

# Decentralised collaborative action: cryptoeconomics in space

Murdoch J. Gabbay

## Abstract

Blockchains and peer-to-peer systems are part of a trend towards computer systems based on *decentralised collaborative action*, by which we mean that they 1) run across many participants, 2) without central control, and 3) are such that qualities 1 and 2 are essential to the system's intended use cases.

We propose a notion of topological space, which we call a *semitopology*, to help us mathematically model such systems. We treat participants as *points* in a space, which are organised into *actionable coalitions*. An actionable coalition is any set of participants who collectively have the resources to collaborate (if they choose) to progress according to the system's rules, independently of the rest of the system.

Mathematicians will recognise semitopology as a generalisation of the notion of point-set topology, where actionable coalitions correspond to open sets.

It turns out that much useful information about the system can be obtained *just* by viewing it as a semitopology and studying its actionable coalitions. For example: we will prove a mathematical sense in which if every actionable coalition of some point  $p$  has nonempty intersection with every actionable coalition of another point  $q$  — note that this is the negation of the famous Hausdorff separation property from topology — then  $p$  and  $q$  must remain in agreement. Remarkably, since this observation depends only on the semitopological structure, it holds for any possible concrete algorithm.

This matters because remaining in agreement is a key correctness property in many distributed systems. For example in blockchain, participants disagreeing is called *forking*, and blockchain designers try hard to avoid it.

We provide an accessible introduction to: the technical context of decentralised systems; why we build them and find them useful; how they motivate the theory of semitopological spaces; and we sketch some basic theorems and applications of the resulting mathematics.

# 1 What is ‘cryptoeconomics’?

Let us begin by proposing a definition:

DEFINITION 1.1. **Cryptoeconomics** is:

the study of socioeconomic systems enabled by modern decentralised computer systems.

Cryptoeconomics is closely related to blockchains, because a blockchain is a decentralised database, and databases store *state* (i.e. data): thus, blockchains make cryptoeconomics possible, and cryptoeconomic outcomes — real or envisaged — make blockchains useful.

There is substantial overlap between cryptoeconomics and many other fields that study aspects of decentralised and distributed systems, including:

- *economics* (the study of value and incentives),
- *game theory* (the study of optimal outcomes for players with choices),
- *social choice theory* (how to synthesise collective decisions from individual votes),
- *blockchains* (decentralised databases),
- *smart contracts* (programs that operate on blockchains), and
- *law* (the interpretation of facts into socially-binding meaning).

We see from this list that in practice, cryptoeconomics touches on nearly everything.

The description of cryptoeconomics in Definition 1.1 uses two adjectives: ‘digital’ and ‘decentralised’. Digitisation is important because it enables decentralisation on a historically unprecedented scale, but it is the decentralisation that is most important to giving cryptoeconomics its particular character.

Concerning digitisation, this has been ongoing since (at least) the digital *mainframe computers* of the 1960s — powerful central computers that acted then, and still act today, as oracles to enable organisations that can afford them to deliver services more efficiently. Note however that this is invisible to the end user, in the sense that the user just sees a bigger, better, faster, and possibly cheaper service. This is a substantive technological advance, and great for efficiency and profits, but it does not necessarily lead to any qualitative structural change in the economics of how the value is created.

We can date the seeds of *decentralised* digitisation to the mid-1970s, when pocket digital calculators took over from slide rules, and desktop computers became available. Putting digital computation in people’s pockets and on people’s desks started a cascade of innovations that, along with the internet, has brought us modern miracles like mobile banking and streaming video. Yet even so, a user doing mobile banking on a mobile phone or watching a video on a streaming service has something fundamental in common with the technician querying a 1960s mainframe computer using a paper card with holes punched in it: the back-end system is still centrally controlled. That mainframe computer system is still around, albeit in an immensely more sophisticated form.<sup>1</sup>

---

<sup>1</sup>In practice, the modern ‘mainframes’ that (for example) serve banking, social media, or streaming services, are usually distributed clusters of servers. But being distributed is not the same as being decentralised: see Definition 2.1(1).

What makes the new breeds of modern computing systems uniquely different is that they are *radically decentralised* and heterogeneous, such that *if they were centrally controlled then they would not even make sense*. This is the start of *cryptoeconomics* as intended in Definition 1.1, and it is new.

It began with Napster (a peer-to-peer filesharing system) in June 1999, which demonstrated how music media could be disseminated independently of (centralised) media companies [Dav16].<sup>2</sup> Then Bitcoin (January 2009) and Ethereum (July 2015) changed everything by showing that money and even contracts could be mediated (albeit imperfectly) independently of banks.

These systems — of which Napster, Bitcoin, and Ethereum are perhaps the best known, but certainly not the only examples — would be meaningless and make no sense if they were implemented in a centralised manner, in much the same way that a pocket calculator or desktop computer would make no sense if it were a mainframe. The decentralisation is not a quality of the thing: it is the *point* of that thing.

This tendency towards radical decentralisation is often discussed in ideological terms, but compelling technical forces also exist to push technology in this direction. Being decentralised gives desirable properties, including: scalability, redundancy, reliability, and resilience. Being decentralised offers unique opportunities for a network of participants to *act collaboratively* to create value by achieving their goals. We can sum this up as follows:

$$\text{cryptoeconomics} = \text{value} + \text{decentralised collaborative action.}$$

We are only beginning to get to grips with the implications of this equation.

## 2 What is ‘decentralised collaborative action’?

### 2.1 The definition

*Decentralised collaborative action* is a feature of *decentralised permissionless heterogeneous computing systems*. Let’s unpack the jargon:

DEFINITION 2.1.

1. A system is *decentralised* when it is *distributed* (meaning that it is composed of several distinct parts), and furthermore the system as a whole is not centrally controlled.

---

<sup>2</sup>The article itself illustrates some of the compromises involved in creating and disseminating knowledge. The article’s author is an Associate Professor at a UK University. The publisher’s version is behind a paywall (since the publisher makes money from publishing); an author’s preprint is made freely online by his University (which makes money by employing the author to educate students); and the author most likely wrote the article in-between teaching obligations, for zero marginal cost to his employer (i.e. ‘for free’).

To be fair: the publisher’s version looks nicer than the author’s preprint, and the publisher’s website makes the work easy to find; the University survives and its students get taught by a well-informed professor; the author’s preprint is accessible to any reader who can dig it out; and the author enjoyed writing the article. In this sense, all parties — the author, the publisher, the students, the university, the article’s readership, and science itself — benefit from the shambling compromise that is academic publishing, though they might also all complain about the division of that benefit.

Most blockchain systems and peer-to-peer networks are decentralised in this sense. The internet is also (mostly) decentralised, at least in principle.<sup>3</sup>

2. A system is *permissionless* (or *unpermissioned*) when participants can leave and join the system at any time.

Nature is naturally permissionless (living things do not need permission to be born or die). National voting systems *are* permissioned (because citizens require government certification to be allowed to vote). Peer-to-peer systems (including filesharing and blockchain protocols) are often, though not always, unpermissioned.

3. A system is *heterogeneous* when participants may legitimately be following different rules.<sup>4</sup>

Ethereum and Tezos are decentralised and permissionless, but they are not heterogeneous in the sense we intend. If you are running a Tezos or Ethereum node, then you are not forced to follow the rules, but if you do not then by definition you are not acting as a legitimate Tezos or Ethereum node.

In contrast, consider the combination of Tezos and Ethereum as a single system connected by a *blockchain bridge*.<sup>5</sup> This is heterogeneous, because Tezos nodes and Ethereum nodes have different rules and different consensus mechanisms. A Tezos node is not a bad node just because it is not following the rules of Ethereum, and vice-versa, but because of the blockchain bridge, they can be considered to be operating within a single (heterogeneous) combined system.

There are many flavours of decentralised system, but in the most general case we have a decentralised heterogeneous permissionless system that consists of *some* participants communicating to do *something*, with no *a priori* restrictions on who, what, or how.

This scenario — with its weak well-behavedness assumptions that do not even assume all participants share a common ruleset — might seem a terrible idea which we should not allow, because it admits crazy networks with bad behaviour. But here the generality is a feature, not a bug:

1. Mathematically speaking, it can be *useful* to admit general models, including both good and bad ones, so that we can formalise their good and bad behaviour<sup>6</sup> and express conditions to include or exclude it.

---

<sup>3</sup>The internet was designed to be an information network that would be resilient to nuclear attack. It did this by being ‘centrifugal’; emphasising node-to-node actions instead of centre-to-centre actions. See [Rya10], summarised by Ars Technica [Rya11].

Note that the boundary between ‘distributed’ and ‘decentralised’ can be fluid. For example, should we consider a system to be decentralised when its parts can act independently most of the time, but every so often they check in with a central controller? This depends on what aspects of the system we care about; e.g. its short-term or long-term behaviour. There is room for a nice discussion here, but it will not be in this particular article.

<sup>4</sup>By ‘different rules’ we include the situation where an algorithm (such as a consensus algorithm) is agreed between participants but a critical parameter may vary substantively across them, e.g. imagine a blockchain in which some participants require a  $>2/3$  majority to act, and others require just a  $>1/2$  majority. By ‘legitimate’ we exclude the case of a hostile participant.

<sup>5</sup>See <https://ethereum.org/en/bridges/> (permalink: <https://web.archive.org/web/20240324090911/https://ethereum.org/en/bridges/>).

<sup>6</sup>... which will vary by application; e.g. sometimes all participants should play by the same rules, but in the case of a blockchain bridge we specifically want to *admit* different rules.

2. Surprisingly, it will turn out that there is still a lot that we can say even about the general case, and we shall see that much useful structure will emerge, even from very weak assumptions.

So granted that the generality of decentralised collaborative action is a feature, not a bug; but how should we approach this mathematical generality?

The key is to look at how groups of participants can *progress* (i.e. *update*) their *local state*. To see this we need a little more discussion.

In a decentralised system, a participant must store local state — if there were a global source of truth for state then whoever controls that truth would *de facto* control the system — and communicate with other participants to decide on how their local states evolve. There must be *some* rules about how this state should be updated, even if these rules may differ across participants in the system, and even though the rules may not always be followed. It turns out that one common feature of decentralised heterogeneous permissionless systems is a notion of what in [Gab24, Gab25] is called an *actionable coalition*, by which we mean

*a set of participants who are legally entitled (but not obliged) to collaborate to progress and update their local state (possibly but not necessarily in identical ways).*

## 2.2 Some informal examples from real life

We will consider some examples of actionable coalitions, taken from real life:

1. Ethereum.  
Ethereum’s consensus protocol is proof-of-stake, so an actionable coalition on Ethereum is any group of participants who hold a majority stake of tokens (this is a bit of a simplification, but it will do).
2. Ethereum and Tezos with a blockchain bridge between them.  
Tezos’s consensus protocol is also proof-of-stake. An actionable coalition in this system is either an actionable coalition of Ethereum, or one of Tezos, *or* the sets union of an actionable coalition from each, along with the bridging node (again, a simplification, but it will do).<sup>7</sup> We return to this in Example 3.6.
3. A Tango dance evening where leaders will only dance with followers and vice-versa.<sup>8</sup>  
An actionable coalition is any set containing equal numbers of leads and followers.
4. A set of people wishing to lift a heavy rock.  
An actionable coalition is any subset of these people who lifting together have enough strength to do so.

---

<sup>7</sup>Typically, participants can update their state if they held a majority of the stake at some time in the past (e.g. two weeks ago) — the idea being that all participants have reached agreement on, and learned, the state of the network two weeks in the past, so this can be treated as immutable common knowledge without undermining the decentralised nature of the system in the present [Goo14, Subsection 3.2.1, final paragraph].

<sup>8</sup>Many dancers can both lead and follow, including this author, but for the sake of the mathematics we will simplify.

If an actionable coalition can communicate to agree on a set of local state updates, e.g. if the Tango lead leads a move and the Tango follower chooses to follow it, then the participants in this coalition are entitled to update their local states accordingly. Note that local state updates need not be literally identical across participants; they just need to be mutually agreed upon and then actioned.

Some important notes:

1. The actionable coalition can progress *without* consulting the rest of the system.
2. Being in an actionable coalition does not imply control. This set describes a potential legal collaboration, but participants can choose what actionable coalition to work with, if any, and they can also choose not to follow the rules.<sup>9</sup>
3. If  $O$  is an actionable coalition for some participant  $p$ , and  $p'$  is another participant in  $O$ , then  $O$  is also an actionable coalition for  $p'$ . Note that this makes actionable coalitions look a bit like open sets in a topology.

So we can now introduce our first mathematical abstraction: we identify participants as *points*, and we let *open sets* be *actionable coalitions*. An actionable coalition is a *coalition of participants with the capacity to act*. They are not obliged to act, and if they do act then their action need not be identical across all participants, but the potential exists for this set to collaborate to progress their states.

1. With reference to our couples dance example: an example of an actionable coalition that is not minimal is a set containing two leads and two followers. There are two ways for the participants to pair off to collaborate (i.e. dance).
2. With reference to our bridged blockchain example: an example of a set that contains an actionable coalition but is not one itself is an actionable coalition from Ethereum, along with the bridging node. The Ethereum coalition on its own is actionable, but the bridging node cannot take any action without also collaborating with an actionable coalition from Bitcoin.

To get a flavour of our mathematical results, consider a fundamental problem in any decentralised system: *consensus*; i.e. the problem of ensuring that participants remain in agreement, for some suitable sense of ‘agree’. To take a simple example from blockchain: if we reach a situation where half of the nodes say that we have paid for a service, and the other half say that we have not — then *everyone* has a problem, because the system has become incoherent and it is not clear how the system as a whole can restore coherence and progress.<sup>10</sup> This phenomenon is called *forking*, and blockchain designers really want to avoid it!

We will call our mathematical abstraction of agreement, *antiseperation*. In a little more detail, antiseperation properties are coherence properties that are guaranteed to hold of a decentralised system *just* by analysing the structure of its actionable coalitions.

---

<sup>9</sup>If you put your elbow into your dance partner’s eye, or simply deliver a poor lead or a poor follow, then the other dancer might stop dancing with you or turn you down if you ask for another dance. But neither of you are *compelled* to dance with one another, and if you do, you are not *compelled* to dance well.

<sup>10</sup>coherent (adj.) 1550s, “harmonious;” 1570s, “sticking together,” also “connected, consistent” (of speech, thought, etc.), from French cohérent (16c.), from Latin cohaerentem (nominative cohaerens), present participle of cohaerere “cohere,” from assimilated form of com “together” (see co-) + haerere “to adhere, stick” (etymologyonline: <https://www.etymonline.com/word/coherent>).

It turns out that we can get surprisingly detailed information about agreement/antiseperation properties, even working from quite weak and abstract mathematical assumptions on the actionable coalitions.

We emphasise this point: sometimes we can predict important macro properties of a system's behaviour without knowing anything about its specifics, so long as we have certain good properties on its actionable coalitions.

### 2.3 Two formal mathematical examples

DEFINITION 2.2. Call **binary consensus** the problem of getting participants in a distributed system to announce a single value  $t$  or  $f$ . This is a simplest possible consensus problem, but note that by running multiple rounds of binary consensus we can get participants to announce *bitstrings* (finite sequences of values), and arbitrary data can be serialised to bitstrings, so this consensus problem — simple as it is — is also complete for all data in a suitable sense.

EXAMPLE 2.3 (A simple majority system). We consider a simple situation where participants are trying to solve binary consensus. Continuing the theme of simplicity, assume some finite nonempty set of participants  $\mathbf{M}$  and let their actionable coalitions be just any set of participants that forms a majority (so it contains strictly more than half of  $\mathbf{M}$  the set of all participants). Now suppose that the participants in some actionable coalition  $O \subseteq \mathbf{M}$  ( $O \subseteq \mathbf{M}$  means that  $O$  is a set of elements from  $\mathbf{M}$ ) have communicated and have progressed to agree on the value  $t$ . Because they form an actionable coalition, they are entitled to act and to announce  $t$ , and so they do. They have now all committed to this state update and they cannot change their minds.

So: can this system fork? Consider some participant  $p \notin O$  ( $p \notin O$  means that  $p$  is a participant that is not in  $O$ ). If  $p$  wants to progress to decide on some value, that value must be  $t$ . This is because all of its actionable coalitions intersect with  $O$ , and so they contain at least one participant that has already committed to  $t$  and cannot change its mind.

This does not mean that  $p$  has to agree on  $t$ ; it could choose not to agree with anything and not progress (i.e. not update its local state with any value), or it could break the rules. But, by definition if  $p$  does want to made a decision legally, then the decision has been made and it must eventually go along with the majority. Thus, we have proved that any progress that is made by one participant within the rules (... must be shared with some actionable coalition of that participant, and since all such coalitions intersect it ...) must eventually be followed any other participant that also progresses within the rules. Thus forking is impossible.

The reader may already be familiar with Example 2.3, but note that this antiseperation property comes simply *from the structure of the actionable coalitions*. There is no need to consider the protocol, or even how values are interpreted. It turns out that antiseperation-style behaviour is common, and arises even if we do not require actionable coalitions that are simple majorities. For example:

EXAMPLE 2.4. Let participants be  $\mathbb{Z} = \{0, 1, -1, 2, -2, \dots\}$  and let actionable coalitions be generated by sets of three consecutive numbers starting at an even number  $\{2i, 2i+1, 2i+2\}$  — so for example  $\{0, 1, 2\}$  and  $\{2, 3, 4\}$  are actionable coalitions,

but not  $\{2\}$  or  $\{1, 2, 3\}$  — and suppose again that we are trying to agree on  $\mathbf{t}$  or  $\mathbf{f}$ .

Note that unlike for Example 2.3, actionable coalitions need not intersect. Yet, once one triplet of participants commits to  $\mathbf{t}$ , the rest of the system is obliged to eventually agree, if all participants play by the rules. Now this example system is not necessarily particularly safe or desirable in practice, because we can imagine that  $\{0, 1, 2\}$  agree on  $\mathbf{t}$ , and  $\{4, 5, 6\}$  acting independently but in good faith agree on  $\mathbf{f}$ , and then 3 cannot legally progress, because within  $\{2, 3, 4\}$ , 2 has announced  $\mathbf{t}$  and 4 has announced  $\mathbf{f}$  and 3 cannot agree with both. But, we know that *if* all participants do legally progress, then they announce the same value. So this example illustrates how antiseperation can arise even when actionable coalitions are rather small.

The two examples above are quite different. In one, all actionable coalitions intersect, and in the other they mostly do not. This suggests that a ‘general mathematics of (anti)separation’ is possible, based on the study of actionable coalitions. In a nutshell, that mathematical story is what we will develop.

### 3 What is a semitopology?

#### 3.1 The definition

So at a high level, what do we have? Points are synonymous with participants, and:

1. There is a notion of an *actionable coalition* (or just: *open set*). This is a set  $O \subseteq \mathbf{P}$  of participants with the capability, though not the obligation, to act collaboratively to advance (= update / transition) the local state of the elements in  $O$ , possibly but not necessarily in the same way for every  $p \in O$ .
2. The empty set  $\emptyset$  (containing no points) is trivially an actionable coalition. Also we assume that  $\mathbf{P}$  (containing all the points) is actionable, effectively assuming that every point is a member of at least one actionable coalition.
3. A sets union of actionable coalitions, is an actionable coalition.

This leads us to the definition of a semitopology.

**DEFINITION 3.1.** Suppose  $\mathbf{P}$  is a set. Write  $\text{pow}(\mathbf{P})$  for the powerset of  $\mathbf{P}$  (the set of subsets of  $\mathbf{P}$ ). Then a **semitopological space**, or **semitopology** for short, consists of a pair  $(\mathbf{P}, \text{Open}(\mathbf{P}))$  of

- a (possibly empty) set  $\mathbf{P}$  of **points**, and
- a set  $\text{Open}(\mathbf{P}) \subseteq \text{pow}(\mathbf{P})$  of **open sets**,

such that:

1.  $\emptyset \in \text{Open}(\mathbf{P})$  and  $\mathbf{P} \in \text{Open}(\mathbf{P})$ .  
In words: the empty set of points, and the set of all points, are both open sets.
2. If  $X \subseteq \text{Open}(\mathbf{P})$  then  $\bigcup X \in \text{Open}(\mathbf{P})$ .<sup>11</sup>  
In words: a sets union of open sets, is an open set.

---

<sup>11</sup>There is a little overlap between this clause and the first one: if  $X = \emptyset$  then by convention  $\bigcup X = \emptyset$ . Thus,  $\emptyset \in \text{Open}(\mathbf{P})$  follows from both clause 1 and clause 2. If desired, the reader can just remove the condition  $\emptyset \in \text{Open}(\mathbf{P})$  from clause 1, and no harm would come of it.

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

We recognise a semitopology as being like a *topology* [Eng89, Wil70], but without the condition that the intersection of two open sets necessarily be an open set. This reflects the fact that the intersection of two actionable coalitions need not itself be an actionable coalition.

NOTATION 3.2. Suppose  $X$  and  $X'$  are sets. Then write  $X \not\sim X'$  when  $X$  and  $X'$  are not disjoint; i.e. they have a nonempty sets intersection:  $X \cap X' \neq \emptyset$ .

We can now state a key definition:

DEFINITION 3.3. Suppose  $(\mathbf{P}, \text{Open})$  is a semitopology and  $p, p' \in \mathbf{P}$ . Then:

1. Write  $p \not\sim p'$  and call  $p$  and  $p'$  **intertwined** when

$$\forall O, O' \in \text{Open}. p \in O \wedge p' \in O' \implies O \not\sim O'$$

In words:  $p$  and  $p'$  are intertwined when they have no pair of disjoint open neighbourhoods.

2. Write  $p \not\sim^* p'$  and call  $p$  and  $p'$  **transitively intertwined** when they are related by the transitive closure of  $\not\sim$ ; thus there exists some (possibly zero length) chain  $p_0, p_1, \dots, p_{n-1}, p_n \in \mathbf{P}$  such that

$$p = p_0 \not\sim p_1 \dots p_{n-1} \not\sim p_n = p'.$$

This recalls the *Hausdorff separation* property, typical in topology, that any two distinct points should have a disjoint pair of open neighbourhoods. As we shall see from Theorem 3.4, for the study of consensus we are particularly interested in semitopologies with antiseperation properties, of which being intertwined is a canonical such property; and it is the precise negation of being Hausdorff separated.

## 3.2 Two mathematical results

Recall the notion of *binary consensus* from Definition 2.2. Now that we have built some mathematical machinery, we can represent binary consensus as the problem of defining a function  $f : \mathbf{P} \rightarrow \{\mathbf{t}, \mathbf{f}\}$ . We will call such a function a **value assignment**.

If the value assignment is constant (so it maps all points to just one value) then this represents system-wide consensus across all of  $\mathbf{P}$ .

Call a value assignment **continuous** at  $p \in \mathbf{P}$  when there exists an open neighbourhood  $p \in O \in \text{Open}$  such that  $\forall p' \in O. f(p) = f(p')$ . The reader can check that this coincides with the usual notion of topological continuity, if we give  $\{\mathbf{t}, \mathbf{f}\}$  the discrete topology (so  $\{\mathbf{t}\}$  and  $\{\mathbf{f}\}$  are open sets); a proof is in [Gab24, Lemma 2.2.4, page 22]. It also coincides with our intuition that if  $p$  declares some value, then it must do so in collaboration with an actionable coalition of other participants. Thus we can write

$$\text{consensus} = \text{continuity},$$

and we can prove:

THEOREM 3.4. Suppose  $(\mathbf{P}, \text{Open})$  is a semitopology and  $p, p' \in \mathbf{P}$ , and suppose  $f : \mathbf{P} \rightarrow \{\mathbf{t}, \mathbf{f}\}$  is a value assignment. Then:

1. If  $p$  and  $p'$  are intertwined and  $f$  is continuous at  $p$  and  $p'$ , then  $f(p) = f(p')$ .
2. If  $p$  and  $p'$  are transitively intertwined and  $f$  is continuous at all points, then  $f(p) = f(p')$ .

*Proof.* For part 1, suppose  $f$  is continuous at  $p$  and  $p'$ , and suppose  $p \bowtie p'$ . Then by assumption there exist open neighbourhoods  $p \in O \in \text{Open}$  and  $p' \in O' \in \text{Open}$  such that  $f$  is constant on  $O$  and on  $O'$ . Since  $O$  and  $O'$  intersect,  $f(p) = f(p')$ .

For part 2, suppose  $f$  is continuous at every point, and suppose  $p \bowtie^* p'$ . By assumption there exists a chain of intertwinedness relations  $p = p_0 \bowtie p_1 \dots p_{n-1} \bowtie p_n = p'$ , and also by assumption  $f$  is continuous at each of these points. By part 1 of this result  $f(p) = f(p_0) = f(p_1) = \dots = f(p_{n-1}) = f(p_n) = f(p')$ .  $\square$

Simple as Theorem 3.4 is, it explains the consensus behaviour we observed of Examples 2.3 and 2.4, via the following easy Lemma:

LEMMA 3.5. *All points in the semitopology in Example 2.3 are intertwined. All points in the semitopology in Example 2.4 are transitively intertwined.*

*Proof.* Left to the reader to check.  $\square$

EXAMPLE 3.6. Consider our previous example of Ethereum and Tezos, connected by a bridging node. What would this look like in abstract semitopological terms?

Assume two semitopologies  $(\mathbf{E}, \text{Open}(\mathbf{E}))$  and  $(\mathbf{T}, \text{Open}(\mathbf{T}))$  such that (for simplicity)  $\mathbf{E} \cap \mathbf{T} = \emptyset$ . Assume some other point  $r \notin \mathbf{E} \cup \mathbf{T}$ , which we will call the **bridging node**. Then define a semitopology  $(\mathbf{B}, \text{Open}(\mathbf{B}))$  by:

- $\mathbf{B} = \mathbf{E} \cup \{r\} \cup \mathbf{T}$ .
- $\text{Open}(\mathbf{B})$  is the closure under arbitrary unions of

$$\text{Open}(\mathbf{E}) \cup \text{Open}(\mathbf{T}) \cup \{O \cup \{r\} \cup O' \mid O \in \text{Open}(\mathbf{E}), O' \in \text{Open}(\mathbf{T})\}.$$

Intuitively, a quorum in the combined system is either a quorum from  $\mathbf{E}$ , or one from  $\mathbf{T}$ , or it is a pair of quorums along with the bridging node  $r$ . The key point about this structure is that the bridging node can only make progress if it is in consensus with some quorum from  $\mathbf{E}$  *and* at the same time some quorum from  $\mathbf{T}$ . This is just what we would expect a bridging node to do.

Now suppose that  $(\mathbf{E}, \text{Open}(\mathbf{E}))$  and  $(\mathbf{T}, \text{Open}(\mathbf{T}))$  are intertwined — which is what we would hope, since this indicates a pair of blockchains that will not fork. Then  $(\mathbf{B}, \text{Open}(\mathbf{B}))$  is *transitively* intertwined, via the bridging node.

## 4 Conclusions and open questions

### 4.1 Overview

We have discussed *semitopology*, a generalisation of point-set topology that removes the restriction that intersections of open sets need necessarily be open. The intuition is that points represent participants in a decentralised system, and open sets represent

<sup>12</sup>Hint: check that  $2i+1$  is intertwined with  $2i$  and  $2i+2$ .

collections of participants that collectively have the authority to collaborate to update their local state; we call this an *actionable coalition*.

Examples of actionable coalition include: majority stakes in proof-of-stake blockchains; communicating peers in peer-to-peer networks; and even pedestrians working together to not bump into one another in the street. Where actionable coalitions exist, they have in common that: collaborations are local (updating the states of the participants in the coalition, but not immediately those of the whole system); collaborations are voluntary (up to and including breaking rules); participants may be heterogeneous in their computing power or in their goals (not all pedestrians want to go to the same place); participants can choose with whom to collaborate; and they are not assumed subject to permission or synchronisation by a central authority.

These decentralised systems can be very complex, and without a centralised authority to control behaviour, it is not immediately obvious why they should display order. Semitopologies are a topology-flavoured mathematics that goes some way to explaining how and in what circumstances they can behave well. Semitopology is also interesting in and of itself, having a rich and interesting theory — one which quickly deviates from standard accounts on topological spaces, because the most interesting semitopologies are rather ill-behaved from the usual viewpoint, as their antiseperation properties mean that they are not Hausdorff. Various antiseperation properties — of which we have considered intertwined / transitively intertwined here, but there are many more — becomes central to the story, as they define participants who should decide the same value in a distributed system that tries to achieve consensus.

It is possible to construct a quite extensive theory of semitopological space based on these ideas [Gab24, Gab25], and to relate these results back to practical systems in ways that are not entirely obvious, including:

1. It can be proved that *any* semitopology partitions into disjoint components whose points are pairwise intertwined within each component. This goes some way to explaining why blockchains tend to exhibit order [Gab24, Theorem 3.5.4, page 32].
2. It can be proved that every semitopology has an actionable *kernel* of participants, such that if they make a decision then all other participants must follow [Gab24, Corollary 11.6.11, page 152]. This can be read as a distributed systems version of Arrow’s theorem [Fey14] (though the proof is different).<sup>13</sup>
3. Semitopological logics can be constructed and used to analyse intertwinedness properties of semitopologies (as documented in [Gab24, Chapter 20]).
4. In ongoing work, we are using these logics used to formally specify, reason about, and debug consensus protocols.

## 4.2 Future work

Semitopologies invite many questions. We mention just a few:

1. What are the natural semitopological notions of path and homotopy?

---

<sup>13</sup>Arrow’s theorem proves that dictators exist; the semitopological result, for decentralised systems, is that dictator-sets exist. So the question is: is the dictator-set small relative to the size of the whole semitopology? If so, then this is a measure that the system may be more centralised than desired.

2. What are natural notions of evolution of a semitopology over time? We ask because in practice, actionable coalitions are not static; they evolve. Thus, it is natural to consider ‘deformations’ of a semitopology over time.
3. Following on from the previous point, suppose we are given a semitopology all of whose points are intertwined — implying, as per Theorem 3.4, that all points must agree where algorithms succeed. How close is this semitopology to one such that *not* all of its points are intertwined — meaning, as discussed, that forking would be possible even where algorithms succeed? This is a question that would be of interest, for example, to users managing a blockchain to make sure it evolves safely.
4. Can logics based on semitopologies be used to accelerate development, and increase confidence in, distributed algorithms, by giving new declarative descriptions of consensus algorithms?

This is current work, and so far it has proven useful: most recently we applied a semitopological modal logic to axiomatise the Paxos consensus algorithm [GZ25], and in ongoing research we have used a more advanced version of the logic to formally specify and identify errors in a proposed industrial protocol (*Heterogeneous Paxos* [SWvRM20, SWvRM21]). We are now using the same techniques to help design its replacement.

## References

- [Dav16] Matthew David, *The Legacy of Napster*, Networked Music Cultures: Contemporary Approaches, Emerging Issues (Raphaël Nowak and Andrew Whelan, eds.), Palgrave Macmillan UK, London, 2016, DOI: 10.1057/978-1-137-58290-4\_4, available online at <https://durham-repository.worktribe.com/output/1642018/> (permalink: <https://web.archive.org/web/20250421093331/https://durham-repository.worktribe.com/OutputFile/1642039>), pp. 49–65.
- [Eng89] Ryszard Engelking, *General topology*, Sigma Series in Pure Mathematics, Heldermann Verlag, 1989.
- [Fey14] Mark Fey, *A straightforward proof of Arrow’s theorem*, Economics Bulletin **34** (2014), no. 3, 1792–1797.
- [Gab24] Murdoch J. Gabbay, *Semitopology: decentralised collaborative action via topology, algebra, and logic*, College Publications, August 2024, ISBN 978-1848904651.
- [Gab25] ———, *Semitopology: a topological approach to decentralised collaborative action*, The Journal of Logic and Computation (2025), <https://doi.org/10.1093/logcom/exae050>.
- [Goo14] L. M. Goodman, *Tezos – a self-amending crypto-ledger (white paper)*, September 2014, <https://tezos.com/whitepaper.pdf>.

- [GZ25] Murdoch J. Gabbay and Luca Zanolini, *A declarative approach to specifying distributed algorithms using three-valued modal logic*, 2025, <https://arxiv.org/abs/2502.00892> (submitted for publication).
- [Rya10] Johnny Ryan, *A history of the internet and the digital future*, Reaktion Books, 2010, ISBN 978-1861897770.
- [Rya11] ———, *How the atom bomb helped give birth to the internet*, <https://arstechnica.com/tech-policy/2011/02/how-the-atom-bomb-gave-birth-to-the-internet/>, 2 2011, Permalink: <http://web.archive.org/web/20240622221756/https://arstechnica.com/tech-policy/2011/02/how-the-atom-bomb-gave-birth-to-the-internet/>.
- [SWvRM20] Isaac Sheff, Xinwen Wang, Robbert van Renesse, and Andrew C. Myers, *Heterogeneous Paxos: Technical report*, 2020, <https://arxiv.org/abs/2011.08253>.
- [SWvRM21] Isaac Sheff, Xinwen Wang, Robbert van Renesse, and Andrew C. Myers, *Heterogeneous Paxos*, 24th International Conference on Principles of Distributed Systems (OPODIS 2020) (Dagstuhl, Germany) (Quentin Bramas, Rotem Oshman, and Paolo Romano, eds.), Leibniz International Proceedings in Informatics (LIPIcs), vol. 184, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, pp. 5:1–5:17.
- [Wil70] Stephen Willard, *General topology*, Addison-Wesley, 1970, Reprinted by Dover Publications.