Distributed Agreement in the Arrovian Framework

Kenan Wood
Davidson College, NC, USA

Hammurabi Mendes
Davidson College, NC, USA

Jonad Pulaj
Davidson College, NC, USA

Abstract

Preference aggregation is a fundamental problem in voting theory, in which public input rankings of a set of alternatives (called *preferences*) must be aggregated into a single preference that satisfies certain soundness properties. The celebrated Arrow Impossibility Theorem is equivalent to a distributed task in a synchronous fault-free system that satisfies properties such as respecting unanimous preferences, maintaining independence of irrelevant alternatives (IIA), and non-dictatorship, along with *consensus* since only one preference can be decided.

In this work, we study a weaker distributed task in which crash faults are introduced, IIA is not required, and the consensus property is relaxed to either k-set agreement or ϵ-approximate agreement using any metric on the set of preferences. In particular, we prove several novel impossibility results for both of these tasks in both synchronous and asynchronous distributed systems. We additionally show that the impossibility for our ϵ-approximate agreement task using the *Kendall tau* or *Spearman footrule* metrics holds under extremely weak assumptions.

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Mathematics of computing → Discrete mathematics

Keywords and phrases Approximate Agreement, Set Agreement, Preference Aggregation, Voting Theory, Impossibility

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2024.32

1 Introduction

Preference aggregation is a classical problem in voting theory where every voter publishes an input *vote* (e.g. a single most-favorable candidate, a linear ranking of all the candidates, weighted rankings of candidates, etc.) and a typically centralized system aggregates the votes into a single decision outcome, according to some choice rule. Various soundness properties of preference aggregation algorithms are desirable, such as preservation of unanimous votes (unanimity), the notion that the outcome is not totally dictated by a small set of voters (non-dictatorship), and many more [8].

One of the first and most-celebrated fundamental results in voting theory was proved by Kenneth Arrow in 1951 [3]. His theorem considers the preference aggregation problem where input votes and the decision outcomes consist of weak linear rankings of the *alternatives* (candidates). In addition to unanimity and non-dictatorship, Arrow considered *independence of irrelevant alternatives* (IIA), which requires that the outcome of the preference aggregation rule with respect to the relative ordering of any two alternatives should only depend on the voters' initial relative ordering of those two alternatives, and none of the other "irrelevant" pairs. Arrow's Impossibility Theorem asserts that no deterministic preference aggregation algorithm can satisfy unanimity, non-dictatorship, and IIA simultaneously.

28th International Conference on Principles of Distributed Systems (OPODIS 2024).
Editors: Silvia Bonomi, Letterio Galletta, Etienne Rivière, and Valerio Schiavoni; Article No. 32; pp. 32:1–32:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

This result has received significant attention from the voting theory community. There are many generalizations of Arrow's theorem and a plethora of different proofs, including elementary proofs that directly exploit the unanimity and IIA axioms [24], as well as others using more advanced mathematics like algebraic topology [6], and fixpoint methods in metric space topology [41, 17].

There are recent advancements stemming from the methods used in distributed computing that yield better geometric/topological understanding of Arrow's Impossibility Theorem. Specifically, Lara, Rajsbaum, and Raventós-Pujo [40, 30] use techniques from combinatorial topology to prove a generalization of Arrow's theorem that sheds light on the result's geometric and topological structure. In general, combinatorial topology has also been successful in proving a myriad of impossibility and complete characterization results in distributed computing [27, 36, 23, 32].

In this work, we introduce novel distributed tasks that combine aspects from well-studied problems such as *set agreement* [9, 34, 37, 12, 21, 10] and *approximate agreement* [33, 35, 20, 38, 2, 31], as well as voting theoretic properties in the Arrovian framework. Specifically, we study two tasks that require unanimity *on the correct processes* along with either k-set agreement – meaning at most k different preferences are decided – or ϵ-approximate agreement – meaning any two decisions are within a distance ϵ apart, with respect to a specified metric. It is important to observe that the tasks presented here are significantly weaker than a naïve translation of the properties in Arrow's Impossibility Theorem; both the IIA and non-dictatorship properties are absent from our problem definition, and consensus is weakened to set or approximate agreement. In fact, we show in Section 4 that a natural formulation of generalized non-dictatorship properties is actually implied by unanimity on the correct processes.

The main contributions of this paper are the following.

1. In our main result, Theorem 33, we formally show that under the relatively weak property that there exists some set of input preferences with a variably fine-grained cyclic structure (see Definitions 30 and 31), neither of our Arrovian tasks are solvable for reasonable agreement parameters and sufficiently many alternatives, despite the apparent weakening of consensus.
2. We also prove that for the special cases of ϵ-approximate preference aggregation, when the distances between preferences are measured using the well-established Kendall tau [28, 29] or Spearman footrule [13] metrics, no algorithm exists if ϵ is less than a certain large quantity, expressed in terms of the metric diameter, presented in Definition 7.
3. We present a unified analysis that captures synchronous and asynchronous systems simultaneously, meant intentionally to shed light on the fundamental structural properties of the problem – such as cyclic preference patterns – that cause the impossibility.

The rest of this paper is organized as follows. Section 2 provides the necessary background on relations (Subsection 2.1) and introduces *distributed aggregation maps* in the context of synchronous fault-free distributed algorithms (Subsection 2.2). In Section 3, we prove some initial observations which are simple generalizations of Arrow's theorem in the context of distributed aggregation maps. Section 4 formally introduces the intersection of the Arrovian framework and set and approximate agreement tasks discussed above, and proves our remark that a non-dictatorship property is nonrestrictive. The main results of this paper are in Section 5, and specifically, can be found in Theorem 33 and its corollaries.

2 Background

In this section, we present the necessary background on relations, and introduce the notation of a *distributed* aggregation map and some of its relevant properties. Throughout this paper, we use the following notation: for a positive integer n, let $[n] \triangleq \{1, \ldots, n\}$.

2.1 Relations

A *relation* on a set X is a subset of $X \times X$. If R is a relation we usually write $x \; R \; y$ as shorthand for $(x, y) \in R$. A relation $R \subseteq X \times X$ is *reflexive* provided $x \; R \; x$ for all $x \in X$; it is *antisymmetric* or *strict* provided that $x \; R \; y$ and $y \; R \; x$ imply $x = y$ for all $x, y \in X$; R is *complete* if for every pair $x, y \in X$, it follows that $x \; R \; y$ or $y \; R \; x$; R is said to be *transitive* provided that for every $x, y, z \in X$, $x \; R \; y$ and $y \; R \; z$ imply $x \; R \; z$. Given a relation R and elements a, b, we often write $a \succsim_R b$ for $a \; R \; b$ and $a \succ_R b$ for $a \; R \; b$ but $a \not R \; b$. The *strict part of* R is defined by

$$\mathrm{st}(R) \triangleq \{(a, b) \in R : b \not R \; a\},$$

that is, the pairs in R without equivalence. So, $a \succ_{\mathrm{st}(R)} b$ if and only if $a \succ_R b$. If R is any relation on X and $Y \subseteq X$, we define the *restriction of* R *to* Y to be the subrelation $R|_Y \triangleq R \cap (Y \times Y)$ of R. Similarly, if $\mathbf{R} = (R_1, \ldots, R_n)$ is a vector of relations on X and $Y \subseteq X$, we define the restriction $\mathbf{R}|_Y \triangleq (R_1|_Y, \ldots, R_n|_Y)$.

The set of all reflexive complete transitive relations, called *preferences*, on a set X is denoted $P(X)$; the set of all antisymmetric reflexive complete transitive relations, called *strict preferences*, on X is denoted $L(X)$. A subset of $P(X)$ is called a *domain*. A *profile* on X is an element of $\bigcup_{k=1}^{\infty} P(X)^k$; usually, we consider only profiles in $P(X)^n$ or $P(X)^{n-t}$, where n is the number of processes and t is the maximum number of faulty processes.

2.2 Distributed Aggregation Maps

Consider a synchronous distributed message-passing system with n processes $\{p_1, \ldots, p_n\}$. In this subsection, we consider only deterministic fault-free models of computation. In particular, we are interested in algorithms of the following type. Let X be a finite set of *alternatives* (also called candidates) and let $W_I, W_O \subseteq P(X)$ be domains of preferences; we denote the size of X by m. Suppose processes take inputs from W_I, and after communicating, decide a preference in W_O satisfying certain desirable properties of a distributed voting system. Given these assumptions, distributed algorithms where processes have inputs in W_I and decide values in W_O are completely characterized by functions $F : W_I^n \to W_O^n$. To see this, since such algorithms are fault-free, every process has complete information after only one round of communication, so that the determinism of the algorithm guarantees the given map; the opposite direction is similarly trivial. This motivates the following definition.

▶ **Definition 1** (Distributed Aggregation Map). *A* distributed aggregation map *on an n-process system with input domain $W_I \subseteq P(X)$ and output domain $W_O \subseteq P(X)$ is a map from W_I^n to W_O^n.*

The properties we are interested in obtaining are characterized in the following definitions. For the rest of this section, let F be an n-process distributed aggregation map with input domain W_I and output domain W_O.

▶ **Definition 2** (Unanimity). *We say that* F *satisfies* unanimity *if the following holds: for all* $a, b \in X$ *and for all* $\mathbf{R} \in W_I^n$ *such that* $a \succ_{R_i} b$ *for all* $i \in [n]$*, we have* $a \succ_{F(\mathbf{R})_i} b$ *for all* $i \in [n]$*.*

▶ **Definition 3** (IIA). *The map* F *satisfies* independence of irrelevant alternatives *(IIA) if for all* $a, b \in X$ *and all* $\mathbf{R}, \mathbf{S} \in W_I^n$ *such that* $\mathbf{R}|_{\{a,b\}} = \mathbf{S}|_{\{a,b\}}$*, it follows that* $F(\mathbf{R})|_{\{a,b\}} = F(\mathbf{S})|_{\{a,b\}}$*.*

The following definition of decisiveness is standard terminology in the voting theory community [8], but the motivation is as follows: a set S is *decisive* if all of the strict rankings between pairs are always dictated by the strict relations of the inputs in S.

▶ **Definition 4** (Decisive). *A set* $S \subseteq [n]$ *is said to be* decisive *if* S *is nonempty and the following holds: for all* $a, b \in X$ *and all* $\mathbf{R} \in W_I^n$ *such that* $a \succ_{\mathbf{R}_i} b$ *for all* $i \in S$*, it follows that* $a \succ_{F(\mathbf{R})_j} b$ *for all* $j \in [n]$*.*

▶ **Definition 5** (Dictatorship). *If* F *has a decisive set of size at most* k*, we say that* F *is a* k-dictatorship *or that* F *is* k*-dictatorial; if* $k = 1$*, we say that* F *is simply a* dictatorship *or that* F *is* dictatorial*.*

In this paper, we are interested in set agreement and approximate agreement relaxations of the consensus property that is usually assumed in the Arrovian framework. Given a vector or list v, define $\mathrm{set}(v)$ to be the set of entries in v.

▶ **Definition 6** (Set Agreement). *If for all* $\mathbf{R} \in W_I^n$*, we have* $|\mathrm{set}(F(\mathbf{R}))| \leq k$*, then we say that* F *satisfies* k-set agreement*. The property of* 1*-set agreement is called* consensus*.*

Before describing the approximate agreement condition, we introduce some definitions and corresponding notation commonly used throughout this paper.

▶ **Definition 7** (Metric Diameter). *If* d *is a metric on a finite set* Y *and* $A \subseteq Y$*, we define the* diameter *of* A *with respect to* d *by* $\mathrm{diam}_d(A) \triangleq \max_{x,y \in A} d(x, y)$*. The diameter of a list* v *of elements of* Y *is defined to be the diameter of the set of values it contains, and is written* $\mathrm{diam}_d(v)$*.*

▶ **Definition 8** (Approximate Agreement). *If* d *is a metric on a domain containing* W_O*, then we say that* F *satisfies* ϵ-agreement *(for* $\epsilon \geq 0$*) with respect to* d *if every* $\mathbf{R} \in W_I^n$ *satisfies* $\mathrm{diam}_d(F(\mathbf{R})) \leq \epsilon$*.*

Our main results in Section 5 pay special attention to the following natural metrics [28, 29, 13] on $L(X)$ which are useful measures of distance in the context of approximate agreement. The Kendall tau metric measures the number of *pairs* of alternatives that differ, while Spearman's footrule measures that *cumulative* distance between the ranks of each of the alternatives. These metrics are formally described below.

▶ **Definition 9** (Kendall tau). *Define the* Kendall tau *metric on* $L(X)$*, denoted* $\mathtt{KT}$*, by* $\mathtt{KT}(R, S) \triangleq |\{(a, b) \in X \times X : a \succ_R b \wedge b \succ_S a\}|$ *for all* $R, S \in L(X)$*.*

▶ **Definition 10** (Rank and Spearman's footrule). *Given* $R \in L(X)$ *and* $a \in X$*, define the* rank *of* a *in* R *by* $\mathrm{rank}_R(a) \triangleq |\{b \in X : b \succsim_R a\}|$*. The* Spearman footrule *on* $L(X)$ *is a metric* $\mathtt{SF}$ *defined by* $\mathtt{SF}(R, S) \triangleq \sum_{a \in X} |\mathrm{rank}_R(a) - \mathrm{rank}_S(a)|$ *for all* $R, S \in L(X)$*.*

The following results with respect to diameter in the Kendall tau or Spearman footrule metrics can be found in [13], but are otherwise easy to prove.

▶ **Proposition 11.** *If* $|X| = m$*, then* $\mathrm{diam}_{\mathtt{KT}}(L(X)) = \frac{m^2-m}{2}$ *and* $\mathrm{diam}_{\mathtt{SF}}(L(X)) = \left\lfloor \frac{m^2}{2} \right\rfloor$.

Since we are interested in distributed analogues to Arrow's Theorem, we introduce the following definition.

▶ **Definition 12.** *We say that a domain* $W \subseteq P(X)$ *is* Arrow-Complete *(AC) if every distributed aggregation map on* n *processes with input domain* W *and output domain* $P(X)$ *that satisfies unanimity, IIA, and consensus is dictatorial.*

Arrow's Impossibility Theorem then states the following.

▶ **Theorem 13** (Arrow). *If* $n \geq 2$ *and* $m \geq 3$*, then* $P(X)$ *is AC.*

In [8], Arrow's theorem has also been generalized to show that any domain satisfying a certain *chain rule* is AC if $n \geq 2$ and $m \geq 3$.

3 Initial Observations

In this section, we state and prove our initial observations as a starting point for discussing our main theorem and corollaries in Section 5. In particular, we prove two impossibility results related to distributed aggregation maps when the consensus condition is weakened to either set agreement or approximate agreement. The core of both arguments relies on a simple coordinate-by-coordinate reduction of a distributed aggregation map. Although we find these preliminary results interesting, the main contributions of this paper can be found at the end of Section 5 (see Theorem 33 and its corollaries in Section 5). In particular, our main results are impossibility theorems in fault-prone distributed systems with either synchronous or asynchronous communication between processes.

▶ **Proposition 14.** *Let* k *be a positive integer such that* $k < n$ *and* $k < |W|$*, where* W *is an AC domain. Furthermore, assume there is a set* $P \subseteq W$ *of size at least* $k+1$ *such that for any two distinct* $R, R' \in P$*, there exists* $a, b \in X$ *satisfying* $a \succ_R b$ *and* $b \succ_{R'} a$*. Every distributed aggregation map on* W *satisfying* k*-set agreement, unanimity, and IIA is* k*-dictatorial.*

▶ **Proposition 15.** *Let* W *be an AC domain that contains some strict preference. Let* d *be a metric on* W*. If* $\epsilon < \mathrm{diam}_d(W \cap L(X))$*, then every distributed aggregation map on* W *satisfying* ϵ*-agreement (with respect to* d*), unanimity, and IIA is dictatorial.*

The key observation underlying the proofs of both of these results is the following lemma. The proof is based on a simple coordinate-by-coordinate reduction from any distributed aggregation map to a distributed aggregation map satisfying consensus. In observing that this reduction preserves unanimity and IIA, we exploit Arrow-Completeness to show that each reduced map is a dictatorship. This shows that every output coordinate is in a sense "dictated" by an input coordinate, which is unique under a very weak condition. It is important to note that we are not claiming to reprove Arrow's theorem in any way, and instead, we are building on a domain that already satisfies Arrow's theorem.

▶ **Lemma 16.** *Suppose* W *is an AC domain. Let* $F : W^n \to P(X)^n$ *be a distributed aggregation map. Suppose* F *satisfies unanimity and IIA, and let* $j \in [n]$*. Then we have the following.*

1. *Then there exists some* $i \in [n]$ *such that for all* $a, b \in X$ *and* $\mathbf{R} \in W^n$ *such that* $a \succ_{R_i} b$*, we have* $a \succ_{F(\mathbf{R})_j} b$*.*
2. *If* W *contains some two preferences* R, R' *and* $a, b \in X$ *such that* $a \succ_R b$ *and* $b \succ_{R'} a$*, then the* i *above is unique.*

Proof. Construct a map $F^j : W^n \to P(X)^n$ by setting $F^j(\mathbf{R}) = (F(\mathbf{R})_j)_{i\in[n]}$ for all $\mathbf{R} \in W^n$. It is clear F^j satisfies consensus. It is also easy to show that since F satisfies unanimity and IIA, F^j also satisfies unanimity and IIA. Since W is AC, it follows that F^j is dictatorial. Thus (1) follows from the construction of F^j and Definition 4.

To show the second part, suppose i and i' be two elements of $[n]$ satisfying the above criteria. Consider a preference profile $\mathbf{R}$ such that $R_i = R$ and $R_{i'} = R'$, where R and R' are defined in the lemma statement. Let $a, b \in X$ such that $a \succ_R b$ and $b \succ_{R'} a$. By the assumption on i and i', we know that $a \succ_{F(\mathbf{R})_j} b$ (as $R_i = R$ and i satisfies (1)) and $b \succ_{F(\mathbf{R})_j} a$ (as $R_{i'} = R'$ and i' satisfies (1)), which is a contradiction. Thus (2) follows. ◀

We are now ready to prove Proposition 14 and Proposition 15.

Proof of Proposition 14. Let P be defined as in the theorem statement. By Lemma 16, every $j \in [n]$ can be uniquely mapped to some $\delta(j) \in [n]$ such that for all $a, b \in X$ and $\mathbf{R} \in W^n$ such that $a \succ_{R_{\delta(j)}} b$, we have $a \succ_{F(\mathbf{R})_j} b$. Let $S = \{\delta(j) : j \in [n]\}$. It is clear by construction that S is a decisive set for F, so it remains to show that $|S| \le k$. Suppose for contradiction that $|S| \ge k+1$. Let $S' \subseteq S$ such that $|S'| = k+1$. It follows that there exists an injection $g : S' \to P$. Also, fix an injection $\Delta : S' \to [n]$ such that $\Delta(i) \in \delta^{-1}(i)$ for all $i \in [n]$. Construct any profile $\mathbf{R} \in W^n$ such that for all $i \in S'$, we have $R_i = g(i) \in P$. Suppose $i, i' \in S'$ are distinct. Consider $\Delta(i)$ and $\Delta(i')$, which are distinct as Δ is injective. Observe that since g is a bijection, $g(i) \ne g(i')$. By definition of P and noting that $g(i), g(i') \in P$, there exists $a, b \in X$ such that $a \succ_{R_i} b$ and $b \succ_{R_{i'}} a$. As $\delta(\Delta(i)) = i$ and $\delta(\Delta(i')) = i'$, the definition of δ implies that $a \succ_{F(\mathbf{R})_{\Delta(i)}} b$ and $b \succ_{F(\mathbf{R})_{\Delta(i')}} a$. This shows that $F(\mathbf{R})_{\Delta(i)} \ne F(\mathbf{R})_{\Delta(i')}$; that is, every $F(\mathbf{R})_{\Delta(i)}$ is distinct over all $i \in S'$. In particular, this shows that

$$|\operatorname{set}(F(\mathbf{R}))| \ge |\{F(\mathbf{R})_{\Delta(i)} : i \in S'\}| = |S'| = k+1 > k.$$

This contradicts the k-set agreement property of F. Hence $|S| \le k$, which shows F is k-dictatorial. ◀

Proof of Proposition 15. Suppose F is a distributed aggregation map on W that satisfies unanimity and IIA. Suppose for contradiction that F is not dictatorial. For each $j \in [n]$, let $\delta(j) \in [n]$ such that for all $a, b \in X$ and $\mathbf{R} \in W^n$ such that $a \succ_{R_{\delta(j)}} b$, we have $a \succ_{F(\mathbf{R})_j} b$, which is well-defined by Lemma 16. Since F is not dictatorial, not all values of $\delta(j)$ for $j \in [n]$ are equal, so that there exists distinct $j, j' \in [n]$ where $\delta(j) \ne \delta(j')$. Let $i = \delta(j)$ and $i' = \delta(j')$. This implies that for all $\mathbf{R} \in W^n$ if $R_i \in L(X)$, then $F(\mathbf{R})_j = R_i$, and if $R_{i'} \in L(X)$, then $F(\mathbf{R})_{j'} = R_{i'}$.

So let $R, R' \in W \cap L(X)$ such that $d(R, R') = \operatorname{diam}(W \cap L(X))$. Construct a profile $\mathbf{R} \in W^n$ such that $R_i = R$ and $R_{i'} = R'$. As $R_i \in L(X)$, the above remark shows that $F(\mathbf{R})_j = R$. Similarly, $F(\mathbf{R})_{j'} = R'$. By construction of R and R', if $\epsilon < \operatorname{diam}(W \cap L(X))$, then F does not satisfy ϵ-agreement with respect to d, a contradiction. ◀

4 Distributed Set and Approximate Preference Aggregation

In this section and the next, we study distributed aggregation algorithms in the presence of process failures, in both synchronous and asynchronous communication models. The previous impossibility theorems in Section 3 were only in the synchronous fault-free case, so naively introducing failures into the system only makes the task at hand more difficult, and so the impossibility trivially holds. Hence we will discard the IIA property, as it is typically viewed as the most restrictive and least necessary for preference aggregation. Additionally, we discard the dictatorship properties as well, only requiring unanimity and agreement. We will find in this section and the next that we still obtain a plethora of impossibilities, despite this apparent simplification.

For the rest of this paper, consider a set of n processes, $\{p_1, \ldots, p_n\}$, at most t (with $1 \leq t < n$) of them suffering crash failures. Let X be any finite set of $m \geq 2$ alternatives. Let $W_I, W_O \subseteq P(X)$ be input and output domains, respectively. The *identity* of process p_i is defined to be i. Let C be the set of correct process identities in a given execution of a distributed algorithm.

We focus on the following two problems, which consider distributed aggregation functions in a message-passing distributed system in the crash failure model.

▶ **Definition 17** (k-Set Preference Aggregation)**.** *Let $k \geq 1$. The* k-set preference aggregation *task with respect to W_I, W_O has the following specifications. Each process p_i selects a* private *input preference $R_i \in W_I$. Every correct process p_i decides a value $S_i \in W_O$ satisfying:*

- *k-set agreement. At most k different orders are decided: $|S_i : i \in C| \leq k$.*
- *Unanimity. For all $a, b \in X$, if every correct p_i has $a \succ_{R_i} b$, then $a \succ_{S_i} b$ for all $i \in C$.*

▶ **Definition 18** (ϵ-Approximate Preference Aggregation)**.** *Let $\epsilon \geq 0$; let d be a metric on a subset of $P(X)$ containing W_O. The* ϵ-approximate preference aggregation *task with respect to W_I, W_O, d has the following specifications. Each process p_i selects a* private *input preference $R_i \in W_I$. Every correct process p_i decides a value $S_i \in W_O$ satisfying:*

- *ϵ-approximate agreement. All correct decisions are at most ϵ apart:* $\text{diam}_d(\{S_i : i \in C\}) \leq \epsilon$.
- *Unanimity. For all $a, b \in X$, if every correct p_i has $a \succ_{R_i} b$, then $a \succ_{S_i} b$ for all $i \in C$.*

Even without the IIA and non-dictatorship properties seen in Propositions 14 and 15 (the deterministic synchronous fault-free case), a significant amount of structure is still imposed on algorithms solving either of these tasks, particularly because of the strength of this unanimity property together with an agreement condition, as we shall see in Section 5.

Another reason for the lack of non-dictatorship criteria in ϵ-approximate and k-set preference aggregation is that these properties are often implied by the unanimity condition. We make this precise in the following definition and proposition, noting that its proof is based on an indistinguishability argument commonly seen in distributed computing [18, 5].

A distributed algorithm A is k*-dictatorial* provided that the following holds for all admissible executions: there is a nonempty set $T \subseteq [n]$ with $|T| \leq k$ such that if every correct process p_i has input R_i and output S_i and $T \subseteq C$, then $a \succ_{R_i} b$ for all $i \in T$ $(a, b \in X)$ implies $a \succ_{S_i} b$ for all $i \in C$. A domain $W \subseteq P(X)$ is *non-trivial* if there are two preferences $R, S \in W$ and two alternatives $a, b \in X$ such that $a \succ_R b$ and $b \succ_S a$.

▶ **Proposition 19.** *Suppose W_I is non-trivial. Any algorithm in any synchrony model that satisfies unanimity is not k-dictatorial if $1 \leq k \leq t$.*

Proof. Suppose A is an algorithm that satisfies unanimity, and assume $1 \leq k \leq t$. Let $R, R' \in W_I$ and $a, b \in X$ such that $a \succ_R b$ and $b \succ_{R'} a$. Let $T \subseteq [n]$ such that $1 \leq |T| \leq k$. Consider an execution Ξ of A where every process is non-faulty and every process p_i for $i \in T$ has input R and every other process has input R'. Since $|T| \leq k \leq t$, there exists an admissible execution Ξ' that is identical to Ξ except the set of faulty processes is precisely T. By unanimity, in Ξ', processes p_i with $i \in [n] \setminus T$ must decide $S_i' \in W_O$ satisfying $b \succ_{S_i'} a$. Since Ξ and Ξ' are indistinguishable executions for any p_i with $i \in [n] \setminus T$, it follows that each such p_i decides some $S_i \in W_O$ satisfying $b \succ_{S_i} a$. Since a is ranked higher than b for all inputs of processes with identity in T, this shows that A is not k-dictatorial. ◀

We will make use of the following synchrony notation in the next sections. Define the *synchrony* of a distributed system to be SYNC if the system is synchronous and ASYNC if the system is asynchronous. Let the synchrony of the distributed system at hand be denoted $\texttt{tsync}$. Additionally, we say that an execution of a distributed algorithm in a particular model of computation (synchrony and maximum number of faults) is *admissible* if the execution satisfies the synchrony requirements and its number of faults is at most the maximum number of faults permissible by the model of computation.

5 Arrovian Impossibilities in Synchronous and Asynchronous Systems

In this section, we present strong impossibility results for both synchronous and asynchronous systems and both k-set and ϵ-approximate preference aggregation tasks. We begin with a simple definition that allows us to treat both synchrony models simultaneously.

▶ **Definition 20** (Synchronous Process Number). *Define the* synchronous process number $\overline{n}(\texttt{tsync})$ *of a synchrony* $\texttt{tsync} \in \{\textit{SYNC}, \textit{ASYNC}\}$ *by*

$$\overline{n}(\texttt{tsync}) \triangleq \begin{cases} n, & \text{if } \texttt{tsync} = \textit{SYNC} \\ n-t, & \text{if } \texttt{tsync} = \textit{ASYNC}. \end{cases}$$

Next we describe a convenient map that captures the relation between input and output of correct processes in an arbitrary admissible execution. Note that these reductions are not topological in nature as in [27, 10, 34, 36], and are merely convenient tools used in our impossibility results.

▶ **Definition 21** (Execution Map: SYNC). *If A is a distributed algorithm in the SYNC synchrony model with inputs in W_I and outputs in W_O, define a map $\mathtt{F}_A^{\textit{SYNC}} : W_I^n \to W_O^n$ as follows: for each $\mathbf{R} \in W_I^n$, deterministically fix an execution of A where all processes are correct and each p_i has input R_i; let S_i be the output of each p_i, and set $\mathtt{F}_A^{\textit{SYNC}}(\mathbf{R}) \triangleq \mathbf{S} = (S_1, \ldots, S_n)$.*

In the following definition, we choose the set of $n-t$ correct processes in the given executions to be $p_1, \ldots, p_{n-t}$ (and the faulty set to be $p_{n-t+1}, \ldots, p_n$) for notational convenience, but this labeling is rather arbitrary.

▶ **Definition 22** (Execution Map: ASYNC). *Let A be an algorithm in ASYNC synchrony model. Define a map $\mathtt{F}_A^{\textit{ASYNC}} : W_I^{n-t} \to W_O^{n-t}$ by setting $\mathtt{F}_A^{\textit{ASYNC}}(\mathbf{R})$ for each $\mathbf{R} \in W_I^{n-t}$ as follows: deterministically fix an admissible execution of A where $p_{n-t+1}, \ldots, p_n$ are the t faulty processes that crash before sending any messages, and p_i has input R_i for $i \in [n-t]$, and the correct processes $p_1, \ldots, p_{n-t}$ communicate perfectly synchronously for the duration of the execution; let S_i be the decided value of p_i; let $\mathtt{F}_A^{\textit{ASYNC}}(\mathbf{R}) \triangleq (S_1, \ldots, S_{n-t}) = \mathbf{S}$.*

In either synchrony cases for $\texttt{tsync} \in \{\text{SYNC}, \text{ASYNC}\}$ of the above reductions, Definition 20 shows that the reduced map is from $W_I^{\overline{n}(\texttt{tsync})}$ to $W_O^{\overline{n}(\texttt{tsync})}$. Note that even when A is a nondeterministic algorithm in the above definitions, a deterministic execution can still be chosen. For the rest of this section, fix a synchrony model $\texttt{tsync} \in \{\text{SYNC}, \text{ASYNC}\}$.

▶ **Observation 23.** *If A is an algorithm that satisfies k-set agreement in the* $\texttt{tsync}$ *synchrony model, then the map $\mathtt{F}_A^{\texttt{tsync}}$ satisfies k-set agreement.*

▶ **Observation 24.** *If A is an algorithm that satisfies ϵ-approximate agreement in the* $\texttt{tsync}$ *synchrony model, then the map $\mathtt{F}_A^{\texttt{tsync}}$ satisfies ϵ-approximate agreement.*

The following definition of u-unanimity for distributed aggregation maps can be seen as a kind of unanimity that is preserved under u-of-$\overline{n}(\texttt{tsync})$ thresholds.

▶ **Definition 25** (u-Unanimity)**.** *A map* $F : W_I^{\overline{n}(\texttt{tsync})} \to W_O^{\overline{n}(\texttt{tsync})}$ *satisfies* u-unanimity *for an integer* u *if for all* $a, b \in X$ *and* $\mathbf{R} \in W_I^{\overline{n}(\texttt{tsync})}$*, the set* $T \triangleq \{i \in [\overline{n}(\texttt{tsync})] : a \succ_{R_i} b\}$ *satisfying* $|T| \geq u$ *implies that* $a \succ_{F(\mathbf{R})_i} b$ *for all* $i \in T$*.*

The next auxiliary lemma is a simple application of a classical indistinguishability argument in distributed computing [18, 5], and connects distributed unanimity with u-unanimity for appropriate u.

▶ **Lemma 26.** *If* A *is an algorithm that satisfies unanimity in the* $\texttt{tsync}$ *synchrony model, then the map* $\mathtt{F}_A^{\texttt{tsync}}$ *satisfies* $(\overline{n}(\texttt{tsync}) - t)$*-unanimity.*

Proof. First, suppose $\texttt{tsync} = \textsc{async}$. Consider any $a, b \in X$ and $\mathbf{R} \in W_I^{\overline{n}(\texttt{tsync})} = W_I^{n-t}$. Suppose the set $T \triangleq \{i \in [n-t] : a \succ_{R_i} b\}$ satisfies $|T| \geq \overline{n}(\texttt{tsync}) - t = n - 2t$. Let Ξ be the execution of A that defines $\mathtt{F}_A^{\textsc{async}}(\mathbf{R})$. Construct a new execution Ξ' that sends and receives the same messages as Ξ (and in the same order) but the set of faulty process identities is $[n-t] \setminus T$ instead of $\{n-t+1, \dots, n\}$, and the processes with identity greater than $n-t$ have their messages delayed until after every other process has decided; we may assume that each p_i for $i > n-t$ has input R_i satisfying $a \succ_{R_i} b$, so that every $i \in C$ satisfies $a \succ_{R_i} b$. Immediately, we know that processes p_i for $i \in T$ must decide $\mathtt{F}_A^{\textsc{async}}(\mathbf{R})_i$ in Ξ'. Furthermore, there are $|[n-t] \setminus T| = (n-t) - |T| \leq (n-t) - (n-2t) = t$ faulty processes in Ξ' since $|T| \geq n - 2t$. This implies that the execution Ξ' is admissible in the $\textsc{async}$ model since asynchronous communication delays may be unbounded (but still finite). It follows that every $i \in T$ satisfies $a \succ_{\mathtt{F}_A^{\textsc{async}}(\mathbf{R})_i} b$ since p_i decides $\mathtt{F}_A^{\textsc{async}}(\mathbf{R})_i$ in Ξ' and A respects unanimity. Hence $\mathtt{F}_A^{\texttt{tsync}}(\mathbf{R})$ satisfies $(\overline{n}(\texttt{tsync}) - t)$-unanimity.

The proof for the case when $\texttt{tsync} = \textsc{sync}$ is similar, but we include it here for completeness. Suppose $\texttt{tsync} = \textsc{sync}$. Suppose A is an algorithm with inputs from W_I and outputs from W_O that satisfies unanimity. Let $a, b \in X$ and $\mathbf{R} \in W_I^n$; let $T \triangleq \{i \in [n] : a \succ_{R_i} b\}$, and assume $|T| \geq n - t$. Let Ξ be the execution of A that defines $\mathtt{F}_A^{\textsc{sync}}(\mathbf{R})$ from Definition 21. Let Ξ' be an execution of A obtained from Ξ by letting each p_j for $j \in [n] \setminus T$ be faulty but still send and receive exactly the same messages as in Ξ. Immediately by construction, for all $i \in T$, p_i decides $\mathtt{F}_A^{\textsc{sync}}(\mathbf{R})_i$ in Ξ'. Since $|T| \geq n - t$, we know $|[n] \setminus T| \leq t$, which shows that Ξ' is an admissible execution of A. Since, in Ξ', every correct process p_i (for $i \in T$) has $a \succ_{R_i} b$, and since A satisfies unanimity, it follows that every p_i for $i \in T$ decides a preference that ranks a above b. It follows that $a \succ_{\mathtt{F}_A^{\textsc{sync}}(\mathbf{R})_i} b$ for all $i \in T$, as desired. Hence $\mathtt{F}_A^{\textsc{sync}}(\mathbf{R})_i$ satisfies $(\overline{n}(\texttt{tsync}) - t)$-unanimity. ◀

A simple consequence of our auxiliary results thus far is the following. Suppose $t \geq n/2$ and let A be an algorithm that satisfies unanimity. Then by Lemma 26, $F = \mathtt{F}_A^{\textsc{async}}$ satisfies 0-unanimity, so that if $i \in [n-t]$ and $\mathbf{R} \in W_I^{n-t}$ has $R_i \in L(X)$, then $F(\mathbf{R})_i = R_i$. Hence, k-set preference aggregation is impossible in the $\textsc{async}$ synchrony model as long as $k < n - t$ and $|W_I \cap L(X)| \geq n - t$. Similarly, ϵ-approximate preference aggregation is impossible in the $\textsc{async}$ synchrony model if $W_I \cap L(X) \neq \emptyset$ and $\epsilon < \mathrm{diam}_d(W_I \cap L(X))$. Hence most of the interesting cases in the asynchronous communication model are when $t < n/2$.

Our main results in this section rely on the definitions below, inspired by the Mendes–Herlihy algorithm [33] for approximate agreement in $\mathbb{R}^d$, except with minor differences.

▶ **Definition 27** (Unanimity Set)**.** *For an indexed set* $M = \{R_\alpha\}_{\alpha \in J} \subseteq W_I$*, define the* unanimity set *of* M *by*

$$\mathrm{unanimity}(M) \triangleq \{S \in W_O : \forall a, b \in X, [(\forall \alpha \in J, a \succ_{R_\alpha} b) \implies a \succ_S b]\}.$$

Thus, given an indexed set $M = \{R_\alpha\}_{\alpha\in J} \subseteq W_I$ for some $J \subseteq [n]$, the unanimity set of M is the set of admissible output rankings required by unanimity, when process p_α has input R_α. The next lemma introduces the concept of *safe area* in this context.

▶ **Definition 28** (Safe Area). *Let $M = \{R_\alpha\}_{\alpha\in J} \subseteq W_I$, where $J \subseteq [n]$ and $|J| > t$. For each $i \in J$, let*

$$\mathrm{safe}_i(M) \triangleq \bigcap_{\substack{T\subseteq J:\\ |T|=|J|-t\\ i\in T}} \mathrm{unanimity}(\{R_\alpha : \alpha \in T\}),$$

called the safe area *of M with respect to p_i.*

We often slightly abuse notation by treating a tuple $(x_i)_{i=1}^k$ as an indexed set $\{x_i\}_{i\in[k]}$, as we do in the following lemma. This next lemma shows the connection between unanimity of an algorithm and the safe area concept above.

▶ **Lemma 29.** *Let $F : W_I^{\overline{n}(\texttt{tsync})} \to W_O^{\overline{n}(\texttt{tsync})}$ be a distributed aggregation map that satisfies $(\overline{n}(\texttt{tsync}) - t)$-unanimity. Then for all $i \in [\overline{n}(\texttt{tsync})]$ and all $\mathbf{R} \in W_I^{\overline{n}(\texttt{tsync})}$, we have $F(\mathbf{R})_i \in \mathrm{safe}_i(\mathbf{R})$.*

Proof. Let $i \in [\overline{n}(\texttt{tsync})]$ and $\mathbf{R} \in W_I^{\overline{n}(\texttt{tsync})}$. Let us prove that $F(\mathbf{R})_i \in \mathrm{safe}_i(\mathbf{R})$. Suppose $T \subseteq [\overline{n}(\texttt{tsync})]$ such that $i \in T$ and $|T| = \overline{n}(\texttt{tsync}) - t$. It suffices to show that $F(\mathbf{R})_i \in \mathrm{unanimity}(\{R_\alpha : \alpha \in T\})$. To this end, suppose $a, b \in X$ such that for all $\alpha \in T$, $a \succ_{R_\alpha} b$. Since F satisfies $(\overline{n}(\texttt{tsync}) - t)$-unanimity and $|T| \geq \overline{n}(\texttt{tsync}) - t$, it follows that $a \succ_{F(\mathbf{R})_i} b$ as $i \in T$. This shows that $F(\mathbf{R})_i \in \mathrm{unanimity}(\{R_\alpha : \alpha \in T\})$, as desired. ◀

To show an impossibility, we seek cases where $\mathrm{safe}_i(\mathbf{R}) = \{R_i\}$ for all i. This leads to the notion of a *k-cyclic* profile, expressed through Definitions 30 and 31. For convenience in the following definition and in the proof of Lemma 32 below, we use the following shorthand notation. If X_1 and X_2 are disjoint nonempty subsets of X and $R \in W_I$, we write $X_1 \succ_R X_2$ as shorthand for $\forall x_1 \in X_1, \forall x_2 \in X_2, x_1 \succ_R x_2$; that is, we write $X_1 \succ_R X_2$ if and only if R ranks every element of X_1 above every element of X_2.

▶ **Definition 30** (Cyclic Preference List). *Let $1 \leq k \leq m$. A k-*cyclic preference list *in W_I is a list of k distinct preferences $R_1, \dots, R_k$ in W_I such that the following holds: there exists a partition $X_1 \cup \cdots \cup X_k$ of X into nonempty sets and preferences $B_j \in L(X_j)$ for $j \in [k]$ such that every R_i for $i \in [k]$ respects the preferences of every B_j (that is, $R_i|_{X_j} = B_j$ for all j) and satisfies*

$$X_i \succ_{R_i} X_{i+1} \succ_{R_i} \cdots \succ_{R_i} X_k \succ_{R_i} X_1 \succ_{R_i} \cdots \succ_{R_i} X_{i-1}.$$

In this case, the preferences B_j for $j \in [k]$ are called the blocks *of $R_1, \dots, R_k$.*

▶ **Definition 31** (Cyclic Profile). *Let $1 \leq k \leq m$. A k-*cyclic profile *with synchrony* `tsync` *is a profile $\mathbf{R} \in W_I^{\overline{n}(\texttt{tsync})}$ such that W_I has a k-cyclic preference list $R'_1, \dots, R'_k$ and there is an equitable partition*[1] *$A_1 \cup \cdots \cup A_k$ of $[\overline{n}(\texttt{tsync})]$ where for all $i \in [k]$ and $j \in A_i$, we have $R_j = R'_i$.*

The set of all k-cyclic profiles with synchrony `tsync` *is denoted $\mathcal{C}_k^{\texttt{tsync}}$. Finally, let*

$$\mathcal{C}^{\texttt{tsync}} \triangleq \bigcup_{\frac{\overline{n}(\texttt{tsync})}{t} \leq k \leq m} \mathcal{C}_k^{\texttt{tsync}}.$$

[1] A partition $\mathcal{P}$ of a finite set S is *equitable* if $|P| \in \{\lfloor |X|/|\mathcal{P}| \rfloor, \lceil |X|/|\mathcal{P}| \rceil\}$ for all $P \in \mathcal{P}$. We allow empty sets in this partition since we may have $k > \overline{n}(\texttt{tsync})$, forcing at least one of the sets to be empty.

Let us show that this definition of cyclic profiles satisfies the intuition stated above.

▶ **Lemma 32.** *If* $\mathbf{R} \in \mathcal{C}^{\texttt{tsync}}$*, then for all* $i \in [\overline{n}(\texttt{tsync})]$*, we have* $\text{safe}_i(\mathbf{R}) = \{R_i\} \cap W_O$*.*

Proof. Let $\mathbf{R} \in \mathcal{C}^{\texttt{tsync}}$ and let $i^* \in [\overline{n}(\texttt{tsync})]$. It is obvious from the definition of $\text{safe}_{i^*}(\mathbf{R})$ that $\{R_{i^*}\} \cap W_O \subseteq \text{safe}_{i^*}(\mathbf{R})$. Now suppose $S \in \text{safe}_{i^*}(\mathbf{R})$. Clearly $S \in W_O$, so it remains to show $S = R_{i^*}$. Since $\mathbf{R} \in \mathcal{C}^{\texttt{tsync}}$, there exists some integer k with $\frac{\overline{n}(\texttt{tsync})}{t} \leq k \leq m$, and there exists a k-cyclic preference list $R'_1, \ldots, R'_k$ along with an equitable partition $A_1 \cup \cdots \cup A_k$ of $[\overline{n}(\texttt{tsync})]$ satisfying Definition 31; let $j^* \in [k]$ be the unique integer such that $i^* \in A_{j^*}$. Since $R'_1, \ldots, R'_k$ is k-cyclic, there exists a partition $X_1 \cup \cdots \cup X_k$ of X into nonempty sets, and $B_j \in L(X_j)$ for $j \in [k]$ that satisfies the properties in Definition 30.

Since every $i \in [\overline{n}(\texttt{tsync})]$ and $j \in [k]$ satisfy $R_i|_{X_j} = B_j \in L(X)$, it is easy to see that $S|_{X_j} = B_j$, as $S \in \text{safe}_{i^*}(\mathbf{R})$. Let $j_1, j_2 \in [k]$ such that $X_{j_1} \succ_{R_{i^*}} X_{j_2}$ and there exists no $j_3 \in [k]$ such that $X_{j_1} \succ_{R_{i^*}} X_{j_3} \succ_{R_{i^*}} X_{j_2}$. Notice that $R_{i^*} = R'_{j^*}$ since $i^* \in A_{j^*}$. It is easy to show that for all $j \in [k]$, we have $X_{j_1} \succ_{R'_j} X_{j_2}$ if and only if $j \neq j_2$; this implies that for all $i \in [\overline{n}(\texttt{tsync})]$, we have $X_{j_1} \succ_{R_i} X_{j_2}$ if and only if $i \notin A_{j_2}$. In particular, this shows that $i^* \notin A_{j_2}$. Furthermore, since $A_1 \cup \cdots \cup A_k$ is an equitable partition of $[\overline{n}(\texttt{tsync})]$ and $k \geq \frac{\overline{n}(\texttt{tsync})}{t}$, we have

$$|A_{j_2}| \leq \left\lceil \frac{\overline{n}(\texttt{tsync})}{k} \right\rceil \leq \left\lceil \frac{\overline{n}(\texttt{tsync})}{\overline{n}(\texttt{tsync})/t} \right\rceil = t.$$

We now let $T \triangleq [\overline{n}(\texttt{tsync})] \setminus A_{j_2}$. Hence $i^* \in T$ and $|T| \geq \overline{n}(\texttt{tsync}) - t$. It follows from Definition 28 that $S \in \text{unanimity}(\{R_\alpha : \alpha \in T\})$. By Definition 27, it follows that $X_{j_1} \succ_S X_{j_2}$.

Since $R'_{j^*} = R_{i^*}$, the definition of j^* and Definition 30 show that

$$X_{j^*} \succ_{R_{i^*}} X_{j^*+1} \succ_{R_{i^*}} \cdots \succ_{R_{i^*}} X_k \succ_{R_{i^*}} X_1 \succ_{R_{i^*}} \cdots \succ_{R_{i^*}} X_{j^*-1}.$$

Since the choice of j_1, j_2 was arbitrary in the above argument, this implies that

$$X_{j^*} \succ_S X_{j^*+1} \succ_S \cdots \succ_S X_k \succ_S X_1 \succ_S \cdots \succ_S X_{j^*-1}.$$

Thus, because $S|_{X_j} = B_j = R_{i^*}|_{X_j}$ for all $j \in [k]$, we obtain $S = R_{i^*}$, as desired. ◀

We are now ready to state and prove our main impossibility theorem.

▶ **Theorem 33** (Main)**.** *Let* $k < n$*. Then there is no algorithm solving* k*-set preference aggregation in the synchrony model* $\texttt{tsync}$ *if there exists a* j*-cyclic profile in* $\mathcal{C}^{\texttt{tsync}}$ *for some* $j > k$*. Similarly, there is no algorithm solving* ϵ*-approximate preference aggregation in the synchrony model* $\texttt{tsync}$ *if* $\mathcal{C}^{\texttt{tsync}} \neq \emptyset$ *and* $\epsilon < \max_{\mathbf{R} \in \mathcal{C}^{\texttt{tsync}}} \text{diam}_d(\mathbf{R})$*.*

Proof. Suppose $k < \overline{n}(\texttt{tsync})$, which is always true if $\texttt{tsync} = \text{SYNC}$. Suppose $\mathbf{R} \in \mathcal{C}^{\texttt{tsync}}$ is a j-cyclic profile for some $j > k$. Suppose A is an algorithm that solves k-set preference aggregation in the $\texttt{tsync}$ synchrony model. Then the map $\texttt{F}_A^{\texttt{tsync}}$ satisfies k-set agreement and $(\overline{n}(\texttt{tsync}) - t)$-unanimity by Observation 23 and Lemma 26. Since $\mathbf{R}$ is j-cyclic and $j > k$ and $k < \overline{n}(\texttt{tsync})$, there exists a set $S \subseteq [\overline{n}(\texttt{tsync})]$ such that $|S| = k + 1$ and every R_i is distinct over all $i \in S$. Then by Lemma 29, for all $i \in S$, we have $\texttt{F}_A^{\texttt{tsync}}(\mathbf{R})_i \in \text{safe}_i(\mathbf{R})$. By Lemma 32, $\text{safe}_i(\mathbf{R}) \subseteq \{R_i\}$ for all $i \in S$, which implies that $\texttt{F}_A^{\texttt{tsync}}(\mathbf{R})_i = R_i$ for all $i \in S$. Since $|S| = k + 1$, this shows that $\texttt{F}_A^{\texttt{tsync}}$ does not satisfy k-set agreement, a contradiction. This proves the result when $k < \overline{n}(\texttt{tsync})$.

Now, suppose $k \geq \overline{n}(\texttt{tsync})$, so $\texttt{tsync} = \text{ASYNC}$. Suppose there exists a j-cyclic profile in $\mathcal{C}^{\texttt{tsync}} = \mathcal{C}^{\text{ASYNC}}$ for some $k < j \leq n$. This profile can be extended to a j-cyclic profile $\mathbf{R} \in W_I^n$. Since $j > \overline{n}(\text{ASYNC}) = n - t$, we know $j \geq n - t + 1 \geq \frac{n}{t}$, so that $\mathbf{R} \in \mathcal{C}^{\text{SYNC}}$. Since we have already proved the synchronous case, this shows that there is no algorithm that solves k-set preference aggregation in the SYNC synchrony model, so certainly no algorithm solves the task in the ASYNC model (such an algorithm necessarily allows synchronous executions).

For the second part, suppose $\mathcal{C}^{\mathtt{tsync}} \neq \emptyset$ and $\epsilon < \max_{\mathbf{R}\in\mathcal{C}^{\mathtt{tsync}}} \mathrm{diam}_d(\mathbf{R})$. Suppose A is an algorithm that solves ϵ-approximate preference aggregation in the $\mathtt{tsync}$ synchrony model. Then the map $\mathtt{F}_A^{\mathtt{tsync}}$ satisfies ϵ-approximate agreement and $(\overline{n}(\mathtt{tsync}) - t)$-unanimity by Observation 24 and Lemma 26. Pick any $\mathbf{R} \in \arg\max_{\mathbf{R}'\in\mathcal{C}^{\mathtt{tsync}}} \mathrm{diam}_d(\mathbf{R}')$. Let $i, j \in [\overline{n}(\mathtt{tsync})]$ such that $d(R_i, R_j) = \mathrm{diam}_d(\mathbf{R})$. Then $d(R_i, R_j) > \epsilon$. By Lemma 29, $\mathtt{F}_A^{\mathtt{tsync}}(\mathbf{R})_i \in \mathrm{safe}_i(\mathbf{R})$ and $\mathtt{F}_A^{\mathtt{tsync}}(\mathbf{R})_j \in \mathrm{safe}_j(\mathbf{R})$. Lemma 32 implies $\mathrm{safe}_i(\mathbf{R}) \subseteq \{R_i\}$ and $\mathrm{safe}_j(\mathbf{R}) \subseteq \{R_j\}$, so that that $\mathtt{F}_A^{\mathtt{tsync}}(\mathbf{R})_i = R_i$ and $\mathtt{F}_A^{\mathtt{tsync}}(\mathbf{R})_j = R_j$. It follows that

$$d(\mathtt{F}_A^{\mathtt{tsync}}(\mathbf{R})_i, \mathtt{F}_A^{\mathtt{tsync}}(\mathbf{R})_j) = d(R_i, R_j) > \epsilon,$$

which contradicts the ϵ-agreement property of $\mathtt{F}_A^{\mathtt{tsync}}$. The theorem follows. ◀

Let us now show some consequences of this theorem on special cases. First, we consider k-set preference aggregation on *full domains*, which contain all strict preferences on X.

▶ **Corollary 34** (k-Set Impossibility: Full Domain)**.** *Suppose $L(X) \subseteq W_I \subseteq P(X)$. If $m \geq \frac{\overline{n}(\mathtt{tsync})}{t}$ and $k < \min\{m, n\}$, then there is no algorithm solving k-set preference aggregation on W_I, W_O in the synchrony model* $\mathtt{tsync}$.

Proof. Suppose $m \geq \frac{\overline{n}(\mathtt{tsync})}{t}$ and $k < \min\{m, n\}$, so $k < m$ and $k < n$. Write $X = \{a_1, \dots, a_m\}$. Consider an m-cyclic preference list $R_1, \dots, R_m$ in $L(X)$ given by setting each $X_i = \{a_i\}$ and B_i to be the only preference in $L(X_i)$ and using the relations in Definition 30. Let $\mathbf{R} \in L(X)^{\overline{n}(\mathtt{tsync})}$ be an m-cyclic profile constructed using Definition 31 and any equitable partition $A_1 \cup \cdots \cup A_m$ of $[\overline{n}(\mathtt{tsync})]$. Since $m \geq \frac{\overline{n}(\mathtt{tsync})}{t}$, we know that $\mathbf{R} \in \mathcal{C}^{\mathtt{tsync}}$. The result follows from Theorem 33 since $\mathbf{R}$ is m-cyclic and $k < n$ and $k < m$. ◀

For our analysis of the ϵ-approximate preference aggregation problem, we consider the Kendall tau metric (Definition 9) and Spearman's footrule metric (Definition 10) on $L(X)$.

The assumption that j is even in the following result is only for simplicity of algebraic expressions, and is not fundamental to the result itself.

▶ **Corollary 35** (ϵ-Approximate Kendall Tau: General Impossibility)**.** *Suppose there exists a j-cyclic profile $\mathbf{R} \in \mathcal{C}^{\mathtt{tsync}}$ for some $j \geq \frac{\overline{n}(\mathtt{tsync})}{t}$ such that each block of $\mathbf{R}$ is on at least $\ell \geq 1$ alternatives. If $\epsilon < \lfloor j^2/4 \rfloor \ell^2$, then no algorithm solves ϵ-approximate preference aggregation on $W_I, W_O, \mathtt{KT}$ in the $\mathtt{tsync}$ synchrony model. In particular, if j is even and $\delta \triangleq \ell \cdot \frac{j}{m}$, then ϵ-approximate preference aggregation is impossible in the $\mathtt{tsync}$ synchrony model for*

$$\epsilon < \frac{\delta^2}{2} \cdot \mathrm{diam}_{\mathtt{KT}}(W_I).$$

Proof. Since there exists a j-cyclic profile for some $j \geq \frac{\overline{n}(\mathtt{tsync})}{t}$, then W_I has a j-cyclic preference list $R'_1, \dots, R'_j$; let $X_1, \dots, X_j$ be the blocks of $R'_1, \dots, R'_j$, each of size at least ℓ. It follows that there exists a j-cyclic profile $\mathbf{R} \in \mathcal{C}^{\mathtt{tsync}}$ (constructed with the preference list $R'_1, \dots, R'_j$) such that $R'_1, R'_{\lfloor j/2 \rfloor+1} \in \mathrm{set}(\mathbf{R})$. Observe that for all $a, b \in X$, we have $R'_1|_{\{a,b\}} \neq R'_{\lfloor j/2\rfloor+1}|_{\{a,b\}}$ if and only if $a \in X_1 \cup \cdots \cup X_{\lfloor j/2 \rfloor}$ and $b \in X_{\lfloor j/2\rfloor+1} \cup \cdots \cup X_j$ or vice versa (see Figure 1); hence $\mathtt{KT}(R'_1, R'_{\lfloor j/2\rfloor+1}) = \left(\sum_{i=1}^{\lfloor j/2 \rfloor} |X_i|\right) \cdot \left(\sum_{i=\lfloor j/2\rfloor+1}^{j} |X_i|\right)$.

Let $\epsilon < \lfloor j^2/4 \rfloor \ell^2$. Then, since $\ell = \delta \cdot \frac{m}{j}$ and by definition of a j-cyclic preference list,

$$\begin{aligned}\mathrm{diam}_{\mathtt{KT}}(\mathrm{set}(\mathbf{R})) \geq \mathtt{KT}(R'_1, R'_{\lfloor j/2 \rfloor+1}) &= \left(\sum_{i=1}^{\lfloor j/2 \rfloor} |X_i|\right) \cdot \left(\sum_{i=\lfloor j/2 \rfloor+1}^{j} |X_i|\right) \\ &\geq \lfloor j/2 \rfloor \cdot \ell \cdot (j - \lfloor j/2 \rfloor) \cdot \ell = \lfloor j/2 \rfloor \cdot \lceil j/2 \rceil \cdot \ell^2 = \lfloor j^2/4 \rfloor \cdot \ell^2 > \epsilon.\end{aligned}$$

The first part then follows from Theorem 33 since $\epsilon < \max_{\mathbf{R}' \in \mathcal{C}^{\mathtt{tsync}}} \mathrm{diam}_{\mathtt{KT}}(\mathrm{set}(\mathbf{R}'))$. For the second part, suppose j is even and $\ell = \delta \cdot \frac{m}{j}$ for some real $\delta > 0$. Then, by the first part, ϵ-approximate agreement is impossible if $\epsilon < \lfloor j^2/4 \rfloor \ell^2$. By Proposition 11,

$$\lfloor j^2/4 \rfloor \ell^2 = \frac{j^2}{4} \cdot \left(\delta \cdot \frac{m}{j}\right)^2 = \frac{\delta^2}{2} \cdot \frac{m^2}{2} \geq \frac{\delta^2}{2} \cdot \mathrm{diam}_{\mathtt{KT}}(L(X)) \geq \frac{\delta^2}{2} \cdot \mathrm{diam}_{\mathtt{KT}}(W_I).$$

Hence $\epsilon < \frac{\delta^2}{2} \cdot \mathrm{diam}_{\mathtt{KT}}(W_I)$ implies $\epsilon < \lfloor j^2/4 \rfloor \ell^2$, proving the result. ◀

Consider the δ value in the context of Corollary 35 (and in the next corollary). Observe that δ is at least the ratio of the *smallest* block size to the *average* block size $\frac{m}{j}$. We necessarily have $\delta < 1$, but the closer δ is to 1, the more the partition $X_1 \cup \cdots \cup X_j$ of X is evenly distributed and closer to being equitable.

R'_1	$R'_{\lfloor j/2 \rfloor+1}$
B_1	$B_{\lfloor j/2 \rfloor+1}$
B_2	$B_{\lfloor j/2 \rfloor+2}$
$\vdots$	$\vdots$
$B_{\lfloor j/2 \rfloor}$	B_j
$B_{\lfloor j/2 \rfloor+1}$	B_1
$B_{\lfloor j/2 \rfloor+2}$	B_2
$\vdots$	$\vdots$
B_j	$B_{\lfloor j/2 \rfloor}$

Figure 1 Visualization of R'_1 and $R'_{\lfloor j/2 \rfloor+1}$ in the proofs of Corollaries 35 and 36.

▶ **Corollary 36** (ϵ-Approximate Spearman Footrule: General Impossibility)**.** *Suppose there exists a j-cyclic profile $\mathbf{R} \in \mathcal{C}^{\mathtt{tsync}}$ for some $j \geq \frac{\overline{n}(\mathtt{tsync})}{t}$ such that each block of $\mathbf{R}$ is on at least $\ell \geq 1$ alternatives. If $\epsilon < \lfloor j^2/2 \rfloor \ell^2$, then no algorithm solves ϵ-approximate preference aggregation on $W_I, W_O, \mathtt{SF}$ in the $\mathtt{tsync}$ synchrony model. In particular, if j is even and $\delta \triangleq \ell \cdot \frac{j}{m}$, then ϵ-approximate preference aggregation is impossible in the $\mathtt{tsync}$ synchrony model for*

$$\epsilon < \delta^2 \cdot \mathrm{diam}_{\mathtt{SF}}(W_I).$$

Proof. The proof of this result is almost identical to the proof of Corollary 35. Let $R'_1, \ldots, R'_j$ and $\mathbf{R} \in \mathcal{C}^{\mathtt{tsync}}$ and $X_1, \ldots, X_j$ be defined as in the proof of Corollary 35. Observe that for every $a \in X_1 \cup \cdots \cup X_{\lfloor j/2 \rfloor}$, we have $|\mathrm{rank}_{R'_1}(a) - \mathrm{rank}_{R'_{\lfloor j/2 \rfloor}}(a)| = \sum_{i=\lfloor j/2 \rfloor+1}^{j} |X_i|$ and for all $a \in X_{\lfloor j/2 \rfloor+1} \cup \cdots \cup X_j$, we have $|\mathrm{rank}_{R'_1}(a) - \mathrm{rank}_{R'_{\lfloor j/2 \rfloor}}(a)| = \sum_{i=1}^{\lfloor j/2 \rfloor} |X_i|$. See Figure 1 for a visualization of these observations.

If $\epsilon < \lfloor j^2/2 \rfloor \cdot \ell^2$, then the observations above imply

$$\begin{aligned}\mathrm{diam}_{\mathtt{SF}}(\mathrm{set}(\mathbf{R})) \geq \mathtt{SF}(R'_1, R'_{\lfloor j/2 \rfloor+1}) &= 2 \cdot \left(\sum_{i=1}^{\lfloor j/2 \rfloor} |X_i| \right) \cdot \left(\sum_{i=\lfloor j/2 \rfloor+1}^{j} |X_i| \right) \\ &= 2 \cdot \mathtt{KT}(R'_1, R'_{\lfloor j/2 \rfloor+1}) \geq 2 \cdot \lfloor j^2/4 \rfloor \cdot \ell^2 = \lfloor j^2/2 \rfloor \cdot \ell^2 > \epsilon.\end{aligned}$$

The result follows by Theorem 33. The second part of this corollary holds similarly to the proof of Corollary 35. Suppose j is even and write $\ell = \delta \cdot \frac{m}{j}$. The result immediately follows from the first part of this corollary and the inequality

$$\lfloor j^2/2 \rfloor \, \ell^2 = \frac{j^2}{2} \cdot \left(\delta \cdot \frac{m}{j} \right)^2 = \delta^2 \cdot \frac{m^2}{2} \geq \delta^2 \cdot \mathrm{diam}_{\mathtt{SF}}(L(X)) \geq \delta^2 \cdot \mathrm{diam}_{\mathtt{SF}}(W_I),$$

which follows from the equality $\mathrm{diam}_{\mathtt{SF}}(L(X)) = \left\lfloor \frac{m^2}{2} \right\rfloor$ in Proposition 11. ◀

We conclude this section by analyzing approximate preference aggregation on the full domain consisting of all strict preferences on X.

▶ **Corollary 37** (ϵ-Approximate Kendall Tau: Full Domain Impossibility)**.** *Assume $W_I = W_O = L(X)$. If $m \geq \frac{\overline{n}(\mathtt{tsync})}{t}$, then no algorithm solves ϵ-approximate preference aggregation on $W_I, W_O, \mathtt{KT}$ in the* `tsync` *synchrony model, if $\epsilon < \lfloor m^2/4 \rfloor$, and in particular, if $\epsilon < \frac{1}{2} \cdot \mathrm{diam}_{\mathtt{KT}}(L(X))$.*

Proof. Using the construction in the proof of Corollary 34, there exists a m-cyclic profile $\mathbf{R} \in \mathcal{C}^{\mathtt{tsync}}$. Since the blocks of $\mathbf{R}$ form an equitable partition of X, each block contains exactly one alternative. Since $m \geq \frac{\overline{n}(\mathtt{tsync})}{t}$, Corollary 35 shows that there is no ϵ-approximate preference aggregation algorithm (on $W_I, W_O, \mathtt{KT}$) in the `tsync` synchrony model for $\epsilon < \lfloor m^2/4 \rfloor \cdot 1^2 = \lfloor m^2/4 \rfloor$. The second part of the result follows from the following inequality: $\frac{1}{2} \cdot \mathrm{diam}_{\mathtt{KT}}(L(X)) = \frac{m^2-m}{4} \leq \left\lfloor \frac{m^2}{4} \right\rfloor$. The first equality holds by Proposition 11 and the last inequality is obvious if $m \geq 4$, and one may easily verify that it is true when $m \leq 3$. ◀

▶ **Corollary 38** (ϵ-Approximate Spearman Footrule: Full Domain Impossibility)**.** *Assume $W_I = W_O = L(X)$. If $m \geq \frac{\overline{n}(\mathtt{tsync})}{t}$, then no algorithm solves ϵ-approximate preference aggregation on $W_I, W_O, \mathtt{SF}$ in the* `tsync` *synchrony model, if $\epsilon < \lfloor m^2/2 \rfloor$, and in particular, if $\epsilon < \mathrm{diam}_{\mathtt{SF}}(L(X))$.*

Proof. Using a similar argument as the previous corollary and by Corollary 36, we have the following. If $m \geq \frac{\overline{n}(\mathtt{tsync})}{t}$, then there is no ϵ-approximate preference aggregation algorithm (on $W_I, W_O, \mathtt{SF}$) in the `tsync` synchrony model for $\epsilon < \lfloor m^2/2 \rfloor$. The second part of the result follows from the equality $\mathrm{diam}_{\mathtt{SF}}(L(X)) = \lfloor m^2/2 \rfloor$, which holds by Proposition 11. ◀

6 Related Work

The Arrovian framework in voting theory has already received a lot of attention; however, its intersection with distributed computing appears to be recent. To the best of our knowledge, distributed combinatorial topology techniques such as the *index lemma* on simplicial complexes were first used in 2022 [40] to prove the base case ($m = 3$, $n = 2$) of Arrow's theorem topologically, and then extended via induction. This work was later improved in 2024 [30] by proving a domain-generalization of Arrow's theorem using only distributed combinatorial topology on *all* cases. Other work has studied social welfare and social choice distributed algorithms satisfying consensus and significantly weaker validity conditions (e.g. preserving unanimity of only the highest ranked alternative) than our unanimity condition [11].

The theory in the Arrovian framework itself is full of rich results. In particular, Arrow's theorem and domain-generalizations thereof have been proved via many combinatorial methods [8, 24], as well as techniques from algebraic topology [6] and fixed-point methods in metric spaces [41, 17]. More generally, (algebraic) topological techniques have proven successful in voting theory. For example, the Gibbard–Satterthwaite theorem – which shows the non-existence of a *social choice function* (inputs being preferences and output being a single alternative) satisfying surjectivity and strategy-proofness – has been proven using combinatorial methods and algebraic topological methods [26, 7].

On the distributed computing side, the k-set agreement task (in which processes must decide on at most k values with a notion of decision validity) was introduced in [9]. Other formulations of this task may be found in [39]. The ϵ-approximate agreement distributed task on the real line was well-studied in [14, 16] and was proven to be solvable in fully asynchronous systems with relatively high resilience, despite the consensus unsolvability in such systems with only one crash fault [19]. These results were later generalized to *multidimensional* approximate agreement in [33, 35]. Multidimensional approximate agreement has been studied extensively since then [1, 25, 15, 22, 4]. Discrete versions of approximate agreement have been formulated on graphs and simplicial complexes [38, 2, 31]; these tasks are somewhat similar to the ϵ-approximate preference aggregation task discussed in this paper.

7 Conclusion

In this paper, we propose two novel distributed tasks in the Arrovian framework and prove strong impossibility results on these tasks in crash-prone, synchronous and asynchronous systems. We adapt previously well-studied distributed tasks – namely, set and approximate agreement – to the context of preference aggregation, by replacing the respective validity properties with unanimity on the correct processes. Our impossibility results are very general, and apply to k-set preference aggregation on a full domain, and ϵ-approximate preference aggregation with both the Kendall tau and Spearman footrule metrics.

Using the Kendall tau metric on a domain of strict preferences illuminates a particularly fascinating connection to a more general "metric space approximate agreement" task. One may embed the given domain into a higher dimensional Euclidean space by examining the order of each pair of alternatives and mapping it to a binary real. With such an embedding, one may think of the approximate preference aggregation task (with the Kendall tau metric) as a more traditional multidimensional approximate agreement problem, where the convexity and agreement properties are with respect to the Euclidean L^1 "taxicab" metric (instead of the usual L^2 metric) and the definition of convexity is adjusted to a more "total" convexity. Many of the lemmas in this paper easily generalize to this metric space framework. Exploration of this more general approximate agreement task would be interesting future work.

Another interesting direction for future work is the following. In our execution map of an asynchronous algorithm, we make a rather arbitrary choice for the set of silent (initially crashed) processes (see the sentence before Definition 22). Using a more rich view and including all possible sets of silent processes in the map (or other techniques) could potentially generate a simplicial complex that captures this information. It would be highly insightful to determine if topological methods could be used in this way to obtain stronger asynchronous results in our Arrovian framework, but also in the more general metric space framework discussed above.

References

1 Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman. Asymptotically optimal validated asynchronous byzantine agreement. *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, 2019. URL: https://api.semanticscholar.org/CorpusID:197660727.

2 Dan Alistarh, Faith Ellen, and Joel Rybicki. Wait-free approximate agreement on graphs. In *Structural Information and Communication Complexity: 28th International Colloquium, SIROCCO 2021, Wrocław, Poland, June 28 – July 1, 2021, Proceedings*, pages 87–105, Berlin, Heidelberg, 2021. Springer-Verlag. doi:10.1007/978-3-030-79527-6_6.

3 Kenneth J. Arrow. *Social Choice and Individual Values.* John Wiley & Sons, 1951.

4 Hagit Attiya and Faith Ellen. The Step Complexity of Multidimensional Approximate Agreement. In Eshcar Hillel, Roberto Palmieri, and Etienne Rivière, editors, *26th International Conference on Principles of Distributed Systems (OPODIS 2022)*, volume 253 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 6:1–6:12, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.OPODIS.2022.6.

5 Hagit Attiya and Sergio Rajsbaum. Indistinguishability. *Commun. ACM*, 63(5):90–99, April 2020. doi:10.1145/3376902.

6 Y.M. Baryshnikov. Unifying impossibility theorems: A topological approach. *Advances in Applied Mathematics*, 14(4):404–415, 1993. doi:10.1006/aama.1993.1020.

7 Yuliy Baryshnikov and Joseph Root. A topological proof of the gibbard–satterthwaite theorem. *Economics Letters*, 234:111447, 2024. doi:10.1016/j.econlet.2023.111447.

8 Donald E. Campbell and Jerry S. Kelly. Chapter 1 impossibility theorems in the arrovian framework. In *Handbook of Social Choice and Welfare*, volume 1 of *Handbook of Social Choice and Welfare*, pages 35–94. Elsevier, 2002. doi:10.1016/S1574-0110(02)80005-4.

9 Soma Chaudhuri. More choices allow more faults: set consensus problems in totally asynchronous systems. *Inf. Comput.*, 105(1):132–158, July 1993. doi:10.1006/inco.1993.1043.

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

11 Himanshu Chauhan and Vijay K. Garg. Democratic elections in faulty distributed systems. In Davide Frey, Michel Raynal, Saswati Sarkar, Rudrapatna K. Shyamasundar, and Prasun Sinha, editors, *Distributed Computing and Networking*, pages 176–191, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. doi:10.1007/978-3-642-35668-1_13.

12 Carole Delporte-Gallet, Hugues Fauconnier, Eli Gafni, and Sergio Rajsbaum. Black art: Obstruction-free k-set agreement with |mwmr registers| < |proccesses|. In Vincent Gramoli and Rachid Guerraoui, editors, *Networked Systems*, pages 28–41, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. doi:10.1007/978-3-642-40148-0_3.

13 Persi Diaconis and R. L. Graham. Spearman's Footrule as a Measure of Disarray. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(2):262–268, December 2018. doi:10.1111/j.2517-6161.1977.tb01624.x.

14 Danny Dolev, Nancy Lynch, Shlomit Pinter, Eugene Stark, and William Weihl. Reaching approximate agreement in the presence of faults. *Journal of the ACM*, 33(3):499–516, May 1986. doi:10.1145/5925.5931.

15 Mose Mizrahi Erbes and Roger Wattenhofer. Asynchronous approximate agreement with quadratic communication. *arXiv preprint*, 2024. doi:10.48550/arXiv.2408.05495.

16 A. Fekete. Asynchronous approximate agreement. In *Proceedings of the sixth annual ACM symposium on principles of distributed computing (PODC)*, pages 64–76, New York, NY, USA, 1987. doi:10.1145/41840.41846.

17 Frank Feys and Helle Hansen. Arrow's theorem through a fixpoint argument. *Electronic Proceedings in Theoretical Computer Science*, 297:175–188, July 2019. doi:10.4204/EPTCS.297.12.

18 Faith Fich and Eric Ruppert. Hundreds of impossibility results for distributed computing. *Distributed Computing*, 16(2):121–163, September 2003. doi:10.1007/s00446-003-0091-y.

19 Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985. doi:10.1145/3149.214121.

20 Pierre Fraigniaud, Ami Paz, and Sergio Rajsbaum. A speedup theorem for asynchronous computation with applications to consensus and approximate agreement. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC'22, pages 460–470, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519270.3538422.

21 Pierre Fraigniaud, Sergio Rajsbaum, Matthieu Roy, and Corentin Travers. The opinion number of set-agreement. In Marcos K. Aguilera, Leonardo Querzoni, and Marc Shapiro, editors, *Principles of Distributed Systems*, pages 155–170, Cham, 2014. Springer International Publishing. doi:10.1007/978-3-319-14472-6_11.

22 Matthias Függer and Thomas Nowak. Fast multidimensional asymptotic and approximate consensus. *ArXiv*, abs/1805.04923, 2018. arXiv:1805.04923.

23 Hugo Rincon Galeana, Sergio Rajsbaum, and Ulrich Schmid. Continuous Tasks and the Asynchronous Computability Theorem. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 73:1–73:27, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2022.73.

24 John Geanakoplos. Three brief proofs of arrow's impossibility theorem. *Economic Theory*, 26(1):211–215, July 2005. doi:10.1007/s00199-004-0556-7.

25 Diana Ghinea, Chen-Da Liu-Zhang, and Roger Wattenhofer. Multidimensional approximate agreement with asynchronous fallback. Cryptology ePrint Archive, Paper 2023/449, 2023. URL: https://eprint.iacr.org/2023/449.

26 Allan Gibbard. *Manipulation of voting schemes: A general result.* s.n., 1998.

27 Maurice Herlihy, D. N. Kozlov, and Sergio Rajsbaum. *Distributed computing through combinatorial topology.* Elsevier/Morgan Kaufmann, 2014.

28 M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938. URL: http://www.jstor.org/stable/2332226.

29 M.G. Kendall. *Rank Correlation Methods.* C. Griffin, 1948. URL: https://books.google.com/books?id=hiBMAAAAMAAJ.

30 Isaac Lara, Sergio Rajsbaum, and Armajac Raventós-Pujol. A generalization of arrow's impossibility theorem through combinatorial topology, July 2024. arXiv:2402.06024.

31 Jérémy Ledent. Brief announcement: Variants of approximate agreement on graphs and simplicial complexes. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, PODC'21, pages 427–430, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3465084.3467946.

32 Hammurabi Mendes. *Byzantine Computability and Combinatorial Topology.* PhD thesis, Brown University, USA, 2016. URL: https://cs.brown.edu/research/pubs/theses/phd/2016/mendes.hammurabi.pdf.

33 Hammurabi Mendes and Maurice Herlihy. Multidimensional approximate agreement in byzantine asynchronous systems. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 391–400. ACM, 2013. doi:10.1145/2488608.2488657.

34 Hammurabi Mendes and Maurice Herlihy. Tight bounds for connectivity and set agreement in byzantine synchronous systems. In Andréa W. Richa, editor, *31st International Symposium on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria*, volume 91 of *LIPIcs*, pages 35:1–35:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICS.DISC.2017.35.

35 Hammurabi Mendes, Maurice Herlihy, Nitin H. Vaidya, and Vijay K. Garg. Multidimensional agreement in byzantine systems. *Distributed Comput.*, 28(6):423–441, 2015. doi:10.1007/S00446-014-0240-5.

36 Hammurabi Mendes, Christine Tasson, and Maurice Herlihy. Distributed computability in byzantine asynchronous systems. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 704–713. ACM, 2014. doi:10.1145/2591796.2591853.

37 Achour Mostefaoui, Sergio Rajsbaum, and Michel Raynal. The combined power of conditions and failure detectors to solve asynchronous set agreement. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing*, PODC '05, pages 179–188, New York, NY, USA, 2005. Association for Computing Machinery. doi:10.1145/1073814.1073848.

38 Thomas Nowak and Joel Rybicki. Byzantine Approximate Agreement on Graphs. In Jukka Suomela, editor, *33rd International Symposium on Distributed Computing (DISC 2019)*, volume 146 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 29:1–29:17, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.DISC.2019.29.

39 R. De Prisco, D. Malkhi, and M. Reiter. On k-set consensus problems in asynchronous systems. *IEEE Transactions on Parallel & Distributed Systems*, 12(01):7–21, January 2001. doi:10.1109/71.899936.

40 Sergio Rajsbaum and Armajac Raventós-Pujol. A distributed combinatorial topology approach to arrow's impossibility theorem. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC'22, page 471–481, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519270.3538433.

41 Yasuhito Tanaka. On the equivalence of the arrow impossibility theorem and the brouwer fixed point theorem. *Applied Mathematics and Computation*, 172(2):1303–1314, 2006. Special issue for The Beijing-HK Scientific Computing Meetings. doi:10.1016/j.amc.2005.02.054.