The semitopology of heterogeneous consensus

Murdoch J. Gabbay
Heriot-Watt University, UK

Giuliano Losa
Stellar Development Foundation, USA

Abstract

An analysis of distributed consensus under heterogeneous agreement requirements reveals a novel mathematical structure which is closely related to event structures and topological spaces.

2012 ACM Subject Classification Mathematics of computing → Continuous functions; Theory of computation → Algebraic semantics; Theory of computation → Axiomatic semantics; Theory of computation → Distributed computing models

Keywords and phrases Topology, Semitopology, Event structures, Heterogeneous consensus, Permissionless/open blockchain systems

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

1 Introduction

Consider the consensus problem as traditionally presented [11]: n processes in a distributed system each propose an arbitrary value and must then arrive (except those that fail) at a *consensus* (i.e. an agreement) on one of them. Solving consensus matters, because it allows a distributed system to function and coordinate its actions.

Now consider an *open* system where participants do not know each other and may have different objectives. In this case, global agreement as per the traditional notion of consensus from [11] might not be relevant, or even desirable: instead, participants might wish to agree with one or more sets of trusted participants whom they care about or otherwise share a common objective with — and participants make independent decisions on whom to trust.

We call such a system *heterogeneous*. So what is a sensible definition of the consensus problem in the heterogeneous setting?

In this paper, we propose to model heterogeneous systems using the new notion of *semitopological space* and we propose to define the consensus problem as the problem of *computing a continuous function on the semitopological space*.

The difference between semitopology and topology is that in semitopologies we drop the requirement that intersections of open sets be open. We develop a theory of semitopologies, thus casting a new light on, and giving a (we would argue) very clear new language for discussions about, the essential distributed-computing problem of consensus. Notably:

1. Whereas topology often studies spaces with strong separability properties between points (like Hausdorff separability), in a semitopological space it seems interesting to study points that cannot be separated. We state and discuss a novel anti-separation axiom which we call *being intertwined* (see Definition 13 and Remark 14).
2. A semitopological space partitions itself naturally into a collection of disjoint sets which we call *topens* (for *transitive open set*; Definition 6 and Remark 12) on which values of continuous functions are strongly correlated. Thus semitopologies articulate, in a clear and familiar topological language, mathematical reasons that a heterogeneous consensus system is likely to self-partition into unanimous communities (Theorems 18 and 27).
3. A substantial body of topology-flavoured results can now be developed. See for example the *characterisation of topen sets* and the *two ways to build a closure from a point* as summarised in Theorems 18 and 27 and Remark 28 (see also Subsection 6.2).

42nd Conference on Very Important Topics (CVIT 2016).
Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:16
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

4. Although semitopologies are not inherently computational — a semitopology is just a set of points and open sets on those points — the definitions support a natural computational structure which we call a *witness function* (Definition 29(1)), which is related to event structures [20]. This gives open and closed sets, and the topens mentioned above, computational content in a way that we make mathematically precise (Propositions 42 and 36), culminating with a compactness result (Theorem 43).
5. We discuss connections with related work in Subsection 6.1 (notably: event structures, consensus tasks, algebraic topology, and fail-prone systems and quorum systems).

Finally, note that semitopology is *practically* motivated: it is in use since 2015 in the Stellar payments network [12], whose notion of Federated Byzantine Agreement System [16] is an example of semitopological space.

2 Semitopology

A semitopology is like a topology, minus the condition that the intersection of two open sets be an open set, and *continuity* can be identified with *consensus*:

▶ **Definition 1.** *A **semitopological space**, or just **semitopology**, is a pair* $(\mathsf{P}, \mathsf{Open}(\mathsf{P}))$ *of a nonempty set* P *of **points**, and a set* $\mathsf{Open}(\mathsf{P}) \subseteq pow(\mathsf{P})$ *of **open sets**, such that:*

1. $\varnothing \in \mathsf{Open}(\mathsf{P})$ *and* $\mathsf{P} \in \mathsf{Open}(\mathsf{P})$.
2. *If* $X \subseteq \mathsf{Open}(\mathsf{P})$ *then* $\bigcup X \in \mathsf{Open}(\mathsf{P})$.

We may write $\mathsf{Open}(\mathsf{P})$ *just as* Open*, if* P *is irrelevant or understood.*

▶ **Definition 2** (Continuity). *If* P *and* P' *are sets and* $f : \mathsf{P} \to \mathsf{P}'$ *and* $O' \subseteq \mathsf{P}'$ *then define the **inverse image** by* $f^{-1}(O') = \{p{\in}\mathsf{P} \mid f(p) \in O'\}$.

If $(\mathsf{P}, \mathsf{Open})$ *and* $(\mathsf{P}', \mathsf{Open}')$ *are semitopological spaces then call* $f : \mathsf{P} \to \mathsf{P}'$ ***continuous*** *when* $O' \in \mathsf{Open}'$ *implies* $f^{-1}(O') \in \mathsf{Open}$.

▶ Remark 3 (Continuity=consensus). We can identify consensus as the instance of continuity in which we map from a semitopology to a discrete semitopology of *values*.

To see why, consider a semitopology $(\mathsf{P}, \mathsf{Open})$ and view $p \in \mathsf{P}$ as *participants* and open neighbourhoods $p \in O \in \mathsf{Open}$ as **quorums** of p — that is, $p \in O \in \mathsf{Open}$ indicates that O is a set that p would be willing to agree or cooperate with. Give some set Val of **values** or **beliefs** the **discrete semitopology** such that $\{v\}$ is open for every $v \in \mathsf{Val}$ (Example 4(1)).

Then having consensus amongst the P regarding a suitable value Val can be identified with having a continuous function f from P to Val because:

- f assigns a value or belief to each $p \in \mathsf{P}$, and
- continuity asserts that for every value or belief $v \in \mathsf{Val}$, each $p \in f^{-1}(v)$ is contained in a (by continuity) *open* set $f^{-1}(v)$ of peers that it is willing to agree or cooperate with, and which (by f) agree with p that v.[1]

(We briefly discuss in Subsection 6.2 how one might set about computing such an f.)

▶ **Example 4.** Examples of semitopologies include:

[1] The astute reader may notice that we sweep some things under the rug. How do we compute these functions? See Subsection 6.2. What about failures and Byzantine participants? Well, our slogan 'continuity=consensus' is a simplification, though a constructive and useful one; e.g. Byzantine behaviour can be modelled with partiality or discontinuity. More in longer paper.

1. The **discrete** semitopology on nonempty P just takes $\mathsf{Open} = pow(\mathsf{P})$. We may silently treat $\mathbb{B} = \{\bot, \top\}$ as a discrete semitopological space.
   Any function from a discrete semitopology is continuous, and intuitively, participants only care to agree with themselves and nobody cares what anybody else thinks.
2. Take P to be any nonempty set. The **trivial** semitopology on P takes $\mathsf{Open} = \{\varnothing, \mathsf{P}\}$.
   Only constant functions are continous, and intuitively, participants want to agree with *everyone*; if someone objects, we do not have an open set nor a continuous function.
3. Let P be people in a town with one cinema and $O \in \mathsf{Open}$ the semitopology generated by groups of friends willing to coordinate to go see a movie together. Then Open describes the sets of people that can be found inside the cinema.
4. Take $\mathsf{P} = \{0, 1, \dots, 41\}$. The **supermajority** semitopology takes $\mathsf{Open} = \{O \subseteq \mathsf{P} \mid \#O \geq 28\}$. So an open set contains at least two-thirds of the points; $2/3$ participation is a typical threshold used for making progress in consensus algorithms.[2]
   The supermajority semitopology captures that consensus is reached when a clear $2/3$ majority of participants are in agreement. This is not a topology: that O and O' contain at least two-thirds of the points in P does not mean their intersection $O \cap O'$ does.
5. Let $O \subseteq \mathsf{P}$ be open when $O = \varnothing$ or $\#O = \#\mathsf{P}$ (e.g. if $\mathsf{P} = \mathbb{N}$ then $\{n \mid n \text{ even}\}$ and $\{n \mid n \text{ odd}\}$ are open). This *many semitopology* is not a topology.
6. Let $O \subseteq \mathsf{P}$ be open when $O = \varnothing$ or $O = \mathsf{P}$ or $O = \mathsf{P} \setminus \{p\}$ for any $p \in P$. Intuitively, in this *lone objector semitopology* (which is not a topology), participants are deemed to have reached consensus when there is at most one objector.
7. Consider any L-labelled automaton A (by which here we mean: a rooted directed graph with labels from L). Let P be finite (possibly empty) lists of elements from L and let a set be open when it is a union of sets of finite initial segments of an infinite path through A.
   To make this concrete: take A to have one node, and two edges labelled 0 and 1. Then $\{[], [0], [0,1], [0,1,0], [0,1,0,1], \dots\}$ is an open set, obtained as finite approximations to the path $0, 1, 0, 1, \dots$. In this semitopology, 'participants' are finite approximations, and a set is open when it is a union of sequences of participants, with each sequence appoximating some infinite limit.

3 Transitive sets and (maximal) topen sets

▶ **Definition 5.** *Suppose X, Y, and Z are sets. Write $X \between Y$ and say that X and Y* ***intersect*** *when $X \cap Y \neq \varnothing$. We may chain $\between$, writing e.g. $X \between Y \between Z$ for $X \between Y \wedge Y \between Z$.*

▶ **Definition 6.** *Suppose $(\mathsf{P}, \mathsf{Open})$ is a semitopology and $S \subseteq \mathsf{P}$. Call S* ***transitive*** *when $\forall O, O' {\in} \mathsf{Open}. O \between S \between O' \Longrightarrow O \between O'$ and call S a* ***(maximal) topen*** *when S is a (maximal) nonempty open and transitive set.*[3]

Values of continuous functions are strongly correlated on transitive sets (thus topens):

▶ **Proposition 7.** *Suppose $(\mathsf{P}, \mathsf{Open})$ is a semitopology and Val is a set of* values *(e.g. $\mathbb{B}$ or $\mathbb{N}$) with the discrete semitopology (Example 4(1)), and suppose $f : \mathsf{P} \to \mathsf{Val}$ is continuous (Definition 2) and $S \subseteq \mathsf{P}$ is transitive (usually, S will be topen). Then f is constant on S.*

Proof. Suppose $p, p' \in S$ and write $v = f(p)$ and $v' = f(p')$. By construction $f^{\text{-}1}(v) \between S \between f^{\text{-}1}(v')$. Therefore $f^{\text{-}1}(v) \between f^{\text{-}1}(v')$, by transitivity of S. This means precisely that there exists p'' such that $v = f(p'') = v'$, and so $v = v'$. ◀

[2] The procedural threshold in the US Senate is often set to two-thirds of the Senators present and voting.
[3] 'Transitive open' → 'topen', like 'closed and open' → 'clopen'.

▶ **Example 8.** **1.** $\{p\}$ and $\varnothing$ are (trivially) transitive, for any $p \in \mathsf{P}$.
2. If $S \subseteq \mathsf{P}$ is *topologically indistinguishable* ($\forall O{\in}\mathsf{Open}.(S \between O \Longrightarrow S{\subseteq}O)$) then S is transitive.
3. Take $\mathsf{P} = \{0, 1, 2, \dots\}$ and let open sets be $\varnothing$, P, and sets $O_n = \{n \times i \mid i \geq 1\}$ for every $n \geq 1$. This has one maximal topen $O_1 = \{1, 2, \dots\}$, and one isolated point 0.
4. Take $\mathsf{P} = \{\text{-}1, 0, 1\}$, with open sets $\varnothing$, P, $\{\text{-}1\}$, $\{\text{-}1, 0\}$, $\{1\}$, $\{0, 1\}$, and $\{\text{-}1, 1\}$. This has two maximal topens $\{\text{-}1\}$ and $\{1\}$, and 0 is not in any topens.

▶ **Lemma 9.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology.*
1. *If* $S, S' \subseteq \mathsf{P}$ *are topen then* $\forall O, O' \in \mathsf{Open}.O \between S \between S' \between O' \Longrightarrow O \between O'$.
2. *If* $\mathcal{S}$ *is a set of topens that are pairwise intersecting (so* $\forall S, S'{\in}\mathcal{S}.S{\between}S'$*) then* $\bigcup \mathcal{S}$ *is topen.*

Proof. **1.** We simplify using Definition 6:

$$\begin{aligned} O \between S \between S' \between O' &\Longrightarrow O \between S' \between O' && S \text{ transitive, } S' \text{ open} \\ &\Longrightarrow O \between O' && S' \text{ transitive.} \end{aligned}$$

2. $\bigcup \mathcal{S}$ is open by Definition 1(2). Also, if $O \between \bigcup \mathcal{S} \between O'$ then there exist $S, S' \in \mathcal{S}$ such that $O \between S$ and $S' \between O'$. We assumed $S \between S'$, so by part 1 of this result we have $O \between O'$. ◀

▶ Remark 10. We care about topens (rather than sets that are just transitive) because they have somewhat better closure properties. E.g. Lemma 9 fails for transitive sets in general: if $\mathsf{P} = \{1, 2, 3\}$ and $\mathsf{Open} = \{\emptyset, \mathsf{P}, \{2\}, \{3\}, \{2, 3\}\}$ then $\{1, 2\}$ and $\{1, 3\}$ are transitive, but their union $\{1, 2, 3\}$ is not. There is fine structure to these results, which we will document in a longer paper.

▶ **Corollary 11.** *If* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology then every topen* $S \subseteq \mathsf{P}$ *is contained in a unique maximal topen* $M \supseteq S$.

Proof. Consider $\mathcal{S} = \{S \cup S' \mid S' \text{ topen} \wedge S \between S'\}$. By Lemma 9(2) this is a set of topens and by Lemma 9(2) again so is $\bigcup \mathcal{S}$. It is easy to check that this is a unique maximal transitive open set that contains S. ◀

▶ Remark 12. We see from Corollary 11 above that a semitopology $(\mathsf{P}, \mathsf{Open})$ naturally partitions itself into some disjoint collection of maximal topens, and other points not contained in any topen.[4]

Combining this with Proposition 7 we see that consensus on a semitopology self-organises into partitions of strongly correlated points acting together, along with some isolated points. In the special case of a space that is a single finite topen, then all participants must agree.

▶ **Definition 13.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $p, p' \in \mathsf{P}$.
1. *Call* p *and* p' ***intertwined*** *when* $\{p, p'\}$ *is transitive. Unpacking Definition 6 this means* $\forall O, O'{\in}\mathsf{Open}.(p \in O \wedge p' \in O') \Longrightarrow O \between O'$. *By a mild abuse of notation, write* $p \between p'$ *when* p *and* p' *are intertwined.*
2. *Define* $p_{\between} = \{p' \in \mathsf{P} \mid p \between p'\}$. *So* $p_{\between}$ *is the points intertwined with* p.

▶ Remark 14. The reader can check that the usual Hausdorff separation axiom can be succinctly written as $\forall p.p_{\between} = \{p\}$. Conversely, $p \between p'$ for $p \neq p'$ is the very *opposite* to being Hausdorff: that p and p' they *cannot* be separated by pairwise disjoint open sets.[5]

---

[4] This raises the question of what those other points can look like topologically. One answer is implicit in Theorem 18, if we consider the topological boundary of a maximal topen. Or, a point can simply be isolated. See Example 8, items 3 and 4. A more detailed analysis is possible but out of scope here.
[5] One might call this an *anti-Hausdorff* property.

For semitopologies as applied to consensus, Hausdorff makes a space separated and liable to non-consensus. Conversely, to maximise consensus and minimise separation — the literature might call this *avoiding forking* — we may prefer a space to be very intertwined.

▶ **Lemma 15.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $S \subseteq \mathsf{P}$*. Then* S *is transitive if and only if* $\forall p, p' \in S. p \between p'$*. In words:* a set is transitive when it is pointwise intertwined.

▶ **Corollary 16.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $S \subseteq \mathsf{P}$*. Then the following assertions are equivalent:*
1. S *is topen.*
2. S *is nonempty, open, and* $p \between p'$ *for every* $p, p' \in S$*.*
3. S *is nonempty, open, and* $S \subseteq p_{\between}$*, for some/any element* $p \in S$*.*

In words: A topen set is a nonempty open set of intertwined points.

Proof. By Definition 6, S is topen when it is nonempty, open, and transitive. By Lemma 15 this last condition is equivalent to $p \between p'$ for every $p, p' \in S$. Thus parts 1 and 2 are equivalent. By Definition 13(2) $p_{\between} = \{p' \mid p \between p'\}$, so part 3 just rephrases part 2. ◀

▶ **Definition 17.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $R \subseteq \mathsf{P}$*. Define the* ***interior*** *of* R *by* $\mathit{interior}(R) = \bigcup\{O \in \mathsf{Open} \mid O \subseteq R\}$*.*

▶ **Theorem 18** (Characterisation of topens)**.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $S \subseteq \mathsf{P}$*. Then the following are equivalent:*
- S *is a* maximal *topen.*
- S *is nonempty and* $S = \mathit{interior}(p_{\between})$ *for some/any element* $p \in S$*.*

In words: A maximal topen is the nonempty open interior of $p_{\between}$.

Proof. From Corollary 16 using Definition 6. ◀

4 Closed sets and interiors

4.1 Basic definitions (TL;DR: this part is just like topology)

▶ **Definition 19.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $p \in \mathsf{P}$ *and* $R \subseteq \mathsf{P}$*. Define the* ***closure*** *of* R *by* $|R| = \{p' \in \mathsf{P} \mid \forall O' \in \mathsf{Open}. p' \in O' \Longrightarrow R \between O'\}$*.*

We may write $|p|$ *for* $|\{p\}|$*, so* $|p| = \{p' \in \mathsf{P} \mid \forall O' \in \mathsf{Open}. p' \in O' \Longrightarrow p \in O'\}$*.*

Call C ***closed*** *when* $C = |C|$*, and write* $\mathit{Closed}(\mathsf{P})$ *for the set of closed sets.*

Closed sets are complements of open sets, and open/closed sets are interiors/closures of closed/open sets — just like in topologies. We check that this works as expected in Lemma 20, Corollary 21, and Lemma 22:

▶ **Lemma 20.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology. Then* $C \in \mathit{Closed}(\mathsf{P})$ *is closed if and only if* $\mathsf{P} \setminus C$ *is open, and* $O \in \mathsf{Open}$ *is open if and only if* $\mathsf{P} \setminus O$ *is closed.*

Proof. 1. Suppose $p \in \mathsf{P} \setminus C$. Since $C = |C|$, $p \in \mathsf{P} \setminus |C|$. From Definition 19 there exists $O \in \mathcal{P}$ with $p \in O$ and $O \cap C = \varnothing$, and this is the openness condition from Definition 30.
2. Suppose $O \in \mathsf{Open}$. It follows from Definition 19 that $O \cap |\mathsf{P} \setminus O| = \varnothing$. But (as can be checked from routine calculations) $\mathsf{P} \setminus O \subseteq |\mathsf{P} \setminus O|$. ◀

▶ **Corollary 21.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology. Then* $\mathit{Closed}(P)$ *contains* $\varnothing$ *and* P *and is closed under arbitrary intersections. Furthermore,* $|R|$ *equals the intersection of the closed sets that contain it:* $|R| = \bigcap\{C \in \mathsf{Closed} \mid R \subseteq C\}$*.*

Proof. The first assertion is immediate from Lemma 20. The second follows from Lemma 20 and Definition 1(1&2). The third assertion follows from the second. ◀

▶ **Lemma 22.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology. Then* $O \in \mathsf{Open}$ *is open if and only if* $interior(|O|) = O$*, and* $C \in \mathsf{Closed}$ *is closed if and only if* $|interior(C)| = C$*.*

Proof. Routine from Definitions 17 and 19. ◀

4.2 Intertwined elements and topens

Recall from Definition 13 the notions of $p \between p'$ and $p_{\between}$.

▶ **Lemma 23.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $p, p' \in \mathsf{P}$*. Then:*

1. $p \between p'$ *when* $\forall O{\in}\mathsf{Open}.p \in O \Longrightarrow p' \in |O|$*.*
2. $p_{\between} = \bigcap\{|O| \mid p \in O \in \mathsf{Open}\} = \bigcap\{C \in \mathsf{Closed} \mid p \in interior(C)\}$*.*
3. $p_{\between}$ *is closed and (by Lemma 20)* $\mathsf{P} \setminus p_{\between}$ *is open.*

Proof. **1.** We just rearrange Definition 13:

$$\begin{array}{ll} \forall O, O' \in \mathsf{Open}.(p \in O \wedge p' \in O') \Longrightarrow O \between O' & \text{rearranges to} \\ \forall O \in \mathsf{Open}.p \in O \Longrightarrow \forall O' \in \mathsf{Open}.p' \in O' \Longrightarrow O \between O' & \text{and by Definition 19 this is} \\ \forall O \in \mathsf{Open}.p \in O \Longrightarrow p' \in |O|. & \end{array}$$

2. Using part 1, then Lemma 22.
3. We combine part 2 of this result with Corollary 21. ◀

▶ **Lemma 24.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $S \subseteq \mathsf{P}$ *is topen and* $p \in S$*. Then* $|p| \subseteq p_{\between} = |S|$*, and the subset inclusion may be strict (that is,* $|p| \neq p_{\between}$ *is possible).*

Proof. $p_{\between} = |S|$ follows using Theorem 18. For the subset inclusion, by Corollary 21(3) $|p| = |\{p\}| = \bigcap\{C \in \mathsf{Closed} \mid p \in C\}$ and also by Lemma 23(2) $p_{\between} = \bigcap\{C \in \mathsf{Closed} \mid p \in interior(C)\}$. We note that if $p \in interior(C)$ then $p \in C$, so that

$$\{C \mid p \in interior(C)\} \subseteq \{C \mid p \in C\} \quad \text{and so} \quad \bigcap\{C \mid p \in interior(C)\} \supseteq \bigcap\{C \mid p \in C\}.$$

Example 25 shows that $|p| \neq p_{\between}$ can hold: ◀

▶ **Example 25.** Take $\mathsf{P} = \{0, 1\}$ and $\mathsf{Open} = \{\varnothing, \{0\}, \{0, 1\}\}$. Then $|1| \neq 1_{\between}$:

- $|1| = \{1\}$ ($\{0\}$ is open), but
- $1_{\between} = \{0, 1\}$ (open neighbourhoods of 0 and 1 all intersect).

Lemma 26 complements Lemma 24:

▶ **Lemma 26.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $S \subseteq \mathsf{P}$ *is topen and* $p \in S$ *and* $O \subseteq S$*. Then* $|O| = p_{\between} = |S|$*.*

Proof. $p_{\between} = |S|$ follows using Theorem 18. We now consider the left-hand equality. Unpacking Definition 19 we have:

$$\begin{array}{ll} p \in |O| \Longleftrightarrow \forall O'{\in}\mathsf{Open}.p \in O' \Longrightarrow O' \between O & \text{and} \\ p \in |S| \Longleftrightarrow \forall O'{\in}\mathsf{Open}.p \in O' \Longrightarrow O' \between S. & \end{array}$$

We see that it suffices to prove $O' \between O \Longleftrightarrow O' \between S$ for any $O' \in \mathsf{Open}$. But this is routine:

- Suppose $O' \between S$. By assumption $S \between O$ and by transitivity of S (Definition 6) $O' \between O$.
- Suppose $O' \between O$. By assumption $O \subseteq S$, and $O' \between S$ follows. ◀

▶ **Theorem 27.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a semitopology and* $p \in \mathsf{P}$*. Then:*

1. $\mathsf{P} \setminus |p| = \bigcup\{O \in \mathsf{Open} \mid p \notin O\}$*, and this is an open set.*
2. $\mathsf{P} \setminus p_{\between} = \bigcup\{C \in \mathsf{Closed} \mid p \notin C\}$*, and this is an open set.*
3. $\mathsf{P} \setminus p_{\between} \subseteq \mathsf{P} \setminus |p|$ *and the inclusion may be strict.*
4. *If* $interior(|p|) \neq \varnothing$ *(so* $|p|$ *has a nonempty open interior) then* $|p| = p_{\between}$ *and* $\mathsf{P}\setminus|p| = \mathsf{P}\setminus p_{\between}$*.*

Proof. **1.** Immediate from Definition 19.[6] Openness is from Definition 1(2).

2. We reason using Lemma 23(2): $\mathsf{P} \setminus p_{\between} = \bigcup\{\mathsf{P} \setminus |O| \mid p \in O\} = \bigcup\{C \in \mathsf{Closed} \mid p \notin C\}$. Openness is Lemma 23(3).

3. From Lemma 24.

4. From Lemma 26. ◀

▶ Remark 28 (Summary).
- $|p|$ is a closed set and is the *closure* of p (Definition 19: the p' whose every open neighbourhood $p' \in O'$ intersects with $\{p\}$).
- $\mathsf{P} \setminus |p|$ is the union of *open* sets that avoid p.
- $p_{\between}$ is also a closed set and is the points *intertwined* with p (Definition 13(2): the p' whose every open neighbourhood $p' \in O'$ intersects with every open neighbourhood $p \in O$).
- $\mathsf{P} \setminus p_{\between}$ is the union of the *closed* sets that avoid p.
- The open interior of $p_{\between}$, if non-empty, is a topen (as studied above, culminating with Theorem 27), and $p_{\between}$ is equal to the closure of any nonempty open that it contains.

Thus we have *two* ways to build a closed set from $p \in \mathsf{P}$: from its closure $|p|$ (Definition 19) which is the set of points all of whose open neighbourhoods contain p; or we can take $p_{\between}$ (Definition 13(2)) which is the set of points that are intertwined with p, which is closed by Lemma 23(3). Furthermore: $|p| \subseteq p_{\between}$ and the reverse inclusion holds if $|p|$ has an open interior (Lemmas 24 and 26); and $interior(p_{\between})$ is a maximal topen if it is nonempty (Theorem 18).

5 The witness function: computable semitopologies

Let us recap: semitpologies are topologies without the restriction that the intersection of two opens be open; notions of continuity and closure carry over from topology and continuity = consensus; we note an anti-Hausdorff property which we call *being intertwined*; we characterise open interiors of maximal sets of intertwined sets as *maximal topens* which partition the space into blocks of consensus, in the formal sense that values of continuous functions are strongly correlated on each topen. This is descriptively nice, but is it compatible with algorithmic content? We consider this next.

5.1 The witness function

Write $pow(X)$ for the powerset of X, and $pow_{\neq\varnothing}(X)$ for the nonempty powerset of X, and $fin(X)$ for the finite powerset of X, and $fin_{\neq\varnothing}(X)$ for the nonempty finite powerset of X.

▶ **Definition 29.** *Suppose* P *is a set. Then:*

1. *A* ***witness function*** *on* P *is a function* $W : \mathsf{P} \to fin_{\neq\varnothing}(pow_{\neq\varnothing}(\mathsf{P}))$*. Call a pair* (P, W) *of a set and a witness function on that set, a* ***witnessed set****.*
2. *If* (P, W) *is a witnessed set and* $p \in \mathsf{P}$ *then call* $w \in W(p)$ *a* ***witness*** *for* p *and say that* w ***witnesses*** p*.*

[6] A longer proof via Corollary 21(3) and Lemma 20 is also possible.

In words: a witnessed set is a set along with a witness function that assigns to each element of that set a nonempty finite set of nonempty (possibly infinite) witnesses.

▶ **Definition 30.** *Suppose* (P, W) *is a witnessed set. Define the* ***witness semitopology*** *by*

$$O \in \mathsf{Open}(W) \quad \textit{when} \quad \forall p \in \mathsf{P}.\big(p \in O \Longrightarrow \exists w{\in}W(p).w \subseteq O\big).$$

▶ Remark 31. Witness functions matter because they yield semitopologies with computational meaning, as we shall see. But before we jump into the details, we pause to reflect on how witness functions can be interpreted:

1. *Consensus interpretation:* W represents groups of 'immediate friends': $w \in W(p)$ is a group of elements that p personally trusts. An open set O is a community of participants such that every $p \in O$ is accompanied by *some* group of immediate friends.
2. *Computational interpretation:* W represents a nondeterministic parallel computation. Each $p \in \mathsf{P}$ is a process and $w \in W(p)$ a parallel computation to which p can nondeterministically evolve. An open set O is a computation trace in which each p is accompanied by (at least one) choice of evolution $w \in W(p)$.
   Example 4(7) illustrates this: e.g. for $p \in \{0,1\}^*$ set $W(p) = \{\{p+0\}, \{p+1\}\}$. Thus p computes its next step by evolving either to $p+0$ or $p+1$, and open sets are generated by computations of infinite streams (this example has nondeterminism but no parallelism).
3. *Modal / event structures interpretation:* W is an *enabling modality.* Each $p \in \mathsf{P}$ is an event and $w \in W(p)$ is a combination of events that enable p to be possible. An open set O is a computation trace in which each $p \in O$ is enabled by at least one $w \in W(p)$.

▶ Remark 32. Continuing the modal interpretation above, a witnessed set (P, W) from Definition 29 is an infinitary generalisation of a special case of an *event structure* [20, Definition 1.1.1] — namely, it is an event structure in which the witness function plays the role of the enabling relation, and sets of events are generalised so they may be infinite (and the consistency predicate is trivial).[7]

This does not make semitopologies a special case of event structures; not only because of the infinitary generalisation, but because we take the definitions in a new direction. It is an exciting possibility for future work to use this new maths to transfer ideas and algorithms between the two worlds— e.g. minimisation algorithms or bisimulation properties from event structures, or concrete algorithmics and applications from Stellar.

5.2 The witness function and open sets

▶ **Definition 33.** *Suppose that* (P, W) *is a witnessed set (Definition 29) and* $X, X' \subseteq \mathsf{P}$*. Define the* ***witness (partial) ordering*** *by* $X \preceq X'$ *when* $X \subseteq X' \ \wedge\ \forall p{\in}X.\exists w{\in}W(p).w \subseteq X'$*. If* $X \preceq X$ *then call* X *a* ***$\preceq$-fixedpoint****.*

In words: $X \preceq X'$ *when* X' *extends* X *with (at least) one witnesses for every* $p \in X$*.*

▶ **Lemma 34.** $\preceq$ *is indeed a partial order (transitive possibly irreflexive relation), and* $\preceq\ \subseteq\ \subseteq$*.*

▶ **Lemma 35.** *Suppose* (P, W) *is a witnessed set. Then* O *is open in the witness semitopology (Definition 30) if and only if* O *is a* $\preceq$*-fixedpoint. In symbols:* $\mathsf{Open} = \{X \subseteq \mathsf{P} \mid X \preceq X\}$*.*

Proof. Being a $\preceq$-fixedpoint from Definition 33 — every point in O is witnessed by a subset of O — reformulates the openness condition of the witness semitopology from Definition 30. ◀

---

[7] A clear overview is online at `https://depend.cs.uni-saarland.de/fileadmin/user_upload/depend/neuhaeusser/concurrency_seminar_2011/event_structures.pdf`; see in particular Definition 9.

▶ **Proposition 36.** *Suppose* (P, W) *is a witnessed set and suppose* $\mathcal{X} = (X_0 \preceq X_1 \preceq \dots)$ *is a countably ascending* $\preceq$*-chain. Write* $\bigcup \mathcal{X}$ *for the sets union* $\bigcup_i X_i$*. Then:*

1. $\bigcup \mathcal{X}$ *is a* $\preceq$*-limit for* $\mathcal{X}$*. In symbols:* $\forall i.X_i \preceq \bigcup \mathcal{X}$*.*
2. $\bigcup \mathcal{X}$ *is a* $\preceq$*-fixedpoint. In symbols:* $\bigcup \mathcal{X} \preceq \bigcup \mathcal{X}$*.*

Proof. 1. We must show that if $p \in X_i$ then $w \subseteq \bigcup \mathcal{X}$ for some $w \in W(p)$. But this is automatic from the fact that $X_i \preceq X_{i+1} \subseteq \bigcup \mathcal{X}$.
2. From part 1 noting that if $p \in \bigcup \mathcal{X}$ then $p \in X_i$ for some i. ◀

Proposition 36 and Lemma 35 above are not complicated (note that this is a feature, which required conscious design effort) and they say something important: in the *witness* semitopology, open sets can be computed using a simple iterative algorithm which we can sum up as 'just iteratively add witnesses'.

5.3 The witness function and closed sets

▶ **Definition 37.** *Suppose* R *is a set and* $\mathcal{W}$ *is a set (or a sequence) of sets. Define* $R \between \mathcal{W}$ *when* $\forall W{\in}\mathcal{W}.R \between W$*. In words:* $R \between \mathcal{W}$ when R intersects with every $W \in \mathcal{W}$.

▶ **Definition 38.** *Suppose* (P, W) *is a witnessed set and* $R \subseteq \mathsf{P}$*. Define* $lim_w(R)$ *by*

$$lim_w(R) = R \cup \{p \in \mathsf{P} \mid R \between W(p)\}.$$

In words: $lim_w(R)$ is the set of points p whose every witness contains an R-element.
We iterate this

$$\begin{aligned} lim_0(R) &= R \\ lim_{i+1}(R) &= lim_w(lim_i(R)) \\ lim(R) &= \textstyle\bigcup_{n \geq 0} lim_n(R) \end{aligned}$$

and we call $lim(R)$ *the set of* ***limit points*** *of* R*.*

▶ **Lemma 39.** *Suppose* (P, W) *is a witnessed set and* $R \subseteq \mathsf{P}$*. Then* $R \subseteq lim(R)$*.*

Proof. It is a fact of Definition 38 that $R = lim_0(R) \subseteq lim_1(R) \subseteq lim(R)$. ◀

▶ **Lemma 40.** *Suppose* (P, W) *is a witnessed set and* $p \in \mathsf{P}$ *and* $R \subseteq \mathsf{P}$*. Then:*

1. *If* $lim(R) \between W(p)$ *(Definition 37) then* $p \in lim(R)$*.*
2. *By the contrapositive and expanding Definition 37, if* $p \in \mathsf{P} \setminus lim(R)$ *then* $\exists w{\in}W(p).w \cap lim(R) = \varnothing$*.*

Proof. Suppose $lim(R) \between W(p)$. Unpacking Definitions 37 and 38 it follows that for every $w{\in}W(p)$ there exists $n_w \geq 0$ such that $lim_{n_w}(R) \between w$. Now by Definition 29(1) $W(p)$ is finite, and it follows that for some/any n greater than the maximum of all the n_w, we have $lim_n(R) \between W(p)$. Thus $p \in lim_w(lim_n(R)) \subseteq lim(R)$. ◀

▶ **Lemma 41.** *Suppose* (P, W) *is a witnessed set and* $p \in \mathsf{P}$ *and* $R \subseteq \mathsf{P}$ *and* $O \in \mathsf{Open}$*:*

1. *If* $O \between lim_w(R)$ *then* $O \between R$*.*
2. *If* $O \between lim(R)$ *then* $O \between R$*.*
3. *As a corollary, if* $O \cap R = \varnothing$ *then* $O \cap lim(R) = \varnothing$*.*

Proof. 1. Consider $p \in \mathsf{P}$ such that $p \in O$ and $p \in lim_w(R)$. By assumption there exists $w \in W(p)$ such that $w \subseteq O$. Also by assumption $w \between R$. It follows that $O \between R$.
2. If $O \between lim(R)$ then $O \between lim_n(R)$ for some finite $n \geq 0$. By a routine induction using part 1 of this result, it follows that $O \between R$.

3. This is just the contrapositive of part 2 of this result, noting that $O \between R$ when $O \cap R = \varnothing$ by Definition 5, and similarly for $O \between \mathit{lim}(R)$. ◀

▶ **Proposition 42.** *Suppose* (P, W) *is a witnessed set and* $R \subseteq \mathsf{P}$*. Then* $\mathit{lim}(R) = |R|$*. In words:* the set of limit points of R (Definition 38) equals the closure of R (Definition 19).

Proof.
- *Suppose* $p \notin |R|$*.* Then there exists some $p \in O \in \mathsf{Open}$ such that $O \cap R = \varnothing$. Thus by Lemma 41(3) also $O \cap \mathit{lim}(R) = \varnothing$.
- *Suppose* $p \notin \mathit{lim}(R)$*.* By Definition 19 we need to exhibit an $p \in O \in \mathsf{Open}$ that is disjoint from R, and since $R \subseteq \mathit{lim}(R)$ by Lemma 39, it would suffice to exhibit $p \in O \in \mathsf{Open}$ that is disjoint from $\mathit{lim}(R)$. We set $O = \mathsf{P} \setminus \mathit{lim}(R)$. Lemma 40(2) expresses that this is an open set, and by construction it is disjoint from $\mathit{lim}(R)$. ◀

Proposition 42 above does for closed sets as Proposition 36 and Lemma 35 do for open sets: in the *witness* semitopology, closed sets can be computed using an iterative algorithm which we can sum up as 'iteratively add points all of whose witnesses intersect'.

5.4 Compactness of descending chains of open sets

Intuitively, 'compactness' is used to indicate 'generalising finiteness'. Theorem 43 is a remarkable property, that a descending chain of open sets behaves as if it were finite — even though it isn't:[8]

▶ **Theorem 43** (Compactness of descending chains)**.** *Suppose that:*
- (P, W) *is a witnessed set with the witness semitopology (Definition 30).*
- $\alpha \geq 1$ *is a nonzero ordinal.*
- $\mathcal{O} \subseteq \mathsf{Open}$ *is a descending* α*-chain of open sets.*[9]

Then

$$\bigcap \mathcal{O} \in \mathsf{Open}.$$

In words: in a witness semitopology, the intersection of a descending chain of open sets, is an open set.

Proof.
- *Suppose* $\bigcap \mathcal{O} = \varnothing$*.* Then we note that $\varnothing \in \mathsf{Open}$ (Definition 1(1)) and we are done.
- *Suppose* $\alpha = \alpha' + 1$*, so that* α *is a successor ordinal.* Then the sequence $\mathcal{O}$ has a final element O_α and by facts of sets $\bigcap \mathcal{O} = O_\alpha \in \mathsf{Open}$ and again we are done.
- *Suppose* α *is a limit ordinal and* $\bigcap \mathcal{O} \neq \varnothing$*.* Consider some $p \in \bigcap \mathcal{O}$. By construction of the witness semitopology (Definition 30) for each O_i there exists a witness $w_i \in W(p)$ such that $w_i \subseteq O_i$. Now by Definition 29(1) $W(p)$ is finite, so by the pigeonhole principle, there exists some $w \in W(p)$ such that $w = w_i$ for infinitely many $w_i \in W(p) \wedge w_i \subseteq O_i$. It follows from the fact that $\mathcal{O}$ is a descending chain that $w \subseteq \bigcap \mathcal{O}$.

[8] One might be tempted to call this property *Noetherian*, since it has to do with a descending chain having a terminator, but to us that seems not right: 'Noetherian' means 'well-founded', but infinite descending chains of open sets are possible in a witness semitopology — they just have an open intersection. Note also that this result says something strictly stronger than 'every descending chain of open sets has a greatest lower bound'; the point is that this greatest lower bound is the sets intersection on-the-nose.

[9] ... an α-indexed chain of sets such that $O_{\alpha'} \subseteq O_{\alpha''}$ for every $0 \leq \alpha'' < \alpha' < \alpha$.

Now p in the previous paragraph was arbitrary, so we have shown that if $p \in \bigcap \mathcal{O}$ then also there exists $w \in W(p)$ such that $w \subseteq \bigcap \mathcal{O}$. It follows by construction of the witness semitopology in Definition 30 that $\bigcap \mathcal{O}$ is open as required. ◀

▶ **Corollary 44.** *Suppose* (P, W) *is a witnessed set with the witness semitopology. Then:*

1. *Any nonempty* $\mathcal{O} \subseteq \mathsf{Open}$ *contains a* $\subseteq$*-minimal element.*
2. *If* $p \in O \in \mathsf{Open}$ *then* $\{O' \in \mathsf{Open} \mid p \in O' \subseteq O\}$ *has a* $\subseteq$*-minimal element, which we may call a* ***cover*** *of* p*.*
3. *If* $p \in \mathsf{P}$ *then* $\{O \in \mathsf{Open} \mid p \in O\}$ *has a minimal element.*

Proof. Parts 2 and 3 are immediate from part 1 (noting that the sets are nonempty because they contain O and P respectively). For part 1 we reason as follows: It follows from Theorem 43 that $\mathcal{O}$, ordered by the *superset* relation (the reverse of the subset inclusion relation) contains limits, and so upper bounds, of ascending chains. By Zorn's lemma [10, 4] $\mathcal{O}$ contains a $\supseteq$-maximal element, and this is the required $\subseteq$-minimal element. ◀

▶ Remark 45. It would be nice if the reverse implication in Theorem 43 held but we suspect that there may exist $(\mathsf{P}, \mathsf{Open})$ in which every descending chain of open sets has an open intersection, yet it is not obtainable as a witness semitopology. To see why, consider Example 4(6) and take $X = \mathbb{N}$; then opens have the form $\varnothing$ or $\mathbb{N}$ or $\mathbb{N} \setminus \{n\}$. An infinite set of witnesses to n is $\mathbb{N} \setminus \{n'\}$ for $n' \neq n$, but this is not finite as required in Definition 29. This example or one like it might be used for a *proof* of non-existence, in future work.

We could allow an infinite set of witnesses in Definition 29, but at a price:

- Theorem 43 depends on the pigeonhole principle, which uses finiteness of the witness set.
- Lemma 40 depends on witness sets being finite, and this is required for Proposition 42.

▶ Remark 46. Recall that in our semitopological analysis consensus is continuity, and continuity means that preimages of open sets are open. Thus to understand consensus near a point p, we need to know what the open sets containing p look like; call this informally the *consensus neighbourhood* of p.

Theorem 43 and Corollary 44 above have specific mathematical meaning — but they also tell us something more general: that in a witness semitopology, we can understand the structure of consensus at a point $p \in \mathsf{P}$ by understanding the structure of its open covers, where a **cover** is a minimal set containing p. This is because if a continuous function $f : \mathsf{P} \to \mathsf{P}'$ such that $f(p) = p' \in O'$ is continuous at $p \in \mathsf{P}$, then certainly there exists some open cover $p \in O \subseteq f^{-1}(O')$. Turning this around: if we want to *create* consensus around p (e.g. because we are designing a consensus algorithm) it suffices to find some open cover of p, and convince that cover.

▶ Remark 47 (Computational content of witnessed sets). A semitopology from Definition 1 is just some points and some open sets. This in and of itself carries no computational structure, and a simple example illustrates this point. Let the **uncomputable semitopology** have $\mathsf{P} = \mathbb{N}$ and have open sets generated as unions of *uncomputable subsets* of $\mathbb{N}$. This is a semitopology and by design it is uncomputable. It is not a topology since the intersection of two uncomputable subsets need not be uncomputable.

Witnessed sets (Definition 29) make semitopologies computationally tractable, in the sense made formal by Propositions 36 and 42 (which show that algorithms exist to compute open and closed sets) and by the remarkable Theorem 43 (which shows intuitively that witness semitopologies behave locally like finite sets, even if they are globally infinite).

▶ Remark 48. We take a moment to unpack the algorithms implicit in Propositions 36 and 42. Consider a witnessed set (P, W). Then:

- To compute an open set in the witness semitopology, pick some $p \in \mathsf{P}$ and set $R_0 = \{p\}$. Once each R_i is defined, branch over all $p' \in R_i$ and for each p' pick some witness $w(p') \in W(p')$ and set $R_{i+1} = R_i \cup \bigcup_{p' \in R_i} w(p')$. This algorithm is parallel and may run forever, but it is clearly an algorithm.
- To compute a closed set we proceed much as for the previous case, but (following Proposition 42 and Definition 38) we extend R_i to R_{i+1} by adding those p such that *every* witness to p intersects with R_i.

We make no claims to efficiency (we have not even set up machinery in this paper to measure what that would mean) but what matters is that for witness semitopologies such procedures exist, in contrast e.g. to the uncomputable semitopology from Remark 46.

6 Conclusions

6.1 Related work

Topology

Topologies are everywhere and we have found another one — almost, since semitopologies are not topologies, and the anti-separation properties which we study here seem different in flavour from the separation properties usually imposed in a topological context. Still, it is pleasing to see (yet another) application gain clarity and rigour thanks to topology-flavoured ideas, and to have this new mathematical structure to investigate.

Event structures

We discussed in Remark 32 how our notion of witness function is an infinitary generalisation of a special case of the enabling relation of event structures. This does not make semitopologies a special case of event structures, since the definitions are subtly different and we apply them in very different ways — but it does hold out a prospect in future work of transferring ideas from event structures to semitopologies, and to the instantiation of semitopologies to the Stellar network in particular. Perhaps also ideas may flow in the other direction as a new application of topology-flavoured ideas to event structures.

The Consensus Task

In the traditional Consensus Problem, every process proposes a value and every process must decide a value subject to two conditions:

- *(Agreement)* all processes that decide must decide the same value, and
- *(Non-Triviality)* every decided value must have been proposed by some process.

The Consensus Problem can be identified as a task [7, Section 8.3.1], and in this context we can intuitively identify *computing agreement* with computing a continuous function on a semitopology (possibly starting from some non-continuous starting state), and *non-triviality* with a structural property implicit in Remark 46, that (in the terminology of that Remark) if p outputs v, then some process in a cover of p (see Remark 46) must have received the input v (see also Lemma 26). This suggests:

▶ **Definition 49.** *Suppose* $(\mathsf{P}, \mathsf{Open})$ *is a finite semitopology and* V *is a set of values. Then the* ***semitopological consensus task*** *is the triple* $(\mathsf{I}, \mathsf{O}, \Delta)$ *where:*

- I *is the (pure) simplicial complex with facets simplices* $\{(p_0, v_0), ..., (p_n, v_n)\}$ *where* $n = |\mathsf{P}|$, $p_i \in \mathsf{P}$ *and* $v_i \in \mathsf{V}$ *for every* $0 \leq i \leq n$*, and* $p_i \neq p_j$ *for every* $i \neq j$*.*

- O *is the (pure) simplicial complex with facets simplices* $o = \{(p_0, v_0), ..., (p_n, v_n)\}$ *where* $n = |\mathsf{P}|$, $p_i \in \mathsf{P}$ *and* $v_i \in \mathsf{V}$ *for every* $0 \leq i \leq n$, $p_i \neq p_j$ *for every* $i \neq j$, *and* o, *when seen as a function from* P *to* V, *is a continuous function on the semitopology* $(\mathsf{P}, \mathsf{Open})$.
- Δ *is the function mapping* $i \in \mathsf{I}$ *to the (pure) simplicial complex* $\Delta(i) \in 2^{\mathsf{O}}$ *such that* $\Delta(\{(p_0, v_0), ..., (p_m, v_m)\})$, $0 \leq m \leq |\mathsf{P}|$, *is the simplicial complex with facets simplices* $o = \{(p_0, w_0), ..., (p_m, w_m)\} \in O$ *where, for every* $0 \leq i \leq m$, *there exists a* cover *(minimal open set)* $O \in \mathsf{Open}$ *for* p_i *and* $0 \leq j \leq m$ *such that* $p_j \in O$ *and* $w_i = v_j$.

This definition can be extended to the case in which P is infinite when $(\mathsf{P}, \mathsf{Open})$ is a witness semitopology from Definition 30; Corollary 44 ensures that covers exist.

Note that in contrast to the classic consensus task, the semitopological consensus task is not *colourless* [7, Section 4.1.4] in general: e.g. if we have two disjoint topens, it matters which process is assigned which output value, because the two topens must agree within themselves but may disagree between one another.

Algebraic topology as applied to distributed computing tasks

Continuing the discussion of tasks above, the reader may know that solvability results about distributed computing tasks have been obtained from algebraic topology, starting with the impossibility of k-set consensus and the Asynchronous Computability Theorem [8, 1, 17] in 1993. See [7] for numerous such results.

The basic observation is that states of a distributed algorithm form a simplicial complex, called its *protocol complex*, and topological properties of this complex, like connectivity, are constrained by the underlying communication and fault model. These topological properties in turn can determine what tasks are solvable. For example: every algorithm in the wait-free model with atomic read-write registers has a connected protocol complex, and because the consensus task's output complex is disconnected, consensus in this model is not solvable [7, Chapter 4].

This paper is also topological, but in a different way: we use (semi)topologies to study consensus in and of itself, rather than the solvability of consensus or other tasks in particular computation models. Put another way: the papers cited above use topology to study the solvability of distributed tasks, but this paper shows how the very idea of 'distribution' can be viewed as a (semi)topological structure.

Of course we can now imagine that these might be combined — that there might be interesting and useful things to say about the topologies of distributed algorithms when viewed as algorithms *on* and *in* a semitopological space — and this is an explicit longer-term motivation for our research. Investigating this is future work.

Fail-prone systems and quorum systems

Given a set of processes P in a distributed system, a *fail-prone* system [15] (or *adversary structure* [9]) is a set of *fail-prone sets* $\mathcal{F} = \{F_1, ..., F_n\}$ where, for every $1 \leq i \leq n$, $F_i \subseteq \mathsf{P}$. $\mathcal{F}$ denotes the assumptions that the set of processes that will fail (potentially maliciously) is a subset of one of the fail-prone sets. A *quorum system* for $\mathcal{F}$ is a set $\{Q_1, ..., Q_m\}$ of quorums where, for every $1 \leq i \leq m$, $Q_i \subseteq \mathsf{P}$, and such that

- for every two quorums Q and Q' and for every fail-prone set F, $(Q \cap Q') \setminus F \neq \emptyset$ and
- for every fail-prone set F, there exists a quorum disjoint from F.

Several well-known distributed algorithms such as Bracha Broadcast [2] and PBFT [5] rely on a quorum system for a fail-prone system $\mathcal{F}$ in order to solve problems such as reliable broadcast and consensus assuming (at least) that the assumptions denoted by $\mathcal{F}$ are satisfied.

More recent models generalise fail-prone systems to heterogeneous settings in which processes make different failure assumptions and have different quorums. Those models include Asymmetric Fail-Prone Systems [3], Learner Graphs [19], Federated Byzantine Agreement Systems [16], Federated Byzantine Quorum Systems [6], and Personal Byzantine Quorum Systems [13]. The last three build on Stellar's Federated Byzantine Agreement Systems, where quorums are obtained using quorum slices (in Stellar's terminology), which are a special case of the notion of witness in Definition 29(2). Cobalt, SCP, Heterogeneous Paxos, and the Ripple Consensus Algorithm [14, 16, 19, 18] are consensus algorithms that rely on heterogeneous quorums or variants thereof. The Stellar network [12] and Ripple [18] are two global payment networks that use heterogeneous quorums to achieve consensus among an open set of participants.

The literature on fail-prone systems and quorum systems is most interested in synchronisation algorithms for distributed systems and has been less concerned with their deeper mathematical structure. Some work by the second author and others [13] gets as far as proving an analogue to Lemma 9 (though we think it is fair to say that the presentation in this paper much simpler and more clear), but it fails to notice the connection with topology and the subsequent results which we present in this paper. So we can view this paper as beginning an in-depth mathematical study of heterogeneous quorum systems.

6.2 Comments and future work

Heterogeneous quorum systems are an empirical fact of many distributed systems (see the references in the two paragraphs above), and we believe we can make a strong claim in this paper to have proposed an illuminating and mathematically tractable analysis of what they are: semitopologies are novel but sit in a well-understood mathematical landscape, the proofs come out well, and witness functions go some way to explaining at a high level why heterogeneous quorum systems are empirically practical in the real world.

The next step, which is current work and will be presented in a longer paper, is to study the mathematical and computational content of *arriving at consensus*, starting from a non-consensus state. In the language of this paper: given a possibly non-continuous function out of a semitopology, what does it mean to find a 'nearby' function that is continuous (i.e. represents a consensus state) that is in some sense close to and related to the starting state; and how, and in what conditions, can a nearby continuous function be computed? We hint at this in Remarks 12 and 46 where we note that such an analysis might do well to start locally by studying sets of open covers of points; this is future work.

We also hope this paper may mark a beginning for new discussions, especially based on connections with topologies and perhaps event structures — including importing algorithms to improve implementations of heterogeneous quorum systems, exporting new and interesting applications, and gaining broader and deeper understandings of the mathematical structures and connections that seem to be involved here.

References

1. Elizabeth Borowsky and Eli Gafni. Generalized FLP Impossibility Result for T-resilient Asynchronous Computations. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Theory of Computing*, STOC '93, pages 91–100, New York, NY, USA, 1993. ACM. URL: `http://doi.acm.org/10.1145/167088.167119`, `doi:10.1145/167088.167119`.
2. Gabriel Bracha. Asynchronous Byzantine agreement protocols. *Information and Computation*, 75(2):130–143, November 1987. tex.ids= brachaAsynchronousByzantineAgreement1987.

URL: http://www.sciencedirect.com/science/article/pii/089054018790054X, doi:10.1016/0890-5401(87)90054-X.

3 Christian Cachin and Björn Tackmann. Asymmetric Distributed Trust. *arXiv:1906.09314 [cs]*, June 2019. arXiv: 1906.09314.

4 Paul J Campbell. The origin of "zorn's lemma". *Historia Mathematica*, 5(1):77–89, 1978. URL: https://www.sciencedirect.com/science/article/pii/0315086078901362, doi:https://doi.org/10.1016/0315-0860(78)90136-2.

5 Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems (TOCS)*, 20(4):398–461, 2002. URL: http://dl.acm.org/citation.cfm?id=571640.

6 Álvaro García-Pérez and Alexey Gotsman. Federated Byzantine Quorum Systems. In *22nd International Conference on Principles of Distributed Systems (OPODIS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

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

8 Maurice Herlihy and Nir Shavit. The asynchronous computability theorem for t-resilient tasks. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 111–120, 1993.

9 Martin Hirt and Ueli Maurer. Player Simulation and General Adversary Structures in Perfect Multiparty Computation. *Journal of Cryptology*, 13(1):31–60, January 2000.

10 Thomas Jech. *The Axiom of Choice*. North-Holland, 1973. ISBN 0444104844.

11 Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, jul 1982. doi:10.1145/357172.357176.

12 Marta Lokhava, Giuliano Losa, David Mazières, Graydon Hoare, Nicolas Barry, Eli Gafni, Jonathan Jove, Rafał Malinowsky, and Jed McCaleb. Fast and secure global payments with Stellar. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, SOSP '19, page 80–96, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3341301.3359636.

13 Giuliano Losa, Eli Gafni, and David Mazières. Stellar consensus by instantiation. In Jukka Suomela, editor, *33rd International Symposium on Distributed Computing (DISC 2019)*, volume 146 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:15, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: http://drops.dagstuhl.de/opus/volltexte/2019/11334, doi:10.4230/LIPIcs.DISC.2019.27.

14 Ethan MacBrough. Cobalt: BFT Governance in Open Networks. February 2018. URL: https://arxiv.org/abs/1802.07240v1, doi:10.48550/arXiv.1802.07240.

15 Dahlia Malkhi and Michael Reiter. Byzantine quorum systems. *Distributed computing*, 11(4):203–213, 1998.

16 David Mazieres. The Stellar consensus protocol: a federated model for Internet-level consensus. Technical report, Stellar Development Foundation, 2015. https://www.stellar.org/papers/stellar-consensus-protocol.pdf.

17 Michael Saks and Fotios Zaharoglou. Wait-free k-set agreement is impossible: The topology of public knowledge. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 101–110, 1993.

18 David Schwartz, Noah Youngs, and Arthur Britto. The ripple protocol consensus algorithm. *Ripple Labs Inc White Paper*, 5(8):151, 2014.

19 Isaac Sheff, Xinwen Wang, Robbert van Renesse, and Andrew C. Myers. Heterogeneous Paxos. In Quentin Bramas, Rotem Oshman, and Paolo Romano, editors, *24th International Conference on Principles of Distributed Systems (OPODIS 2020)*, volume 184 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:17, Dagstuhl, Germany, 2021. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. ISSN: 1868-8969. URL: https://drops.dagstuhl.de/opus/volltexte/2021/13490, doi:10.4230/LIPIcs.OPODIS.2020.5.

20 Glynn Winskel. Event structures. In W. Brauer, W. Reisig, and G. Rozenberg, editors, *Petri Nets: Applications and Relationships to Other Models of Concurrency*, pages 325–392, Berlin, Heidelberg, 1987. Springer Berlin Heidelberg. URL: `https://web.archive.org/web/20220120181851/https://www.cl.cam.ac.uk/~gw104/Winskel1987_Chapter_EventStructures.pdf`.