Semitopologies and semiframes for distributed heterogeneous consensus

Murdoch J. Gabbay (joint work with Giuliano Losa)

The first workshop on
Heterogeneous Trust in Distributed Systems (HTDS 2023)

26 October 2023

Thanks

Thank you to the organisers for organising this event. I am delighted we can be here together.

Semitopologies

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

- P a nonempty set of points and
- Open a set of open sets such that:
  1. $\mathsf{P} \in \mathsf{Open}$
  2. $\mathcal{O} \subseteq \mathsf{Open} \Rightarrow \bigcup \mathcal{O} \in \mathsf{Open}$

Semitopologies abstract notions of heterogeneous trust, such as Federated Byzantine Quorum systems and Fail-Prone Systems:

- Open set $\leftrightarrow$ Quorum.
- Open neighbourhoods of point $p \leftrightarrow$ Quorums of participant p.
- Closed set $=$ complement of an open set $\leftrightarrow$ fail-prone set.

(Quorum = set of participants with potential to collaborate to locally progress, where algorithms succeed. See Slide 5.)

Semitopologies and topologies compared

Semitopology generalises topology by removing the condition that intersections of opens must necessarily open.

Example difference:

- In topology, minimal open neighbourhoods are least.
- In semitopology, a point may have multiple minimal open neighbourhoods, as illustrated for 0 (left) and 1 (right):

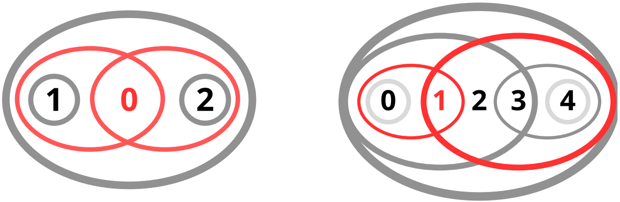

Semitopologies applied to model consensus

Imagine a set of participants P following a protocol to progress as some transition $t_1 \to t_2$, and consider $p, p' \in \mathsf{P}$.

Desirable properties include:

1. Quorum Intersection:
   Any two quorums $p \in Q_p \subseteq \mathsf{P}$ and $p' \in Q'_{p'}$ intersect.
2. Quorum Provision:
   p may only progress as $t_1 \to t_2$ provided it does so in agreement with some quorum $p \in Q_p$ of other participants that also progress as $t_1 \to t_2$.

What does this correspond to semitopologically?

Quorum Intersection = ¬ Hausdorff Separation

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

Recall the Hausdorff property that $p \neq p' \in \mathsf{P}$ have disjoint open neighbourhoods:

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

Consider the negation of the Hausdorff condition — *all* open neighbourhoods intersect:

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

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

Call $P \subseteq \mathsf{P}$ is transitive when $\forall p, p' \in P . p \between p'$.

Lemma. $P \subseteq \mathsf{P}$ is transitive iff $\forall O, O' \in \mathsf{Open}. O \between P \between O' \Rightarrow O \between O'$.

Quorum intersection for P means precisely that P is transitive.

(Non-)Intertwined points

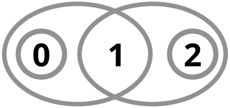

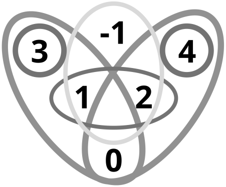

On the left, points are intertwined only with themselves.

On the right, $\{\text{-}1, 0, 1, 2\}$ are all intertwined with each other.

Definition: A topen set is an transitive open. Above right, $\{\text{-}1, 1, 2\}$, $\{1, 2\}$, $\{3\}$, and $\{4\}$ are topen.

Lemma. If $T, T' \subseteq \mathsf{P}$ are topen and $T \between T'$ then $T \cup T'$ is topen.

Corollary. Any semitopology partitions into disjoint maximal topens, plus isolated points (cf. corollary on Slide 9).

Quorum Provision = Continuity

Consider a semitopology $(\mathsf{P}, \mathsf{Open})$. Represent 'next transition step of' as a function $f : (\mathsf{P}, \mathsf{Open}) \to (\mathsf{T}, pow(\mathsf{T}))$; values for next states are treated as a discrete semitopology.

Quorum Provision corresponds to continuity at p:

"there exists $O \in \mathsf{Open}$ with $O \subseteq f^{-1}(f(p))$"

is *exactly the same condition* as

"there exists a quorum that progresses with p".

Lemma: A function to a discrete semitopology agrees on intertwined points, where continuous.

Proof: By continuity at p, f is constant on some open neighbourhood $p \in Q_p$. Similarly f is constant on some open neighbourhood $p' \in Q'_{p'}$. But $Q_p \between Q'_{p'}$, so $f(p) = f(p')$.

A topen set is a quorum of intertwined points, thus a set that can a) progress and b) will agree where algorithms succeed.

Semitopological classification of points

Suppose $(\mathsf{P}, \mathsf{Open})$ is a semitopology and $p \in \mathsf{P}$.

Definition:

1. Write $p_{\between} = \{p' \mid p' \between p\}$ and call this intertwined of p.
2. Write $K(p) = \mathit{interior}(p_{\between})$ and call this the community of p.
3. Call p regular when $p \in K(p) \in \mathsf{Topen}$.
4. Call p weakly regular when $p \in K(p) \in \mathsf{Open}$.
5. Call p quasiregular when $K(p) \neq \varnothing$.

A regular point p is a *good* point. It has a quorum of points $K(p)$ that (because it is open) can progress and (because it is transitive) remains in agreement where algorithms succeed.

Theorem. p is regular $\Leftrightarrow$ p has a topen neighbourhood.

Corollary. Any semitopology partitions into disjoint topen communities of regular points, plus non-regular points.

Semitopological classification of points

1. A regular point p is in a quorum of points $K(p)$ that (because it is open) can progress and (because it is transitive) remains in agreement where algorithms succeed.
2. A weakly regular non-regular point p is such that $K(p)$ includes multiple topens. If those communities agree with one another, then p can progress as part of $K(p)$.
3. A quasiregular non-weakly-regular point p is intertwined with some topen. It cannot necessarily progress as part of a quorum, but it can follow points that do.

Pictures follow.

Examples of (weak /quasi-)regularity

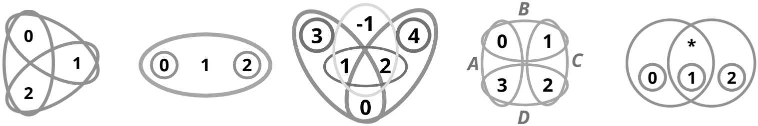

Left-to-right, we have:

1. A topen set $\{0, 1, 2\}$. All points are regular.
2. $\{0\}$ and $\{2\}$ are topen. 1 is intertwined with 0 and 2, and is weakly regular.
3. $\{3\}$, $\{4\}$, $\{1, 2\}$, and $\{\text{-}1, 1, 2\}$ are topen.
   Point 0 is weakly regular; its community includes topens $\{3\}$, $\{4\}$, and $\{1, 2\}$.
4. Points are intertwined only with themselves and no points are even quasiregular.
5. $\{0\}$, $\{1\}$, and $\{2\}$ are topen. $*$ is quasiregular with $K(*) = \{1\}$.

Semiframes

A semiframe is an algebraic version of a semitopology; a *compatible complete semilattice*:

Definition. A semiframe is a tuple $(\mathsf{X}, \leq, *)$ such that

1. $(\mathsf{X}, \leq)$ is a complete semilattice.
2. $*$ is a compatibility relation on $(\mathsf{X}, \leq)$:
   2.1 $*$ is commutative.
   2.2 $*$ distributes over $\bigvee$, thus $x * \bigvee X' = \bigvee\{x * x' \mid x' \in X'\}$.
   2.3 $*$ is properly reflexive: $x * x$ if $x \neq \bot$.

Theorem. There is a categorical duality between the category SemiTop of semitopologies and continuous functions between them, and the category SemiFrame of semiframes and semiframe morphisms between them; see paper on semiframes for details and full definitions. Idea: $\leq \leftrightarrow \subseteq$ and $* \leftrightarrow \between$.

This exhibits semitopologies as a generalisation of topologies; semitopologies are topologies with a *generalised intersection*.

The logic

Consider this semitopology:

$\mathbf{3}' \quad =$ 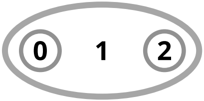

Consider a partial function f from $\mathbf{3}'$ to $\mathrm{Bool} = \{\bot, \top\}$ with the discrete semitopology, such that $f(0) = \bot$ and $f(2) = \top$.

This cannot be completed to a total continuous function; setting $f(1) = \bot$ or $f(1) = \top$ yields discontinuity at 1.

In fact the only continuous functions from $\mathbf{3}'$ to Bool are constant! In the continuous functions $\mathrm{SemiTop}(\mathbf{3}', \mathrm{Bool})$, $\mathbf{3}'$ behaves like a singleton semitopology.

As a corollary, given a semitopology S, the set of continuous functions $\mathrm{SemiTop}(\mathrm{S}, \mathrm{Bool})$ does not necessarily determine S: counterexample $\mathrm{S} = \mathbf{3}'$.

The logic

This means that:

1. We cannot reason about semitopologies using continuous functions to Bool.
2. In categorical language: Bool does not classify SemiTop.
3. In logical language: the natural propositional language of semitopologies is not two-valued.

So?

In fact, $\mathbf{3}'$ classifies SemiTop: the set of continous functions $\mathsf{SemiTop}(\mathsf{S}, \mathbf{3}')$ uniquely determines S.

The logic

$\mathsf{SemiTop}(\mathsf{S}, \mathbf{3}')$ is a complete lattice (because $\mathbf{3}'$ is), with $\wedge$, $\vee$, and also $\neg$ (swaps 0 and 2). This suggests a *3-valued paraconsistent logic*. Note the claim implicit in this observation:

Claim. Any approach to semitopologies, and by extension to heterogeneous consensus, must necessarily contain a three-valued logic. The third truth-value may be disguised, but it *must be there* because our example of $\mathbf{3}'$ demonstrates that without it, we cannot distinguish $\mathbf{3}'$ from singleton semitopology by continuous functions.

It turns out that we can build such a logic, and it is fruitful:

1. We can characterise desirable properties like quorum intersection and regularity.
2. Stellar's notion of *slices* corresponds to a Horn clause theory.

Paraconsistent logic

Definition. Define $\mathbf{3} = \{\mathbf{t}, \mathbf{b}, \mathbf{f}\}$, with semitopology

$$\begin{aligned} \mathsf{Open} &= \{\varnothing, \{\mathbf{t}\}, \{\mathbf{f}\}, \{\mathbf{t}, \mathbf{f}\}, \mathbf{3}\} \\ \mathsf{Closed} &= \{\mathbf{3}, \{\mathbf{t}, \mathbf{b}\}, \{\mathbf{f}, \mathbf{b}\}, \{\mathbf{b}\}, \varnothing\}. \end{aligned}$$

Make $\mathbf{3}$ into a lattice, ordered as $\mathbf{t} \geq \mathbf{b} \geq \mathbf{f}$, with operators

- $\wedge$ (glb),
- $\vee$ (lub),
- and $\neg$ (reverses the lattice, thus swaps $\mathbf{t}$ and $\mathbf{f}$ and fixes $\mathbf{b}$).

Returning to our example, given $f(0) = \mathbf{t}$ and $f(2) = \mathbf{f}$, we can safely set $f(1) = \mathbf{b}$ and the valuation becomes continuous.

(Yes, $\mathbf{3}'$ and $\mathbf{3}$ are isomorphic, but the classification $\mathsf{SemiTop}(\mathbf{3}', \mathbf{3})$ is canonical and generalises to *any* semitopology S on the left-hand side.)

Some truth tables

$\neg p$	
t	**f**
b	**b**
f	**t**

$p \wedge q$	**t**	**b**	**f**
t	**t**	**b**	**f**
b	**b**	**b**	**f**
f	**f**	**f**	**f**

$p \vee q$	**t**	**b**	**f**
t	**t**	**t**	**t**
b	**t**	**b**	**b**
f	**t**	**b**	**f**

$(\neg p) \vee q = p \rightsquigarrow q$	**t**	**b**	**f**
t	**t**	**b**	**f**
b	**t**	**b**	$\underline{\mathbf{b}}$
f	**t**	**t**	**t**

$p \leftrightsquigarrow q = p \rightsquigarrow q \wedge q \rightsquigarrow p$	**t**	**b**	**f**
t	**t**	**b**	**f**
b	**b**	**b**	**b**
f	**f**	**b**	**t**

$p \Rightarrow q$	**t**	**b**	**f**
t	**t**	**b**	**f**
b	**t**	**b**	$\underline{\mathbf{f}}$
f	**t**	**t**	**t**

$p \Leftrightarrow q$	**t**	**b**	**f**
t	**t**	**b**	**f**
b	**b**	**b**	**f**
f	**f**	**f**	**t**

Current & future work

1. Gabbay and Losa have a theory of point-set semitopologies: https://arxiv.org/abs/2303.09287.
2. Semitopologies have a corresponding point-free theory of semiframes, semifilters, and a categorical duality between them: https://arxiv.org/abs/2310.00956.
3. A three-valued paraconsistent logic is naturally suggested by the mathematics, and turns out to correspond to some engineering realities of Sellar — e.g. Sellar's consensus algorithm solves a Horn clause theory and progresses when it detects a quroum (= open set). Paper in preparation.

Semitopologies are a new connection between heterogeneous consensus and pure mathematics, via topology. *And* they are a source of new ideas, *and* they connect nicely to logic.

If semitopologies might help you too then be in touch!