $\phi \in \mathcal{P}_N$ satisfying $\sigma \subseteq \phi$.

Note that Equation (5) and Equation (6) are consistent, in the sense that the image $g(\sigma)$ of a facet $\sigma \in \mathcal{P}_N$ is the same for both definitions, since $T_\sigma = \{\sigma\}$ here.

The facets of our source complex $\mathcal{P}_N$ need to satisfy the following additional conditions:

▶ **Definition 4.2.** *We define conditions* ***C1*** *and* ***C2*** *as follows:*
C1: *There is a facet $\phi \in \mathcal{P}_N$ with* $\mathsf{names}(\phi) = \Pi \setminus S$*, for every subset S of t processes.*
C2: *For every $\sigma \in \mathcal{P}_N$, $\sigma = \bigcap_{\phi \in T_\sigma} \phi$.*

Condition **C1** is satisfied for our source complexes, since both $\mathcal{P}_N = \mathrm{skel}_{n-t-1}\big(\Psi(n,k+1)\big)$ and $\mathcal{P}_N = \mathcal{K}_N$ contain every possible $(n-t-1)$-face by definition/construction.

The special context in which strictness of g is actually utilized, namely, [12, Lem. 13.4.2] (see Section 3.2 for details) is restricted to faces which are solely facet intersections, which actually makes **C2** in Definition 4.2 superfluous. However, since it is guaranteed for our source complexes, we can safely require it. Indeed, the regularity of the source complex $\mathcal{P}_N = \mathrm{skel}_{n-t-1}\big(\Psi(n,k+1)\big)$ trivially guarantees (C2), and for the source complex $\mathcal{P}_N = \mathcal{K}_N$, condition (C2) is easy to prove since every k-subset of the alive processes in $\mathcal{K}_{N-1}$ is crashed in round N in order to produce some facet $\phi \in \mathcal{P}_N$. So if $v = (p, \lambda_q) \in \bigcap_{\phi \in T_\sigma} \phi \setminus \sigma$ would exist, consider any facet $\phi \in T_\sigma$ where some process $q \notin \mathsf{names}(\phi)$ has crashed uncleanly after successfully sending his λ_q to everybody in round N in ϕ. There must also be a facet $\phi' \in \mathcal{P}_N$, which is the same as ϕ, except that $(q, \lambda_q) \in \phi'$ but $v \notin \phi'$ since p has crashed uncleanly after successfully sending λ_p to everybody in round N in ϕ'. Since $\sigma \subseteq \phi$ and hence $\phi' \in T_\sigma$ as well, we get the desired contradiction.

▷ Claim 4.3. g is a carrier map.

Proof. Let σ_1 and σ_2 be two faces of $\mathcal{P}_N$ with $\sigma_1 \subseteq \sigma_2$. We have $T_{\sigma_2} \subseteq T_{\sigma_1}$, and $\mathsf{names}(\sigma_1) \subseteq \mathsf{names}(\sigma_2)$, from which it follows that $g(\sigma_1) \subseteq g(\sigma_2)$. ◁

In Claim 4.4 below, we will prove that g is also strict. Our proof will rely on an essential property of g, namely, that if two facets ϕ_1 and ϕ_2 share a vertex $y = (p, \lambda_p) \in \phi_1 \cap \phi_2$, then the information p propagates to other common vertices in $\phi_1 \cap \phi_2$ in rounds $N+1, \ldots, M$ is the same in G_{ϕ_1} and G_{ϕ_2}. Recall that the graph G_ϕ is the product $G_{N+1,\phi} \circ G_{N+2,\phi} \circ \cdots \circ G_{M,\phi}$, where $G_{N+1,\phi}, \ldots, G_{M,\phi}$ are the graphs induced by $\mathsf{names}(\phi)$ in the sequence of graphs $G_{N+1}, G_{N+2}, \ldots, G_M$ that govern rounds $N+1, \ldots, M$. Whereas this is trivially satisfied when the source complex is $\mathrm{skel}_{n-t-1}\big(\Psi(n,k+1)\big)$, as there are no unclean crashes since all absent processes are initially dead, it needs to be secured by "discarding" vertices in $\breve{\phi}_1$ (resp., $\breve{\phi}_2$) in Equation (5) when unclean crashes may have happened. Indeed, p could have a predecessor q that sent its value λ_q to p in round N because $q \in \mathsf{names}(\phi_1)$ in ϕ_1, whereas it did not so in ϕ_2 because it crashed uncleanly, so that $q \notin \mathsf{names}(\phi_2)$.

▷ Claim 4.4. g is a strict carrier map.

Proof. Let ϕ_1, ϕ_2 be facets of $\mathcal{P}_N$, and let $\sigma = \phi_1 \cap \phi_2$. We show that $g(\phi_1) \cap g(\phi_2) = g(\sigma)$ by proving that $g(\phi_1) \cap g(\phi_2) \subseteq g(\sigma)$. Assume by contradiction that there exists a vertex $x \in g(\phi_1) \cap g(\phi_2)$, but $x \notin g(\sigma)$. Since $\sigma = \phi_1 \cap \phi_2$, and thanks to the definition of T_σ, $\sigma = \cap_{\phi \in T_\sigma} \phi$ must obviously hold. By the definition of g, we then get $g(\sigma) = \cap_{\phi \in T_\sigma} g(\phi)$. Hence, there must be some $\phi_3 \in T_\sigma$ such that $x \notin g(\phi_3)$, as well as a unique $x' \in \sigma$ with $\mathsf{names}(x') = \mathsf{names}(x)$.

Since $x \in g(\phi_1) \cap g(\phi_2)$, node $\mathsf{names}(x)$ can only hear from nodes in $\mathsf{names}(\sigma)$ in the subgraphs G_{ϕ_1} and G_{ϕ_2} of G induced by $\mathsf{names}(\phi_1)$ and $\mathsf{names}(\phi_2)$, respectively. On the other hand, $\mathsf{names}(x)$ must hear from a process outside $\mathsf{names}(\sigma)$ in G_{ϕ_3}. Indeed, assume for

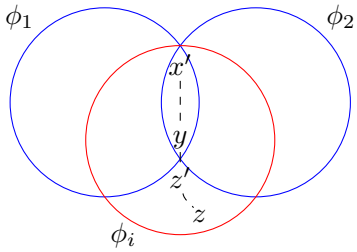

Figure 1 Illustration of the strictness proof in Claim 4.4.

a contradiction that this is not the case. Then, in any of the graphs $G_{\phi_1}, G_{\phi_2}, G_{\phi_3}$, node $\mathsf{names}(x)$ hears only from nodes in $\mathsf{names}(\sigma)$, and all nodes in $\mathsf{names}(\sigma)$ are of course included in $G_{\phi_1}, G_{\phi_2}, G_{\phi_3}$. Node $\mathsf{names}(x)$ hence has the same heard-of history in $G_{\phi_1}, G_{\phi_2}, G_{\phi_3}$, contradicting the assumption $x \notin g(\phi_3)$. Therefore, there must exist some $z \in \phi_3$ with $z \notin \phi_1$ and $z \notin \phi_2$ such that $\mathsf{names}(x)$ hears from $\mathsf{names}(z)$ in G_{ϕ_3} via some path P in round $N+1, \dots, M$ (see Figure 1 for an illustration). Let $y \in \sigma$ be such that

- $\mathsf{names}(y)$ is a node in P that is the closest to $\mathsf{names}(x') = \mathsf{names}(x)$ in G_{ϕ_3}, and
- $\mathsf{names}(y)$ has a neighbor $\mathsf{names}(z') \in P$ and $z' \notin \sigma$.

We must have $z' \notin \phi_1$. Indeed, if $z' \in \phi_1$, then the path suffix $P' \subseteq P$ leading from $\mathsf{names}(z') \to \mathsf{names}(y) \to \mathsf{names}(x)$ would be in G_{ϕ_1}, so $\mathsf{names}(x)$ would hear from $\mathsf{names}(z') \notin \mathsf{names}(\sigma)$ in G_{ϕ_1} in rounds $N+1, \dots, M$, which contradicts that it can only hear from nodes in $\mathsf{names}(\sigma)$ as established above. Analogously, $z' \notin \phi_2$ must hold. Note carefully, however, that the path suffix $P'' \subseteq P' \subseteq P$ leading from $\mathsf{names}(y) \to \mathsf{names}(x)$ is contained in any of $G_{\phi_1}, G_{\phi_2}, G_{\phi_3}$.

Since $z' \notin \phi_1 \cup \phi_2$ but $y \in \phi_1 \cap \phi_2 \cap \phi_3$, process $\mathsf{names}(z')$ must have crashed uncleanly in round N after sending to $\mathsf{names}(y)$ in both ϕ_1 and ϕ_2 (note that it is here where we need condition **C1** in Definition 4.2). However, in that case, Equation (5) would guarantee that $\mathsf{names}(x) \notin g(\phi_1) \cup g(\phi_2)$, which contradicts our initial assumption $x \in g(\phi_1) \cap g(\phi_2)$.

By monotonicity of g, it follows that $g(\phi_1) \cap g(\phi_2) = g(\phi_1 \cap \phi_2)$ for every two facets ϕ_1, ϕ_2 of $\mathcal{P}_N$.

We also need to prove strictness for faces, so let σ_1, σ_2 be two faces of $\mathcal{P}_N$ with $\sigma = \sigma_1 \cap \sigma_2$, and assume for a contradiction that there is some $x \in g(\sigma_1) \cap g(\sigma_2)$ but $x \notin g(\sigma)$. According to Equation (6), this implies that $x \in \phi_1 \cap \phi_2$ for any two facets $\phi_1 \in T_{\sigma_1}$ and $\phi_2 \in T_{\sigma_2}$, but that there is some facet $\phi_3 \in T_\sigma$ with $x \notin g(\phi_3)$. If we pick any such ϕ_1 and ϕ_2, we might observe $\phi' = \phi_1 \cap \phi_2 \supset \phi$, but still $x \in g(\phi_1) \cap g(\phi_2)$ but $x \notin g(\phi_3)$. It is easy to see, in particular, from Figure 1, that the above contradiction proof applies also here, since its arguments are not affected by assuming $\phi' \supset \phi$. ◁

We will now utilize our carrier map g for establishing our desired agreement overhead lower bound. As an appetizer, we will first provide a simple-to-prove eccentricity-based definition of a graph radius, which requires a static communication graph, i.e., $G = G_{N+1} = G_{G+2} = \dots = G_M$ in rounds $N+1, \dots, M$. Note that here, G_ϕ is equal to the $(M-N)$-th power of the subgraph of G induced by the processes present in ϕ.

▶ **Definition 4.5.** *For a node set $D \subseteq V$, the* eccentricity *of D in the graph $G = (V, E)$, denoted* $\mathrm{ecc}(D, G)$*, is the smallest integer d such that for every node $v \in V$ there is a path from some node $u \in D$ to v in G consisting of at most d hops.*
If G is a dynamic graph, $\mathrm{ecc}(D, G)$ is similarly defined, but with a temporal path from u to v; that is, if all nodes in D broadcast the same message in G by flooding, then all nodes in V receive the message in at most d rounds.

By this definition, if all nodes in D broadcast for $\mathrm{ecc}(D,G)$ rounds in the distributed message-passing model, every node in $V(G)$ hears from at least one node in D. By considering the set D that minimizes the eccentricity, we define a corresponding *Radius* $\mathrm{Rad}(G,t+k) = \min_{D\subseteq V,|D|=t+k} \mathrm{ecc}(D,G)$. This allows us to state the following essential property of the corresponding carrier map g:

▶ **Lemma 4.6.** *Let $R = M - N$. For a static communication graph G, if $R < \mathrm{Rad}(G,t+k)$, then $g : \mathcal{P}_N \to \mathcal{P}_M$ is a $(k-1)$-connected carrier map.*

▶ **Theorem 4.7.** *For every graph G, $t \geq 0$ and $k \geq 1$, there are no algorithms solving k-set agreement in the t-resilient model in G in less than $R = \lfloor \frac{t}{k} \rfloor + \mathrm{Rad}(G,t+k)$ rounds.*

Now we will finally turn to our ultimately desired lower bound, which essentially follows from the lower bound given in Theorem 4.7, by replacing $\mathrm{Rad}(G,t+k)$ with the (t,k)-radius $\mathrm{rad}(G,t,k)$ defined as follows:

▶ **Definition 4.8** ((t,k)-radius of a graph sequence G)**.** *For an n-node graph sequence $G = G_{N+1},\ldots,G_M$ and any two integers t,k with $t \geq 0$ and $k \geq 1$, we define the (t,k)-radius $\mathrm{rad}(G,t,k)$ as follows:*

$$\mathrm{rad}(G,t,k) = \min_{D,|D|=t+k} \max_{D'\subseteq D,|D'|=t} \mathrm{ecc}(D\setminus D', G\setminus D'). \tag{7}$$

Recall that $\mathrm{ecc}(D\setminus D', G\setminus D')$ is the number of rounds needed for $D\setminus D'$ to collectively broadcast in the subgraph sequence of G induced by $\Pi\setminus D'$.

Note carefully that this definition generalizes the definition of the (t,k)-radius of a static graph already given in Equation (1) to our graph sequences.

▶ **Lemma 4.9.** *Let $R = M-N$. If $R < \mathrm{rad}(G,t,k)$, then $g : \mathcal{P}_N \to \mathcal{P}_M$ is a $(k-1)$-connected carrier map.*

Exactly the same proof as for Theorem 4.7 thus yields the refined lower bound stated in the following theorem:

▶ **Theorem 4.10.** *For every graph G, $t \geq 0$ and $k \geq 1$, there are no algorithms solving k-set agreement in G in the t-resilient model in strictly less than $R = \lfloor \frac{t}{k} \rfloor + \mathrm{rad}(G,t,k)$ rounds.*

Our analysis also provides a lower bound for systems with t initially dead processes, by starting from the source complex $\mathcal{P}_N = \mathrm{skel}_{n-t-1}\big(\Psi(\{(p_i,[k+1]) \mid i\in[n]\})\big)$, which is shellable according to Theorem 3.6. Analogous to Theorem 4.10, this concludes in the following theorem.

▶ **Theorem 4.11.** *For every graph G, $t \geq 0$ and $k \geq 1$, there are no algorithms solving k-set agreement in G with t initially dead processes in less than $\mathrm{rad}(G,t,k)$ rounds.*

Note that we establish in [11] that Theorem 4.11 for $t = 0$ (almost) coincides with the dominance number based lower bound for k-set agreement in the KNOW-ALL model established in [5].

5 Upper Bound for Fixed Graphs

Let $G = (V,E)$ be an n-node graph with vertex connectivity $\kappa(G)$. Let $t < \kappa(G)$ be a non-negative integer, and let $k \geq 1$ be an integer. We are interested in solving k-set agreement in G with at most t crash failures. For $S \subseteq V$, let $G\setminus S$ denote the subgraph of G induced by

the nodes in $V \setminus S$, i.e., $G \setminus S$ is an abbreviation for $G[V \setminus S]$. For every graph H, let $D(H)$ denote its diameter. We define $D(G,t) = \max_{S \subseteq V, |S| \leq t} D(G \setminus S)$. Note that since $t < \kappa(G)$, and the maximization is over all sets S of size at most t, $D(G,t)$ is finite.

▶ **Theorem 5.1.** *There exists an algorithm solving k-set agreement in G in $\lfloor \frac{t}{k} \rfloor + D(G,t)$ rounds.*

Note that the bound in Theorem 5.1 matches the bound $\lfloor \frac{t}{k} \rfloor + 1$ rounds for k-set agreement in the n-node clique K_n under the synchronous t-resilient model (see [8]), as $D(K_n, t) = 1$.

Examples

- Let us consider the n-node cycle, i.e., $G = C_n$, with $t = 1$, and $k = 1$ (i.e., consensus). We have $D(C_n, 1) = n - 2$, as, for every node v, $C_n \setminus \{v\}$ is a path with $n - 1$ nodes. The algorithm of Theorem 5.1 must thus perform min-flooding for $1 + (n - 2) = n - 1$ rounds to solve consensus in C_n. Intuitively, this appears to be the best that can be achieved as the node with the smallest input value may crash at the first round, by sending its value to just one of its neighbors, and then $n - 2$ additional rounds will be needed for this value to reach all nodes.
- Let us consider the d-dimensional hypercube Q_d, $d \geq 1$, with $n = 2^d$ nodes. We have $\kappa(Q_d) = d$, and there are d internally-disjoint paths of length at most $d + 1$ between any two nodes, which implies that $D(Q_d, d - 1) = d + 1$. The algorithm of Theorem 5.1 must thus perform min-flooding for $\lfloor \frac{t}{k} \rfloor + (d + 1)$ rounds to solve k-set agreement in the t-resilient hypercube Q_d.

6 Conclusions

We provided novel lower bounds for k-set agreement in synchronous t-resilient systems connected by an arbitrary directed communication network. Our lower bound combines the $\lfloor t/k \rfloor$ lower bound (which we generalized to arbitrary communication graphs) obtained for rounds where exactly t processes crash with an additional novel lower bound on the agreement overhead caused by an arbitrary network, i.e., different from the complete graph. Our results use the machinery of combinatorial topology for studying the (high) connectivity properties of the round-by-round protocol complexes obtained by some novel and strikingly simple carrier maps, which we firmly believe to have applications also in other contexts. Whereas we also provided some upper bound result, the challenging question of possible tightness is deferred to future research.

References

1 Dan Alistarh, James Aspnes, Faith Ellen, Rati Gelashvili, and Leqi Zhu. Why extension-based proofs fail. *SIAM Journal on Computing*, 52(4):913–944, 2023. doi:10.1137/20M1375851.

2 Hagit Attiya, Armando Castañeda, and Sergio Rajsbaum. Locally solvable tasks and the limitations of valency arguments. *J. Parallel Distributed Comput.*, 176:28–40, 2023. doi:10.1016/J.JPDC.2023.02.002.

3 Hagit Attiya, Pierre Fraigniaud, Ami Paz, and Sergio Rajsbaum. On the existence of extension-based proofs of impossibility for set-agreement. In *32nd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, volume 15671 of *LNCS*, pages 56–73. Springer, 2025. doi:10.1007/978-3-031-91736-3_4.

4 Anders Björner and Michelle L. Wachs. Shellable nonpure complexes and posets. i. *Transactions of the American Mathematical Society*, 348(4):1299–1327, 1996. URL: http://www.jstor.org/stable/2155143.

5 Armando Castañeda, Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum, Matthieu Roy, and Corentin Travers. A topological perspective on distributed network algorithms. *Theoretical Computer Science*, 849:121–137, 2021. doi:10.1016/J.TCS.2020.10.012.

6 Armando Castañeda, Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum, Matthieu Roy, and Corentin Travers. Synchronous t-resilient consensus in arbitrary graphs. *Inf. Comput.*, 292:105035, 2023. doi:10.1016/J.IC.2023.105035.

7 Soma Chaudhuri. More *choices* allow more *faults*: Set consensus problems in totally asynchronous systems. *Information and Control*, 105(1):132–158, July 1993. doi:10.1006/INCO.1993.1043.

8 Soma Chaudhuri, Maurice Herlihy, Nancy A. Lynch, and Mark R. Tuttle. Tight bounds for k-set agreement. *J. ACM*, 47(5):912–943, 2000. doi:10.1145/355483.355489.

9 Pierre Fraigniaud, Minh Hang Nguyen, and Ami Paz. Agreement tasks in fault-prone synchronous networks of arbitrary structure. In *42nd International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 327 of *LIPIcs*, pages 34:1–34:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICS.STACS.2025.34.

10 Pierre Fraigniaud, Minh-Hang Nguyen, and Ami Paz. A simple lower bound for set agreement in dynamic networks. In *8th Symposium on Simplicity in Algorithms (SOSA)*, pages 253–262. SIAM, 2025. doi:10.1137/1.9781611978315.20.

11 Pierre Fraigniaud, Minh Hang Nguyen, Ami Paz, Ulrich Schmid, and Hugo Rincon Galeana. Lower bounds for k-set agreement in fault-prone networks, 2025. arXiv:2508.15562.

12 Maurice Herlihy, Dmitry Kozlov, and Sergio Rajsbaum. *Distributed computing through combinatorial topology*. Morgan Kaufmann, 2013.

A Proofs for Section 3 (Connectivity-Based Topological Analysis of Synchronous Systems)

▶ **Theorem 3.8.** *If $\mathcal{P}[J]$ is $(\dim(J)-1)$-connected for all $J \subseteq [k+1]$, then k-set agreement is not solvable with respect to $\mathcal{P}$.*

Proof. Let us assume, for the purpose of contradiction, that the k-set agreement task $\mathcal{T}$ is solvable with respect to $\mathcal{P}$. This implies there exists a simplicial map $\delta : \mathcal{P} \to \mathcal{O}$ that agrees with Δ. Let $\mathcal{K} = [k+1]$ viewed as a complex, and let $\Theta : \mathcal{K} \to 2^{\mathcal{P}}$ be defined as $\Theta(J) = \mathcal{P}[J]$ for every $J \subseteq [k+1]$ viewed as a simplex.

We claim that Θ is a carrier map. Indeed, if $J' \subseteq J$, then $\Psi(\{(p_i, J') \mid i \in [n]\}) \subseteq \Psi(\{(p_i, J) \mid i \in [n]\})$. That is, every $\sigma \in \Psi(\{(p_i, J') \mid i \in [n]\})$ belongs to $\Psi(\{(p_i, J) \mid i \in [n]\})$, which implies $\mathcal{P}[J'] \subseteq \mathcal{P}[J]$.

Thanks to Theorem 3.7.7(2) in [12], since Θ is a carrier map, and since, for each J, $\mathcal{P}[J]$ is $(\dim(J)-1)$-connected, we get that Θ has a simplicial approximation $(\mathrm{Div}(\mathcal{K}), g)$. That is:

- $\mathrm{Div}(\mathcal{K})$ is a chromatic subdivision of $\mathcal{K}$,
- $g : \mathrm{Div}(\mathcal{K}) \to \mathcal{P}$ is simplicial and chromatic, and
- for every $J \subseteq [k+1]$, and every $\rho \in \mathrm{Div}(J)$, $g(\rho) \in \Theta(J)$, where $\mathrm{Div}(J)$ is the subdivision of the face J of $\mathcal{K}$ induced by the global subdivision $\mathrm{Div}(\mathcal{K})$.

Let $f : \mathcal{P} \to \partial\mathcal{K}$ be defined as $f = \mathrm{val} \circ \delta$, where val is the trivial mapping that discards process IDs, and where $\partial\mathcal{K}$ is the boundary of $\mathcal{K}$. As the combination of two simplicial maps, we get that f is simplicial. Let $h : \mathrm{Div}(\mathcal{K}) \to \partial\mathcal{K}$ be defined as $h = f \circ g$.

We claim that h is a Sperner coloring of $\mathrm{Div}(\mathcal{K})$. For $J \subseteq [k+1]$ and $\rho \in \mathrm{Div}(J)$, $h(\rho) = f \circ g(\rho)$. Since $g(\rho) \in \Theta(J) = \mathcal{P}[J]$, there exists $\sigma \in \Psi(\{p_i, J) \mid i \in [n]\})$ such that $g(\rho) \in \mathcal{P}_\sigma$. The validity condition implies when processes in σ run alone, each output must be in $\mathrm{val}(\sigma) \subseteq J$, that is, for every $\tau \in \mathcal{P}_\sigma$, $\mathrm{val}(\delta(\tau)) \subseteq J$. Therefore $\mathrm{val} \circ \delta(g(\rho)) \subseteq J$.

Since there are no Sperner colorings of $\mathrm{Div}(\mathcal{K})$ that can avoid facets that are colored with $k+1$ colors, we get a contradiction to the assumption that $\mathcal{T}$ is solvable w.r.t. $\mathcal{P}$. ◀

▶ **Lemma 3.12.** *Let $q \geq 0$ be an integer, and let us consider a sequence of pure complexes, and carrier maps $\mathcal{K}_0 \xrightarrow{f_0} \mathcal{K}_1 \xrightarrow{f_1} \mathcal{K}_2$, where f_0 is a q-shellable carrier map, f_1 is a q-connected carrier map, and, for every $\sigma \in \mathcal{K}_0$ with $\mathrm{codim}(\sigma) \leq q+1$, $\mathrm{codim}(\sigma) \geq \mathrm{codim}(f_0(\sigma))$. Then, $f_1 \circ f_0$ is a q-connected carrier map.*

Proof. First, $g = f_1 \circ f_0$ is a strict carrier map because it is a composition of two strict carrier maps. It remains to check that $g(\sigma)$ is $(q - \mathrm{codim}\,\sigma)$-connected. Considering an arbitrary $\sigma \in \mathcal{K}_0$ satisfying $\mathrm{codim}(\sigma) \leq q+1$, we have:

- **(i)** $f_0(\sigma)$ is shellable and pure,
- **(ii)** $\mathrm{codim}(f_0(\sigma)) \leq \mathrm{codim}(\sigma)$,
- **(iii)** for every simplex $\tau \in f_0(\sigma)$, the co-dimension of τ in $f_0(\sigma)$, denoted by $\mathrm{codim}(\tau, f_0(\sigma))$, satisfies $\mathrm{codim}(\tau, f_0(\sigma)) = \dim(f_0(\sigma)) - \dim(\tau)$, and
- **(iv)** $\mathrm{codim}(\tau, \mathcal{K}_1) = \mathrm{codim}(\tau, f_0(\sigma)) + \mathrm{codim}(f_0(\sigma), \mathcal{K}_1)$.

Let $q' = q - \mathrm{codim}(f_0(\sigma))$. Since f_1 is a q-connected carrier map, $f_1(\tau)$ is $(q - \mathrm{codim}(\tau, \mathcal{K}_1))$-connected. Equivalently, $f_1(\tau)$ is $(q' - \mathrm{codim}(\tau, f_0(\sigma)))$-connected. By applying Lemma 3.10, $f_1(f_0(\sigma))$ is q' connected. Thus, $g(\sigma)$ is $(q - \mathrm{codim}\,\sigma)$-connected since $q' \leq q - \mathrm{codim}(\sigma)$ by (ii). ◀

▶ **Lemma 3.13.** *Let $q \geq 0$ and $\ell \geq 0$ be integers, and let us consider a sequence of pure complexes, and carrier maps $\mathcal{K}_0 \xrightarrow{f_0} \mathcal{K}_1 \xrightarrow{f_1} \ldots \xrightarrow{f_\ell} \mathcal{K}_{\ell+1}$, such that the carrier maps $f_0, \ldots, f_{\ell-1}$ are q-shellable, the carrier map f_ℓ is q-connected, and, for every $i \in \{0, \ldots, \ell-1\}$, and every $\sigma \in \mathcal{K}_i$ with $\mathrm{codim}(\sigma) \leq q+1$, $\mathrm{codim}(\sigma) \geq \mathrm{codim}(f_i(\sigma))$. Then, $f_\ell \circ \cdots \circ f_0$ is a q-connected carrier map.*

Proof. We use induction on $k \geq 0$ to prove that $g_k = f_\ell \circ \cdots \circ f_{\ell-k}$ is a q-connected carrier map. Note that $g_\ell = f_\ell \circ \cdots \circ f_0$. For the base case $k = 0$, the claim is immediate from our assumption on f_ℓ. For the induction step from k to $k+1$, we note that $g_{k+1} = f_\ell \circ \cdots \circ f_{\ell-k-1} = g_k \circ f_{\ell-k-1}$. Since $f_{\ell-k-1}$ is q-shellable and guarantees $\mathrm{codim}(\sigma) \geq \mathrm{codim}(f_{\ell-k-1}(\sigma))$ for all $\sigma \in \mathcal{K}_{\ell-k}$ with $\mathrm{codim}(\sigma) \leq q+1$ by our assumptions, and since g_k is q-connected by the induction hypothesis, we can apply Lemma 3.12, which ensures that $g_{k+1} = g_k \circ f_{\ell-k-1}$ is q-connected as needed. ◀

▶ **Lemma 3.14.** *For every $i \in \{0, \ldots, N-1\}$, f_i is $(k-1)$-shellable, and, for every $\sigma \in \mathcal{K}_i$ with $\mathrm{codim}(\sigma) \leq k$, $\mathrm{codim}(f_i(\sigma)) \leq \mathrm{codim}(\sigma)$.*

Proof. Lemma 13.5.5 in [12] shows that one can define a shelling order on the facets of $f_i(\sigma)$, which is a pure complex by Equation (4), for every $\sigma \in \mathcal{K}_i$ that yields $f_i(\sigma) \neq \varnothing$. All that remains to be proved is hence strictness, and the additional condition $\mathrm{codim}\, f_i(\sigma) \leq \mathrm{codim}\,\sigma$ for every $\sigma \in \mathcal{K}_i$ with $\mathrm{codim}\,\sigma \leq k$.

For the latter, note that, for every $\sigma \in \mathcal{K}_i$, if $\dim(\sigma) < n - k(i+1)$, then $f_i(\sigma) = \varnothing$. Since Equation (4) implies that $\dim(\mathcal{K}_{i+1}) = \dim(\mathcal{K}_i) - k$, we can indeed guarantee $\mathrm{codim}(f_i(\sigma)) \leq \mathrm{codim}(\sigma)$ for every simplex σ in $\mathcal{K}_i$ satisfying $\mathrm{codim}(\sigma) \leq k$.

For strictness, let ϕ_1, ϕ_2 be simplices of $\mathcal{K}_i$, and let $\phi = \phi_1 \cap \phi_2$. We prove that $f_i(\phi) = f_i(\phi_1) \cap f_i(\phi_2)$. We have

$$f_i(\phi_1) = \bigcup_{\tau \in \mathrm{Face}_{n-1-k(i+1)}\, \phi_1} \Psi(\tau, [\tau, \phi_1]), \quad \text{and} \quad f_i(\phi_2) = \bigcup_{\tau \in \mathrm{Face}_{n-1-k(i+1)}\, \phi_2} \Psi(\tau, [\tau, \phi_2]).$$

If $f_i(\phi_1) \cap f_i(\phi_2) = \varnothing$, then $f_i(\phi) = \varnothing$ by the monotonicity of carrier map f_i. So let us assume that $f_i(\phi_1) \cap f_i(\phi_2) \neq \varnothing$. Let us then consider an arbitrary simplex $\sigma \in f_i(\phi_1) \cap f_i(\phi_2)$. There exists $\tau \in \mathrm{Face}_{n-1-k(i+1)}\, \phi_1$ and $\tau' \in \mathrm{Face}_{n-1-k(i+1)}\, \phi_2$ such that

$$\sigma \in \Psi(\tau, [\tau, \phi_1]) \cap \Psi(\tau', [\tau', \phi_2]).$$

This implies that there exists a simplex $\tau'' \subseteq \tau \cap \tau' \subseteq \phi_1 \cap \phi_2$ such that

$$\sigma \in \Psi(\tau'', [\tau, \phi_1]) \cap \Psi(\tau'', [\tau', \phi_2]) \subseteq \Psi(\tau'', [\tau, \phi_1] \cap [\tau', \phi_2]).$$

We claim that τ and τ' are faces of ϕ. Indeed, if τ (which is a face of ϕ_1) is not a face of ϕ, then τ is not a face of ϕ_2. Then, $[\tau, \phi_1] \cap [\tau', \phi_2] = \varnothing$, which contradicts the fact that $f_i(\phi_1) \cap f_i(\phi_2) \neq \varnothing$. Consequently,

$$\begin{aligned}
\sigma \in \Psi(\tau'', [\tau, \phi_1] \cap [\tau', \phi_2]) &\subseteq \Psi(\tau'', [\tau, \phi] \cap [\tau', \phi]) \\
&= \Psi(\tau'', [\tau \cup \tau', \phi]) && (8) \\
&\subseteq \bigcup_{\rho \in \mathrm{Face}_{n-1-k(i+1)}\, \phi} \Psi(\tau'', [\rho, \phi]) \\
&\subseteq \bigcup_{\rho \in \mathrm{Face}_{n-1-k(i+1)}\, \phi} \Psi(\rho, [\rho, \phi]) \\
&= f_i(\phi). && (9)
\end{aligned}$$

where Equation (8) follows from the fact that τ (resp., τ') is a face of every simplex in $[\tau, \phi]$ (resp., $[\tau', \phi]$). Equation (9) implies that $f_i(\phi_1) \cap f_i(\phi_2) \subseteq f_i(\phi)$, from which $f_i(\phi_1) \cap f_i(\phi_2) = f_i(\phi)$ follows by monotonicity of carrier maps. ◀

▶ **Theorem 3.16.** *For integers $t \geq 0$ and $k \geq 1$, let $\mathcal{P}$ be the protocol complex of k-set agreement after $N = \lfloor \frac{t}{k} \rfloor$ rounds in the synchronous t-resilient model with the complete communication graph, where exactly k processes crash per round. For every $J \subseteq [k+1]$, $\mathcal{P}[J]$ is a $(\dim(J) - 1)$-connected subcomplex.*

Proof. Let $J \subseteq [k+1]$. For every $i = 1, \ldots, \lfloor \frac{t}{k} \rfloor$, let $\mathcal{P}^{(i)}[J]$ be the protocol complex after i rounds, starting from $\mathcal{P}^{(0)}[J] = \mathcal{I}[J] = \Psi(\{(p_i, J) \mid i \in [n]\})$ and $\mathcal{P}[J] = \mathcal{P}^{(N)}[J]$. By construction, $\mathcal{K}_0 = \mathcal{P}^{(0)}[J] = \mathcal{I}[J], \mathcal{K}_1 = \mathcal{P}^{(1)}[J], \ldots, \mathcal{K}_N = \mathcal{P}^{(N)}[J]$ are the complexes induced by the carrier map $f_i : \mathcal{K}_i \to \mathcal{K}_{i+1}$ given by Equation (4), which crashes exactly k processes in round $i+1$, $0 \leq i < N$. Consider the sequence

$$\mathcal{K}_0 \xrightarrow{f_0} \mathcal{K}_1 \xrightarrow{f_1} \ldots \xrightarrow{f_{N-1}} \mathcal{K}_N \xrightarrow{id} \mathcal{K}_N. \tag{10}$$

By Lemma 3.15, we have $f_{N-1} \circ \cdots \circ f_0 : \mathcal{K}_0 \to \mathcal{K}_N$ is a $(k-1)$-connected carrier map. Since the input complex $\mathcal{I}[J]$ is a pseudosphere, and hence pure and shellable, Lemma 3.10 implies that $\mathcal{K}_N = \mathcal{P}[J]$ is $(k-1)$-connected. Therefore, $\mathcal{P}[J]$ is also $(\dim(J) - 1)$-connected. ◀

B Proofs for Section 4 (A Lower Bound for the Agreement Overhead for Arbitrary Graphs)

▶ **Lemma 4.6.** *Let $R = M - N$. For a static communication graph G, if $R < \mathrm{Rad}(G, t+k)$, then $g : \mathcal{P}_N \to \mathcal{P}_M$ is a $(k-1)$-connected carrier map.*

Proof. We show that for every face σ of $\mathcal{P}_N$ with $codim(\sigma) \leq k$, $g(\sigma) \neq \varnothing$. Then, since $g(\sigma)$ is a face of $\mathcal{P}_M$, $g(\sigma)$ is $(k-1)$-connected. Recall that $g(\sigma) = \{\rho\}$, where ρ is the maximal simplex in $\bigcap_{\phi \in T_\sigma} g(\phi)$ satisfying $\mathsf{names}(\rho) \subseteq \mathsf{names}(\sigma)$.

Let $S = V(G) \setminus \mathsf{names}(\sigma)$, $|S| \leq t + k$. Since $R < \mathrm{Rad}(G, t+k)$, there is a node $p \in \mathsf{names}(\sigma)$ such that p does not hear from any node in S after R rounds in G. It implies that, for every $\phi \in T_\sigma$, node p does not hear from S in G_ϕ in rounds $N+1, \ldots, M$. There is hence a vertex x with $\mathsf{names}(x) = p \in \mathsf{names}(\sigma)$ and $x \in \bigcap_{\phi \in T_\sigma} g(\phi)$. Thus, $x \in g(\sigma) \neq \varnothing$ as claimed. ◀

▶ **Theorem 4.7.** *For every graph G, $t \geq 0$ and $k \geq 1$, there are no algorithms solving k-set agreement in the t-resilient model in G in less than $R = \lfloor \frac{t}{k} \rfloor + \mathrm{Rad}(G, t+k)$ rounds.*

Proof. The proof is literally the same as the one for Theorem 3.16, except that it replaces Equation (10) by the following chain of carrier maps, with $N = \lfloor \frac{t}{k} \rfloor$ and $M = N + R$:

$$\mathcal{K}_0 \xrightarrow{f_0} \mathcal{K}_1 \xrightarrow{f_1} \ldots \xrightarrow{f_{N-1}} \mathcal{K}_N = \mathcal{P}_N \xrightarrow{g} \mathcal{P}_M \tag{11}$$

Herein, $\mathcal{K}_0 = \mathcal{P}^{(0)}[J] = \mathcal{I}[J] = \Psi(P_i, J \mid i \in \{0, \ldots, n\})$ is again the input complex, $\mathcal{K}_1 = \mathcal{P}^{(1)}[J], \ldots, \mathcal{K}_N = \mathcal{P}^{(N)}[J]$ are the protocol complexes resulting from the N crashing rounds, $\mathcal{P}_N = \mathcal{K}_N$ is the source complex for our carrier map g, and finally $\mathcal{P}[J] = \mathcal{P}_M$ is the protocol complex reached from $\mathcal{P}_N$ after $M - N = \mathrm{Rad}(G, t+k)$ rounds. ◀

▶ **Lemma 4.9.** *Let $R = M - N$. If $R < \mathrm{rad}(G, t, k)$, then $g : \mathcal{P}_N \to \mathcal{P}_M$ is a $(k-1)$-connected carrier map.*

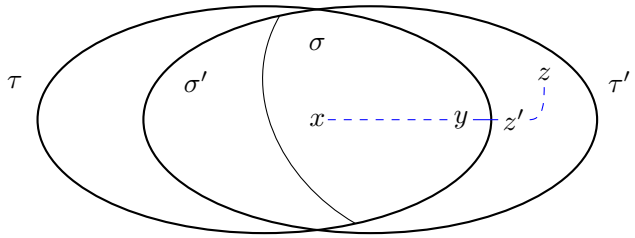

Figure 2 Two facets τ, τ' in $\mathcal{P}_N$ with $\tau \cap \tau' = \sigma'$. The vertices x resp. y resp. $\{z', z\}$ belong to σ resp. $\sigma \subseteq \sigma'$ resp. τ' as shown.

Proof. Let σ be a face of $\mathcal{P}_N$ with $codim(\sigma) \leq k$. It suffices to show that $g(\sigma) \neq \varnothing$: since $g(\sigma)$ is a face of $\mathcal{P}_M$, $g(\sigma)$ must be $(k-1)$-connected.

For every facet τ of $\mathcal{P}_N$, recall that G_τ denotes the subgraph sequence of G induced by $\mathsf{names}(\tau)$. Choose $D = \Pi \setminus \mathsf{names}(\sigma)$, which must satisfy $t \leq |D| \leq t + k$. Due to condition **C1** in Definition 4.2, there is indeed a facet τ containing σ in $\mathcal{P}_N$ such that $\mathsf{names}(\tau) = \Pi \setminus D'$. Since $R < \mathrm{rad}(G, t, k)$, there is hence a process $p \in \mathsf{names}(\sigma)$, and $D' \subseteq D, |D'| = t$ such that p does not hear from any process in $D \setminus D'$ in G_τ in rounds $N+1, \ldots, M$.

Let $x = (p, \lambda_p) \in \sigma$ be the vertex corresponding to p. So even if every node broadcasts in G_τ during rounds $N+1, \ldots, M$, node $p \in \mathsf{names}(\sigma)$ does not hear from any process in $\mathsf{names}(\tau) \setminus \mathsf{names}(\sigma)$.

Now assume that there is a facet $\tau' \supseteq \sigma$ in $\mathcal{P}_N$ such that process $p = \mathsf{names}(x)$ hears from a process $\mathsf{names}(z) \in \mathsf{names}(\tau') \setminus \mathsf{names}(\sigma)$ in $G_{\tau'}$ in rounds $N+1, \ldots, M$, see Figure 2 for an illustration. Define $\sigma' = \tau \cap \tau' \supseteq \sigma$. Let P be a path in $G_{\tau'}$ (of course of length less or equal to R) leading from $\mathsf{names}(z) \to \mathsf{names}(x)$, and let $\mathsf{names}(z') \in P \setminus \mathsf{names}(\sigma)$

be the node closest to $\mathsf{names}(x)$ in P outside σ. If $\mathsf{names}(z')$ belonged to $\mathsf{names}(\sigma')$, then $\mathsf{names}(x)$ would hear from $\mathsf{names}(z')$ also in G_τ within R rounds, through the path suffix $P' \subseteq P$ going from $\mathsf{names}(z') \to \mathsf{names}(x)$, which contradicts our assumption. Thus, $\mathsf{names}(z') \in \mathsf{names}(\tau') \setminus \mathsf{names}(\tau)$. Consequently, the path P' from $\mathsf{names}(z')$ to $\mathsf{names}(x)$ only contains $\mathsf{names}(z')$ and processes from $\mathsf{names}(\sigma)$. Let $\mathsf{names}(y) \in P \cap \mathsf{names}(\sigma)$ be the neighbor of z' in P contained in σ, i.e., $\mathsf{names}(z') \in \mathrm{In}_{\mathsf{names}(y)}(G_N)$ in round N: Indeed, in the scenario corresponding to the facet τ, $\mathsf{names}(z')$ is dead in $\mathcal{P}_N$, but alive in the scenario corresponding to τ'. Therefore, in round N, process $\mathsf{names}(y)$ hears from $\mathsf{names}(z')$ in $G_{\tau'}$, but does not hear from $\mathsf{names}(z')$ in G_τ. But then, according to g's "discarding" of "dirty" source vertices in Equation (5), $y \notin \sigma'$, contradicting our assumption.

Thus, for every facet $\tau' \supseteq \sigma$ in $\mathcal{P}_N$, $\mathsf{names}(x)$ does not hear from any process in $\mathsf{names}(\tau') \setminus \mathsf{names}(\sigma)$ in $G_{\tau'}$ in rounds $N+1, \ldots, M$. Consequently, there must be a vertex x in σ with $\mathsf{names}(x) = p$, and $x \in g(\sigma)$. So, $g(\sigma) \neq \varnothing$. ◀

C Proofs for Section 5 (Upper Bound for Fixed Graphs)

▶ **Theorem 5.1.** *There exists an algorithm solving k-set agreement in G in $\lfloor \frac{t}{k} \rfloor + D(G,t)$ rounds.*

Proof. The algorithm and its proof of correctness are directly inspired from the k-set agreement algorithm for the clique K_n in [8], and from its analysis. The algorithm is merely the min-flooding algorithm for $\lfloor \frac{t}{k} \rfloor + D(G,t)$ rounds. That is, every node sends its input value to all its neighbors at the first round, and, at each round $r \geq 2$, every node forwards the minimum value received so far to all its neighbors. After $\lfloor \frac{t}{k} \rfloor + D(G,t)$ rounds, every node outputs the smallest value it became aware of during the whole execution of the protocol, which may be its own input value, or the input value of another node received during min-flooding.

Termination and validity are satisfied by construction. We now show that at most k values are outputted in total by the (correct) nodes. Let $r \in \{1, \ldots, \lfloor \frac{t}{k} \rfloor\}$, and let us consider the system after $r-1$ rounds of min-flooding have been performed. We focus on the nodes that have not crashed during the first $r-1$ rounds, and, among these nodes, we consider those that are holding the smallest values currently in the system. More precisely, let $U \subseteq V$ be a set of k nodes that have not crashed during the first $r-1$ rounds, and satisfying that, for every value x held by a node $u \notin U$ that has not crashed during the first $r-1$ rounds, x is at least as large as any value currently hold by the nodes in U.

We claim that, if some node $u \in U$ does not crash at round r and holds a value x then, then by the end of round $r + D(G,t)$ each correct node will either know x or a smaller value. Indeed, if u does not crash at round r, then, at this round, u sends x to all its (correct) neighbors. Since $t < \kappa(G) \leq \deg(u)$, we have that, for every suffix of the current execution, at least one neighbor u' of u is correct, i.e., one correct node u' holds x at the end of round r. It follows that all the correct nodes will have received x or smaller values by the end of round $r + D(G,t)$.

As a consequence of the claim, if less than k nodes crash at some round $r \in \{1, \ldots, \lfloor \frac{t}{k} \rfloor\}$, then at the end of round $r + D(G,t)$, every correct node knows at least one value among the smallest k values present in the system at the end of round $r-1$. This guarantees that at most k distinct values are outputted by the nodes.

On the other hand, for all executions in which at least k nodes crash in each of the first $\lfloor \frac{t}{k} \rfloor$ rounds, less that k nodes can crash at round $\lfloor \frac{t}{k} \rfloor + 1$. So, let v be a node that does not crash at round $\lfloor \frac{t}{k} \rfloor + 1$, and that holds one of the smallest k values in the system after

$\lfloor \frac{t}{k} \rfloor$ rounds, say x. Round $\lfloor \frac{t}{k} \rfloor + 1$ can be viewed as the first round of broadcast of value x from node v. This broadcast will complete in $D(G,t)$ rounds in total, no matter which nodes distinct from v crashes at rounds $r \geq \lfloor \frac{t}{k} \rfloor + 1$, and no matter whether v itself crashes at some round $r > \lfloor \frac{t}{k} \rfloor + 1$. Therefore, at the end of round $\lfloor \frac{t}{k} \rfloor + D(G,t)$, all correct nodes have received at least one value among the smallest k values present in the system at the end of round $\lfloor \frac{t}{k} \rfloor$. This guarantees that at most k distinct values are outputted by the nodes. ◀