The (semi)topology of distributed consensus

Murdoch J. Gabbay (joint work with Giuliano Losa)

Presented at GETCO 2022 in Paris. Thanks to the organisers.

30 May 2022

Definition of a semitopology

Definition. A semitopology is a pair $(\mathsf{P}, \mathsf{Open} \subseteq \mathit{pow}(\mathsf{P}))$ of

- a nonempty set of points and
- a set of open sets that contains P and is closed under (possibly infinite, possibly empty) unions.

Think: *"topology, minus the condition that finite intersections of opens are open"*.

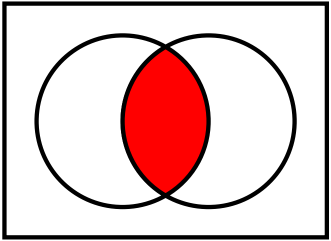

(Image credit: Wikipedia.)

Many definitions transfer smoothly . . .

Definition. Notions of

- closure of $C \subseteq \mathsf{P}$ — when $C = \mathsf{P} \setminus O$ for $O \in \mathsf{Open}$ — and
- continuity of $f : \mathsf{P} \to \mathsf{P}'$ — inverse image of open set is open, or in symbols: $f^{\text{-}1}(O') \in \mathsf{Open}$

— can be defined as usual on semitopologies.

. . . and there are differences too

In a topology, a *minimal* open neighbourhood of p is also *least*: any two minimal open neighbourhoods of p can be intersected to get an open neighbourhood included in both, hence by minimality they are equal.

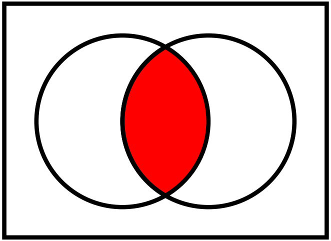

In a semitopology, p may have more than one minimal open neighbourhood, since their intersection need not be a (smaller) open neighbourhood.

Relevance to distributed consensus

Interpret $p \in \mathsf{P}$ as *participants* in a distributed algorithm. Interpret an open neighbourhood $p \in O \in \mathsf{Open}$ as a *quorum* sufficient for progress to be made.

Consensus schemes like 'decision by majority vote' or 'decision by 2/3 majority' give rise to semitopologies as follows:

- A finite set P has the majority semitopology, where $O \subseteq \mathsf{P}$ is open when O contains at least half of the elements of P.
- A finite set P has the supermajority semitopology, in which $O \subseteq \mathsf{P}$ is open when O contains 2/3 of the elements of P.

These semitopologies are not topologies: $O \cap O'$ need not be open.

Semitopologies are *local*

Many consensus algorithms are *homogeneous*, meaning they determine a quorum to agree on a value v by a generalised quantifier referring to a global property:

- Voting power is given by whoever has the most X, where
- X = ‘tokens’, or ‘computational power’, ...

Such globality is expensive, for two reasons:

- Distributed algorithms are naturally local, by definition: a global X can be inferred but not immediately locally computed.
- Deciding a global condition itself requires consensus, but if consensus is the thing we are trying to determine then this adds layers of complexity and iteration and invites attack: we have to agree on who has most X, before we can agree to agree on v!

Technical slides alert

Semitopologies are a natural framework for *local* or *heterogeneous* distributed systems. For example:

- P is a graph G and opens are generated by a local supermajority, this being any p along with 2/3 of its immediate neighbours.
- $\mathsf{P} = \mathbb{N} \times \mathbb{N}$, and opens are unions of rows and columns.

I will now pull some maths out of this deceptively simple setup, so get ready for some technical slides!

Intertwined points: the anti-Hausdorff property

Notation. Write $O \between O'$ when $O \cap O' \neq \varnothing$.

Note: we can write the Hausdorff property as saying for every $p \neq p' \in \mathsf{P}$ that

$$\exists O, O' \in \mathsf{Open}.(p \in O \wedge p' \in O') \wedge \neg(O \between O').$$

Definition. Call p and p' intertwined and write $p \between p'$ when

$$\forall O, O' \in \mathsf{Open}.(p \in O \wedge p' \in O') \Rightarrow O \between O'.$$

In words: $p \between p'$ when all their open neighbourhoods intersect. Think of $p \between p'$ as the essence of *anti-Hausdorff*.

Note: we can impose a metric on $(\mathsf{P}, \mathsf{Open})$ *where* $|p, p'|$ *is the least* i *such that* $p \between p_1 \between p_2 \cdots \between p_i \between p'$, *or* ω *if no such chain exists. A space is Hausdorff when the metric is trivial: points are either* 0 *or* ω *apart.)*

(Continuous) value assignments

Fix a set of *values* Val and give it the discrete semitopology $(\mathsf{Val}, pow(\mathsf{Val}))$ in which $\{v\}$ is open for every $v \in \mathsf{Val}$.

Call $f : \mathsf{P} \to \mathsf{Val}$ a value assignment: fix some continuous value assignment f and some participant p. Then:

- $f^{-1}(f(p)) \in \mathsf{Open}$.
- A quorum of points agree with p on its value $f(p) \in \mathsf{Val}$.

Lemma. If $p \between p'$ then $f(p) = f(p')$.

Proof. $p \in f^{-1}(f(p))$ and $p' \in f^{-1}(f(p'))$ are open (by continuity), so intersect.

Corollary 1. Continuous value assignments *agree* between intertwined points.

Transitive sets

Definition. Call $S \subseteq \mathsf{P}$ transitive when

$$\forall O, O' \in \mathsf{Open}. O \between S \between O' \Rightarrow O \between O'.$$

Call S a transitive and open set a topen set.

Lemma 1. S is topen $\Leftrightarrow$ all its points are pairwise intertwined $(\forall p, p' \in S. p \between p')$.

Proof. Suppose S is topen and $p, p' \in S$ and $p \in O$ and $p' \in O'$. Then $O \between S \between O'$ and so $O \between O'$. Conversely, if all points are pairwise intertwined and $O \between S \between O'$ then $p \in O$ and $p' \in O'$ for $p, p' \in S$ and so $O \between O'$.

Lemma 2. $\mathcal{S}$ is a set of pairwise intersecting topens $\Rightarrow \bigcup \mathcal{S}$ is topen.

Proof. $O \between \bigcup \mathcal{S} \between O'$ implies (wlog) $O \between S \between S' \between O'$ for some $S, S' \in \mathcal{S}$, and by transitivity $O \between O'$.

The fundamental theorem

Theorem 1.

1. Any semitopology $(\mathsf{P}, \mathsf{Open})$ partitions into disjoint maximal topen sets (plus isolated points), and
2. continuous value assignments are constant on each partition.

Proof.

By Lemma 2, if topen S and S' intersect then $S \between S'$ is topen. Also using Lemma 2, an increasing chain $S_0 \subseteq S_1 \subseteq \ldots$ of topens is topen. The partitioning follows.

Continuous value assignments are constant from Lemma 1 (all points intertwined) and Corollary 1 (value assignment is constant on intertwined points).

Theorem 1 = consensus

Theorem 1 provides a high-level account of consensus, in 3 easy steps:

1. Let participants make local choices about who they trust.
2. Derive a semitopological notion of quorum from these local choices by taking suitable unions (for details see *witness functions*, below).
3. Compute locally continuous value assignments.

Theorem 1 guarantees that the system will self-organise into a partition of maximal topens on which local consensus is guaranteed.

Note how this is reminiscent of real life organises itself into communities with (literally) shared values.

Why infinities matter, even in a finite world

Computer systems are finite, but this is (perhaps surprisingly) no reason to assume $(\mathsf{P}, \mathsf{Open})$ is finite:

- A participant $p \in \mathsf{P}$ sees only a local portion of the universe: as far as p is concerned, P may be unbounded.
  Thus, a distributed algorithm running at p must behave *as if* P might be infinite. (Similarly we have only a finite view of $\pi \in \mathbb{R}$, but we behave *as if* a perfect circle exists.)
- Hostile participants may be free to join the network and tell arbitrary lies. Thus, the system may appear arbitrarily large, due to input from hostile players, even if physically it is smaller.
  Thus, algorithms must be resilient and terminating even on potentially hostile, potentially nonterminating inputs.

A theory of infinite semitopologies is therefore relevant even if we can only realise finite approximations in the real world.

Semitopologies in practice

We can sum up the essential challenge of semitopologies for consensus as follows:

> *Use (semi)topology and combinatorics to design, and prove well-behavedness properties of, algorithms to compute continuous local value assignments on general classes of semitopologies.*

The Stellar payments network does this, in practice, to facilitate cross-border payments.

We see broader applications to distributed systems in which notions of quorum and identity are handled locally, rather than by (some emulation of) a central authority.

The mathematics of semitopologies has rich structure . . .

A flavour of the rest of the maths

Definition.

- A witness function is an assignment $w : \mathsf{P} \to \mathit{fin}(\mathit{pow}(\mathsf{P}))$ (finite set of subsets of P).
- Given a w, the witness semitopology sets O to be open when $p \in O$ implies $w \subseteq O$ for at least one witness $w \in w(p)$.

This is how Stellar actually works: participants don't choose *quorums*; they choose *witnesses*, and quorums are dynamically computed as needed.

A flavour of the rest of the maths: *semicompactness*

Proposition. Suppose $O_0 \supseteq O_1 \supseteq O_2 \dots$ is a descending chain of open sets in a witness semitopology.

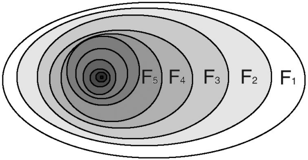

(Image credit: Katzourakis. $F \to O$)

Then $\bigcap O_i$ is open. *(Infinite intersections of opens aren't generally open, but they are here.)* As a corollary, every $p \in O$ has a minimal open neighbourhood $p \in M \subseteq O$.

This expresses a termination property: a *minimal* open neighbourhood for p always exists, for a local value assignment.
Open Problem: Does semicompactness characterise witness semitopologies?

Conclusions

Semitopology generalises topology, and is useful in the study of distributed algorithms, including for consensus. In this world, continuity at p is interpreted as consensus at p.

We are particularly interested in computing continuous functions on anti-Hausdorff spaces having lots of intertwined points. Far from being boring, such intertwined spaces have rich mathematical, combinatorial, and algorithmic structure.

Semitopologies may become of increasing relevance, especially as distributed systems become more common, especially due to recent interest in *heterogeneous* distributed systems, which are truly distributed in the sense that they do not necessarily seek to emulate the effect of a centralised authority.

More examples of semitopologies . . .

. . . just to give a sense of the design space:

- Take $\mathsf{P} = \mathbb{N}$ and let the many semitopology be generated by countable subsets. (This models the 'many' generalised quantifier.)
  This is not a topology: just because there are many points in O and in O' does not mean there are many in $O \cap O'$.
- Take $\mathsf{P} = \mathbb{N}^*$ (finite sequences of natural numbers), and let $O \subseteq \mathsf{P}$ be open when if $p \in O$ then there exists $n \in \mathbb{N}$ such that $p, n \in O$.
  This models a space of nonterminating processes.