Fairness and Consensus in an Asynchronous Opinion Model for Social Networks

Jesús Aranda
Universidad del Valle, Colombia

Sebastián Betancourt
Universidad del Valle, Colombia

Juan Fco. Díaz
Universidad del Valle, Colombia

Frank Valencia
CNRS LIX, École Polytechnique de Paris, France
Pontificia Universidad Javeriana Cali, Colombia

Abstract

We introduce a DeGroot-based model for opinion dynamics in social networks. A community of agents is represented as a weighted directed graph whose edges indicate how much agents influence one another. The model is formalized using labeled transition systems, henceforth called *opinion transition systems (OTS)*, whose states represent the agents' opinions and whose actions are the edges of the influence graph. If a transition labeled (i, j) is performed, agent j updates their opinion taking into account the opinion of agent i and the influence i has over j. We study *(convergence to) opinion consensus* among the agents of strongly-connected graphs with influence values in the interval $(0, 1)$. We show that consensus cannot be guaranteed under the standard *strong fairness* assumption on transition systems. We derive that consensus is guaranteed under a stronger notion from the literature of concurrent systems; *bounded fairness*. We argue that bounded-fairness is too strong of a notion for consensus as it almost surely rules out random runs and it is not a constructive liveness property. We introduce a weaker fairness notion, called *m-bounded fairness*, and show that it guarantees consensus. The new notion includes almost surely all random runs and it is a constructive liveness property. Finally, we consider OTS with *dynamic influence* and show convergence to consensus holds under m-bounded fairness if the influence changes within a fixed interval $[L, U]$ with $0 < L < U < 1$. We illustrate OTS with examples and simulations, offering insights into opinion formation under fairness and dynamic influence.

2012 ACM Subject Classification Theory of computation → Social networks

Keywords and phrases Social networks, fairness, DeGroot, consensus, asynchrony

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2024.7

Related Version *Technical Report with proofs*: `https://arxiv.org/abs/2312.12251` [7]

Supplementary Material *Software (Source python code for simulations)*: `https://github.com/promueva/Fairness-and-Consensus-in-Opinion-Models`
archived at `swh:1:dir:3d1d063e991e22e10ce933f8b060dcb8f1703702`

Funding This work is partly supported by the Colombian Minciencias project PROMUEVA, BPIN 2021000100160.

1 Introduction

Social networks have a strong impact on *opinion formation*, often resulting in polarization. Broadly, the dynamics of opinion formation in social networks involve users expressing their opinions, being exposed to the opinions of others, and potentially adapting their own views based on these interactions. Modeling these dynamics enables us to glean insights into how opinions form and spread within social networks.

The models of social learning aim to capture opinion dynamics in social networks [36]. The DeGroot model [14] is one of the most prominent formalisms for social learning and opinion formation dynamics, and it remains a continuous focus of study in social network theory [21]. A given community is represented as a weighted directed graph, known as the *influence graph*, whose edges indicate how much individuals (*agents*) influence one another. Each agent has an opinion represented as a value in $[0,1]$, indicating the strength of their agreement with an underlying proposition (e.g., "*AI poses a threat to humanity*"). Agents repetitively revise their opinions by averaging them with those of their contacts, taking into account the influence each contact holds. (There is empirical evidence validating the opinion formation through averaging of the model in controlled sociological experiments, e.g., [10].) A fundamental theoretical result of the model states that the agents will *converge to consensus* if the influence graph is *strongly connected* and the agents have non-zero self-influence (*puppet freedom*) [21]. The significance of this result lies in the fact that consensus is a central problem in social learning. Indeed, the inability to reach consensus is a sign of a polarized community.

Nevertheless, the DeGroot model makes at least two assumptions that could be overly constraining within social network contexts. Firstly, it assumes that all the agents update their opinions simultaneously (*full synchrony*), and secondly, it assumes that the influence of agents remains the same throughout opinion evolution (*static influence*). These assumptions may hold in some controlled scenarios and render the model tractable but in many real-world scenarios individuals do not update their opinions simultaneously [29]. Instead, opinion updating often occurs *asynchronously*, with different agents updating their opinions at different times. Furthermore, individuals may gain or lose influence through various factors, such as expressing contrarian or extreme opinions [20].

In this paper, we introduce an *asynchronous* DeGroot-based model with *dynamic influence* to reason about opinion formation, building upon notions from concurrency theory. The model is presented by means of labeled transition systems, here called *opinion transition systems (OTS).* The states of an OTS represent the agents' opinions, and the actions (labels) are the edges of the influence graph. All actions are *always* enabled. If a transition labeled with an edge (i,j) is chosen, agent j updates their opinion by averaging it with the opinion of agent i weighted by the influence that this agent carries over j. A *run* of an OTS is an infinite sequence of (chosen) transitions.

We shall focus on the problem of convergence to opinion consensus in runs of the OTS, *assuming* strong connectivity of the influence graph and puppet freedom. For consensus to make sense, all agents should have the chance to update their opinions. Therefore, we need to make *fairness* assumptions about the runs. In concurrency theory, this means requiring that some actions be performed sufficiently often.

We first show that contrary to the DeGroot model, consensus *cannot* be guaranteed for runs of OTS even under the standard *strong fairness* assumption (i.e., that each action occurs infinitely often in the run) [22, 27]. This highlights the impact of asynchronous behavior on opinion formation.

We then consider the well-known notion of *bounded fairness* in the literature on verification of concurrent systems [16]. This notion requires that every action must be performed not just eventually but within some bounded period of time. We show that bounded-fairness guarantees convergence to consensus. This also gives us insight into opinion formation through averaging, i.e., preventing unbounded delays of actions (opinion updates) is sufficient for convergence to consensus.

Nevertheless, bounded fairness does not have some properties one may wish in a fairness notion. In particular, it is not a *constructive liveness* property in the sense of [34, 33]. Roughly speaking, a fairness notion is a constructive liveness property if, while it may require that a particular action is taken sufficiently often, it should not prevent any other action from being taken sufficiently often. Indeed, we will show that preventing unbounded delays implies preventing some actions from occurring sufficiently often.

Furthermore, bounded-fairness is not *random inclusive*. A fairness notion is random inclusive if any random run (i.e., a run where each action is chosen independently with non-zero probability) is *almost surely* fair under the notion. We find this property relevant because we wish to apply our results to other asynchronous randomized models whose runs are random and whose opinion dynamics can be captured as an OTS.

We therefore introduce a new weaker fairness notion, called *m-bounded fairness*, and show that it guarantees consensus. The new notion is shown to be a constructive liveness property and random inclusive. We also show that consensus is guaranteed under m-bounded fairness even if we allow for *dynamic influence* as long as all the changes of influence are within a fixed interval $[L, U]$ with $0 < L < U < 1$.

All in all, we believe that asynchronous opinion updates and dynamic influence provide us with a model more faithful to reality than the original DeGroot model. The fairness assumptions and consensus results presented in this paper show that the model is also tractable and that it brings new insights into opinion formation in social networks. To the best of our knowledge, this is the first work using fairness notions from concurrency theory in the context of opinion dynamics in social networks.

Furthermore, since *m-bounded fairness* is random inclusive, our result extends with dynamic influence the consensus result in [17] for distributed averaging with randomized gossip algorithms. Distributed averaging is a central problem in other application areas, such as decentralized computation, sensor networks and clock synchronization.

Organization. The paper is organized as follows: In Section 2, we introduce OTS and the consensus problem. Initially, to isolate the challenges of asynchronous communication in achieving consensus, we assume static influence. In Section 3, we identify counter-examples, graph conditions, and fairness notions for consensus to give some insight into opinion dynamics. In Section 4, we introduce a new notion of fairness and state our first consensus theorem. Finally, in Section 5, we add dynamic influence and give the second consensus theorem.

The detailed proofs are included in a related technical report [7]. The Python code used to produce OTS examples and simulations in this paper can be found in the following repository: `https://github.com/promueva/Fairness-and-Consensus-in-Opinion-Models`.

2 The Model

In the standard DeGroot model [14], agents update their opinion *synchronously* in the following sense: at each time unit, all the agents (individuals) update simultaneously their current opinion by listening to the current opinion values of those who influence them. This notion of updating may be unrealistic in some social network scenarios, as individuals may listen to (or read) others' opinions at different points in time.

In this section, we introduce an opinion model where individuals update their beliefs asynchronously; one agent at a time updates their opinion by listening to the opinion of one of their influencers.

2.1 Opinion Transition Systems

In social learning models, a *community* is typically represented as a directed weighted graph with edges between individuals (agents) representing the direction and strength of the influence that one has over the other. This graph is referred to as the *Influence Graph*.

▶ **Definition 1** (Influence Graph). *An* influence graph *is a directed weighted graph* $G = (A, E, I)$ *with* $A = \{1, \ldots, \mathbf{n}\}$, $\mathbf{n} > 1$, *the vertices,* $E \subseteq A^2 - Id_A$ *the edges (where* Id_A *is the identity relation on* A*) and* $I : E \to (0, 1]$ *the weight function.*

The vertices in A represent $\mathbf{n}$ agents of a given community or network. The set of edges E represents the (direct) influence relation between agents; i.e., $(i, j) \in E$ means that agent i influences agent j. The value $I(i, j)$, for simplicity written $I_{(i,j)}$ or I_{ij} , denotes the strength of the influence: a higher value means stronger influence.

Similar to the DeGroot-like models in [21], we model the evolution of agents' opinions about some underlying *statement* or *proposition*, such as, for example, "*human activity has little impact on climate change*" or "*AI poses a threat to humanity*".

The *state of opinion* (or *belief state*) of all the agents is represented as a *vector* in $[0, 1]^{|A|}$. If $\mathbf{B}$ is a state of opinion, $\mathbf{B}[i]$ denotes the *opinion* (*belief*, or *agreement*) value of agent $i \in A$ regarding the underlying proposition: the higher the value of $\mathbf{B}[i]$, the stronger the agreement with such a proposition. If $\mathbf{B}[i] = 0$, agent i completely *disagrees* with the underlying proposition; if $\mathbf{B}[i] = 1$, agent i completely *agrees* with the underlying proposition.

The opinion state is updated as follows: Starting from an initial state, at each time unit, one of the agents, say j, updates their opinion taking into account the influence and the opinion of one of their contacts, say i. Intuitively, in social network scenarios, this can be thought of as having an agent j read or listen to the opinion of one of their influencers i and adjusting their opinion $\mathbf{B}[j]$ accordingly.

The above intuition can be realized as a *Labelled Transition System* (LTS) whose set of states is $S = [0, 1]^{|A|}$ and set of *actions* is E.

▶ **Definition 2** (OTS). *An Opinion Transition System (OTS) is a tuple* $M = (G, \mathbf{B}_{\text{init}}, \to)$ *where* $G = (A, E, I)$ *is an influence graph,* $\mathbf{B}_{\text{init}} \in S = [0, 1]^{|A|}$ *is the initial opinion state, and* $\to \subseteq S \times E \times S$ *is a (labelled) transition relation defined thus:* $(\mathbf{B}, (i, j), \mathbf{B}') \in \to$, *written* $\mathbf{B} \xrightarrow{(i,j)} \mathbf{B}'$, *iff for every* $k \in A$,

$$\mathbf{B}'[k] = \begin{cases} \mathbf{B}[j] + (\mathbf{B}[i] - \mathbf{B}[j])I_{ij} & \text{if } k = j \\ \mathbf{B}[k] & \text{otherwise} \end{cases} \quad (1)$$

If $\mathbf{B} \xrightarrow{e} \mathbf{B}'$ *we say that* $\mathbf{B}$ *evolves into* $\mathbf{B}'$ *by performing (choosing or executing) the* action e.

A labeled transition $\mathbf{B} \xrightarrow{(i,j)} \mathbf{B}'$ represents the opinion evolution from $\mathbf{B}$ to $\mathbf{B}'$ when choosing an action represented by the edge (i, j). As a result of this action, agent j updates their opinion as $\mathbf{B}[j] + (\mathbf{B}[i] - \mathbf{B}[j])I_{ij}$, thereby moving closer to the opinion of agent i. Alternatively, think of agent i as pulling the opinion of agent j towards $\mathbf{B}[i]$. The higher the influence of i over j, I_{ij}, the closer it gets. Intuitively, if $I_{ij} < 1$, it means that agent j is *receptive* to agent i but offers certain *resistance* to fully adopting their opinion. If $I_{ij} = 1$, agent j may be viewed as a *puppet* of i who disregards (or forgets) their own opinion to adopt that of i.

▶ Remark 3. In Def. 1, we do not allow edges of the form (j, j). In fact, allowing them would *not* present us with any additional technical issues, and the results in this paper would still hold. The reason for this design choice, however, has to do with clarity about the intended intuitive meaning of a transition. Suppose that $\mathbf{B} \xrightarrow{(i,j)} \mathbf{B}'$. Since $\mathbf{B}'[j] = \mathbf{B}[j] + (\mathbf{B}[i] - \mathbf{B}[j])I_{ij} = \mathbf{B}[j](1 - I_{ij}) + \mathbf{B}[i]I_{ij}$, agent j gives a weight of I_{ij} to the opinion of i and of $(1 - I_{ij})$ to *their own opinion*. Therefore, the weight that j gives to their opinion may change depending on the agent i. Thus, allowing also a fixed weight I_{jj} of agent j to their own opinion may seem somewhat confusing to some readers. Furthermore, for any $\mathbf{B} \in S$ we would have $\mathbf{B} \xrightarrow{(j,j)} \mathbf{B}$ regardless of the value I_{jj} thus making the actual value irrelevant. Notice also we do not require the sum of the influences over an agent to be 1.

2.2 Runs and Consensus

We are interested in properties of opinion systems, such as convergence to consensus and fairness, which are inherent properties of infinite runs of these systems.

▶ **Definition 4** (e-path, runs and words). *An* execution path (e-path) *of an OTS $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$, where $G = (A, E, I)$, is an infinite sequence $\pi = \mathbf{B}_0 e_0 \mathbf{B}_1 e_1 \ldots$ (also written $\mathbf{B}_0 \xrightarrow{e_0} \mathbf{B}_1 \xrightarrow{e_1} \ldots$) such that $\mathbf{B}_t \xrightarrow{e_t} \mathbf{B}_{t+1}$ for each $t \in \mathbb{N}$. We say that e_t is the action performed at time t and that B_t is the state of opinion at time t. Furthermore, if $\mathbf{B}_0 = \mathbf{B}_{\text{init}}$ then the e-path π is said to be a* run *of M.*

*An ω-*word *of M is an infinite sequence of edges (i.e, an element of E^ω). The sequence $w_\pi = e_0 e_1 \ldots$ is the ω-word generated by π. Conversely, given an ω-word $w = e'_0.e'_1 \ldots$ the (unique) run that corresponds to it is $\pi_w = \mathbf{B}_{\text{init}} \xrightarrow{e'_0} \mathbf{B}_1 \xrightarrow{e'_1} \ldots$*

▶ Remark 5. The uniqueness of the run that corresponds to a given ω-word is derived from the fact that an OTS is a deterministic transition system[1]. This gives us a one-to-one correspondence between ω-words and runs, which allows us to abstract away from opinion states when they are irrelevant or clear from the context. In fact, throughout the paper, *we will use the terms ω-words and runs of an OTS interchangeably when no confusion arises.* It is also worth noting that in OTS, any action (edge) can be chosen at any point in an execution path; that is, *all actions are enabled.*

Consensus is a property of central interest in social learning models [21]. Indeed, failure to reach a consensus is often an indicator of polarization in a community.

▶ **Definition 6** (Consensus). *Let $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$ be an OTS with $G = (A, E, I)$ and $\pi = \mathbf{B}_{\text{init}} \xrightarrow{e_0} \mathbf{B}_1 \xrightarrow{e_1} \ldots$ be a run. We say that an agent $i \in A$* converges *to an opinion value $v \in [0, 1]$ in π if $\lim_{t \to \infty} \mathbf{B}_t[i] = v$. The run π* converges to consensus *if all the agents in A converge to the* same *opinion value in π.*

Furthermore, $\mathbf{B}$ is said to be a consensual state *if it is a constant vector; i.e., if there exists $v \in [0, 1]$ such that for every $i \in A$, $\mathbf{B}[i] = v$.*

▶ **Example 7.** Let $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$ where G is the influence graph in Fig. 1a and $\mathbf{B}_{\text{init}} = (0, 0.5, 1)$. If we perform a on $\mathbf{B}_{\text{init}}$ we obtain $\mathbf{B}_{\text{init}} \xrightarrow{a} \mathbf{B}_1 = (0.0, 0.25, 1.0)$.

Consider the word $w = (abcd)^\omega$. Then $\pi_w = \mathbf{B}_{\text{init}} \xrightarrow{a} (0.0, 0.25, 1.0) \xrightarrow{b} (0.125, 0.25, 1.0) \xrightarrow{c} (0.125, 0.625, 1.0) \xrightarrow{d} (0.125, 0.625, 0.8125) \xrightarrow{a} \ldots$. Fig. 1b suggests that π_w indeed converges to consensus (to opinion value 0.5). A more complex example of the evolution of opinions from a randomly generated graph with eleven agents is illustrated in Fig. 1c.

[1] While the actions in a run can be seen as being chosen non-deterministically by a scheduler, an OTS is a *deterministic* transition system in the sense that given a state $\mathbf{B}$ and an action e, there exists a unique state $\mathbf{B}'$ such that $\mathbf{B} \xrightarrow{e} \mathbf{B}'$.

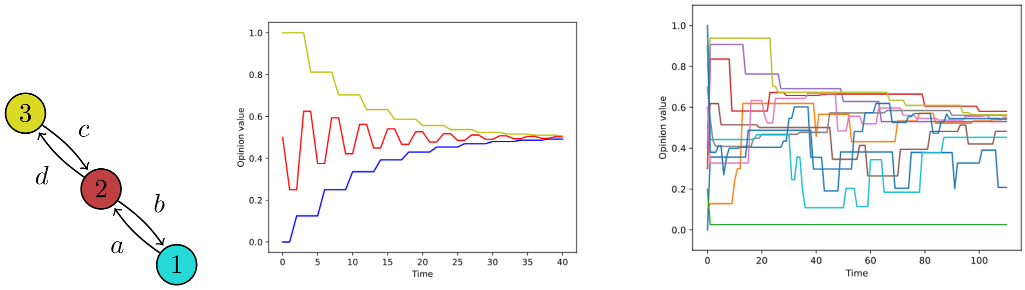

(a) OTS with influence graph with agents $A = \{1, 2, 3\}$, edges $E = \{a, b, c, d\}$, influence $I_e = 1/2$ for all $e \in E$ and $\mathbf{B}_{\text{init}} = (0, 0.5, 1)$.

(b) Opinion evolution for the run that corresponds to $(abcd)^\omega$ of the OTS in Fig. 1a. Each plot corresponds to the opinion evolution of the agent with the same color.

(c) Opinion evolution of a run of an OTS with a $G = (A, E, I)$, $A = \{1, \ldots, 11\}$, $I_e = 0.5$ for each $e \in E$, $\mathbf{B}_{\text{init}} = (0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0)$. Each edge of G was generated with prob. 0.3. The edges of the (partial) run were uniformly chosen from E.

Figure 1 Run examples for OTS in Fig. 1a and randomly-generated OTS in Fig. 1c.

The examples above illustrate runs that may or may not converge to consensus. In the next section, we identify conditions on the influence and topology of graphs and on the runs that guarantee this central property of opinion models.

3 Strong Connectivity, Puppet-Freedom and Fairness

In this section, we discuss graph properties, as well as fairness notions and criteria from the literature on concurrent systems that give us insight into how agents converge to consensus in an OTS. For simplicity, we assume an underlying OTS $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$ with an influence graph $G = (A, E, I)$. We presuppose basic knowledge of graph theory and formal languages.

3.1 Strong Connectivity

As in the DeGroot model, if there are (groups of) agents in G that do not influence each other (directly or indirectly) and their initial opinions are different, these groups may converge to different opinion values. Consider the example in Fig. 2 where the groups of agents $G_1 = \{1, 2\}$ and $G_2 = \{5, 6\}$ do not have external influence (directly or indirectly), but influence the group $G_3 = \{3, 4\}$. Each group is strongly connected within; their members influence each other. The agents in G_1 converge to an opinion, and so do the agents in G_2, but to a different one. Hence, the agents in both groups cannot converge to consensus. The agents in G_3 do not even converge to an opinion because they are regularly influenced by the dissenting opinions of G_1 and G_2.

The above can be prevented by requiring *strong connectivity*, i.e., there must be a path in G from any other to any other. Recall that a *graph path* from i to j of length m in G is a sequence of edges of E of the form $(i, i_1)(i_1, i_2) \ldots (i_{m-1}, j)$, where the agents in the sequence are distinct. We shall refer to graph paths as *g-paths* to distinguish them from e-paths in Def. 4. We say that agent i influences agent j if there is a g-path from i to j in G. The graph G is *strongly connected* iff there is a g-path from any agent to any other in G. Hence, in strongly-connected graphs, all agents influence one another.

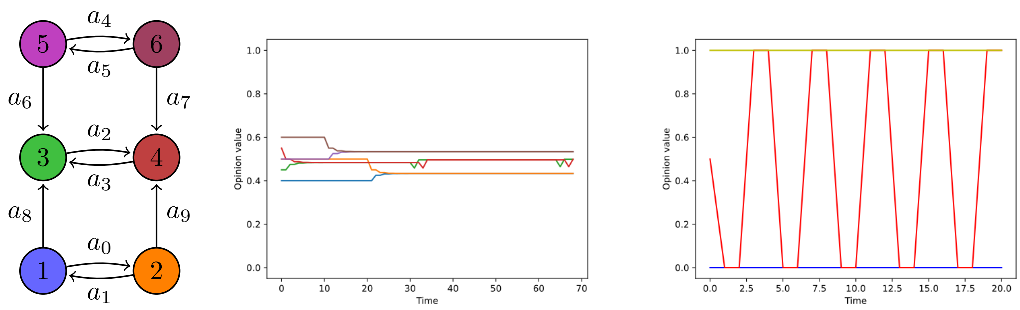

(a) Influence graph with $I_e = 1/2$ for all $e \in E$ and $\mathbf{B}_{\text{init}} = (0.4, 0.5, 0.45, 0.55, 0.5, 0.6)$.

(b) Opinion evolution of the run $((a_2a_3)^5(a_4a_5)^5(a_0a_1)^5(a_6a_7a_8a_9))^\omega$. Each plot corresponds to the opinion evolution of the agent with the same color in Fig. 2a.

(c) Opinion evolution of the run $(abcd)^\omega$ for an OTS with G and $\mathbf{B}_{\text{init}}$ from Fig. 1a but assuming $I_e = 1$ for all $e \in E$.

Figure 2 Run examples for OTS in Fig. 1a and Fig. 2a.

3.2 Puppet-Freedom

Nevertheless, too much influence may prevent consensus. If $\mathbf{B} \xrightarrow{(i,j)} \mathbf{B}'$ and $I_{ij} = 1$, agent j behaves as a *puppet* of i forgetting their own opinion and adopting that of j. Fig. 2c illustrates this for the strongly-connected graph in Fig. 1a but with $I_{ij} = 1$ for each $(i, j) \in E$: Agents 1 and 3 use Agent 2 as a puppet, constantly swaying his opinion between 0 and 1. We therefore say that the influence graph G is *puppet free* if for each $(i, j) \in E$, $I_{ij} < 1$.

3.3 Strong Fairness

In an OTS, if G is strongly connected but a given edge is never chosen in a run (or not chosen sufficiently often), it may amount to not having all agents influence each other in that run, hence preventing consensus. For this reason, we make some fairness assumptions about the runs.

In the realm of transition systems, fairness assumptions rule out some runs, typically those where some actions are not chosen sufficiently often when they are enabled sufficiently often. There are many notions of fairness (see [5, 19, 25] for surveys), but strong fairness is perhaps one of the most representative. As noted above, every action $e \in E$ is always enabled in every run of an OTS. Thus, in our context, strong fairness of a given OTS run (ω-word) amounts to requiring that every action e occurs infinitely often in the run.

▶ **Definition 8** (Strong fairness). *Let w be an ω-word of an OTS. We say that w is* strongly fair *if every $e \in E$ occurs in every suffix of w.*

Notice that the graph from Ex. 7 is strongly connected and puppet free, and the ω-word $w = (abcd)^\omega$ is indeed strongly fair and converges to consensus. Nevertheless, puppet freedom, strong fairness, and strong connectivity are not sufficient to guarantee consensus.

▶ **Proposition 9.** *There exists $(G, \mathbf{B}_{\text{init}}, \rightarrow)$, where G is strongly connected and puppet free, with a strongly-fair run that does not converge to consensus.*

The proof of the existence statement in Prop. 9 is given next.

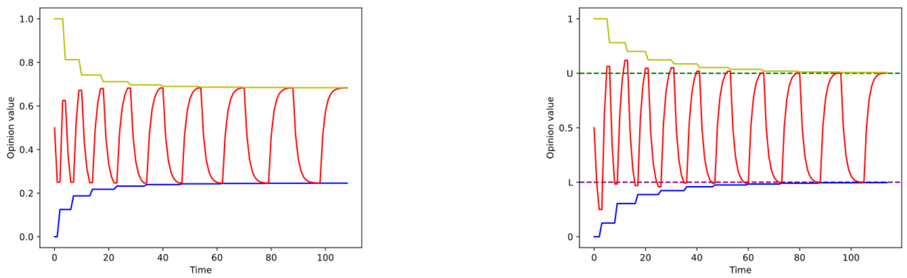

(a) Opinion evolution of the OTS from Fig. 1a for the ω-word $u = (a^n bc^n d)_{n\in\mathbb{N}^+}$.

(b) Opinion evolution of the OTS from Fig. 1a for $U = 0.75$, $L = 0.25$ and the ω-word w from Cons. 10.

Figure 3 Run examples for OTS in Fig. 1a.

▶ **Construction 10** (Counter-Example to Consensus). *Let $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$ be an OTS where G is the strongly-connected puppet-free influence graph in Fig. 1a and $\mathbf{B}_{\text{init}}$ is any state of opinion such that $\mathbf{B}_{\text{init}}[1] < \mathbf{B}_{\text{init}}[2] < \mathbf{B}_{\text{init}}[3]$. We have $A = \{1, 2, 3\}$ and $E = \{a, b, c, d\}$. We construct an ω-word w such that π_w does not converge to consensus with the following infinite iterative process. Let U and L be such that $\mathbf{B}_{\text{init}}[1] < L < \mathbf{B}_{\text{init}}[2] < U < \mathbf{B}_{\text{init}}[3]$.*

Process: *(1) Perform a non-empty sequence of a actions with as many a's as needed until the opinion of Agent 2 becomes smaller than L. (2) Perform the action b. (3) Perform a non-empty sequence of c's with as many c's as needed until the opinion of Agent 2 becomes greater than U. (4) Perform the action d. The result of this iteration is a sequence of the form a^+bc^+d. Repeat steps 1–4 indefinitely.*

The above process produces the ω-sequence $w = w_1 \cdot w_2 \cdot \ldots$ of the form $(a^+bc^+d)^\omega$, where each $w_i = a^{n_i}bc^{m_i}d$ is the result of the i-th iteration of the process and $n_i > 0$ and $m_i > 0$ are the number of a's and c's in such interaction. (The evolution of the opinion of run π_w, with $U = 0.75$, $L = 0.25$ and $\mathbf{B}_{\text{init}} = (0, 0.5, 1)$ is illustrated in Fig. 3b)

Since each action $e \in E$ appears infinitely often in w, w is strongly fair*. Furthermore, right after each execution of Step 2, the opinion of Agent 1 gets closer to L, but it is still smaller than L since the opinion of Agent 2 at that point is smaller than L. For symmetric reasons, the opinion of Agent 3 gets closer to U, but it is still greater than U. Consequently, the opinion of Agent 1 is always below L, while the opinion of Agent 3 is always above U with $L < U$. Therefore, they cannot converge to the same opinion.*

Another ω-word for the OTS in Fig. 1a exhibiting a behavior similar to w in Cons. 10, but whose proof of non-convergence to consensus seems more involved, is $u = (a^n bc^n d)_{n\in\mathbb{N}^+} = u_1 \cdot u_2 \cdot \ldots$, where each $u_n = a^n bc^n d$. (see Fig. 3a). The *delay* in both w and u to execute d after b grows unboundedly due to the growing number of c's. More precisely, let $\#e(v)$ be the number of occurrences of $e \in E$ in a finite sequence v.

▶ **Proposition 11.** *Let $w = w_1 \cdot w_2 \cdot \ldots$ be the ω-word from Cons. 10 where each w_m has the form a^+bc^+d. Then for every $m \in \mathbb{N}$, there exists $t \in \mathbb{N}$ such that $\#c(w_{m+t}) > \#c(w_m)$.*

The above proposition states that the number of consecutive c's in w grows *unboundedly*, and hence so does *the delay for executing d right after executing b*. To prevent this form of *unbounded delay*, we recall in the next section some notions of fairness from the literature that require, at each position of an ω-word, every action to occur within some bounded period of time.

3.4 Bounded Fairness

We start by introducing some notation to give a uniform presentation of some notions of fairness from the literature. We assume $|E| > 1$; otherwise, all the fairness notions are trivial.

A word w is a possibly infinite sequence over E. A *subword* of w is either a suffix of w or a prefix of some suffix of w. Let κ be an ordinal from the set $\omega + 1 = \mathbb{N} \cup \{\omega\}$ where ω denotes the first infinite ordinal. A *κ-word* is a word of length κ. Recall that each ordinal can be represented as the set of all strictly smaller ordinals. We can then view a κ-word $w = (e_i)_{i\in\kappa}$ as a function $w : \kappa \to E$ such that $w(i) = e_i$ for each $i \in \kappa$. A κ-word w is *complete* if $w(\kappa) = E$ (where $w(\kappa)$ denotes the image of the function w). A *κ-window* u of w is a *subword* of w of length κ. Thus, if $\kappa = \omega$ then u is a suffix of w, and if $\kappa \in \mathbb{N}$, u can be thought of as a *finite observation* of κ consecutive edges in w. We can now introduce a general notion of fairness parametric in κ.

▶ **Definition 12** (κ-fairness, bounded-fairness). *Let w be an ω-word over E and $\kappa \in \omega + 1$: w is* κ-fair *if every κ-window of w is complete. Furthermore, w is* bounded fair *if it is k-fair for some $k \in \mathbb{N}$.*

Notice that the notion of strong fairness in Def. 8 is obtained by taking $\kappa = \omega$; indeed, w is ω-fair iff every $e \in E$ occurs infinitely often in w. Furthermore, if $\kappa = k$ for some $k \in \mathbb{N}^+$, then we obtain the notion of k-fairness from [16][2]. Intuitively, if w is k-fair, then at any position of w, every $e \in E$ will occur within a window of length k from that position.

It is not difficult to see that ω-fairness is strictly weaker than bounded-fairness, which in turn is strictly weaker than any k-fairness with $k \in \mathbb{N}$. Let $F(\kappa)$ be the set of all ω-words over E that are κ-fair. We have the following sequence of strict inclusions.

▶ **Proposition 13.** *For every $k \in \mathbb{N}$, $F(k) \subset F(k+1) \subset (\bigcup_{\kappa\in\mathbb{N}} F(\kappa)) \subset F(\omega)$.*

▶ **Example 14.** Let us consider the *fair word* w from Cons. 10, the counter-example to consensus. From Prop. 11, the delay for executing action d immediately after executing action b increases without bound. Thus, for every k, there must be a non-complete k-window u of w such that d does not occur in u. Consequently, w is not bounded fair.

Not only does bounded fairness rule out the counter-example in Cons. 10, but it also guarantees consensus, as shown later, for runs of OTS with strongly-connected, puppet-free influence graphs. Nevertheless, it may be too strong of a requirement for consensus. We, therefore, introduce a weaker notion that satisfies the following criteria and guarantees consensus.

Some Fairness Criteria

Let us briefly discuss some fairness criteria and desirable properties that justify our quest for a weaker notion of fairness that guarantees consensus. An in-depth discussion about criteria for fairness notions, from which we drew some inspiration, can be found in [34, 33, 19, 5].

Machine Closure. Following [1, 26] one of the most important criteria that a notion of fairness must meet is *machine closure* (also called *feasibility* [5]). Fairness properties are properties of infinite runs; hence, a natural requirement is that any finite partial run must have the chance to be extended to a fair run. Thus, we say that a notion of fairness is *machine closed* if every finite word u can be extended to a fair ω-word $u \cdot w$.

[2] This notion is different from the notion of k-fairness from [9]

Clearly, k-fairness with $k \in \mathbb{N}$ is not machine closed; e.g., the word $c^k d$ with $E = \{c, d\}$ cannot be extended to a k-fair ω-word. Nevertheless, bounded fairness is machine closed: Each k-word u can be extended to a $(k+m)$-fair word $u \cdot (e_1 \dots e_m)^\omega$ assuming $E = \{e_1, \dots, e_m\}$.

Constructive Liveness. According to [34], a notion of fairness may require that a particular action is taken sufficiently often, but it should not prevent any other actions from being taken sufficiently often. This concept is formalized in [34, 33] in a game-theoretical scenario, reminiscent of a Banach-Mazur game [28], involving an infinite interaction between a scheduler and an opponent. The opponent initiates with a word w_0, then the scheduler appends a finite word w_1 to w_0. This pattern continues indefinitely, resulting in an ω-word $w = w_0 \cdot w_1 \cdot w_2 \dots$. A given fairness notion is said to be a *constructive liveness* property if, regardless of what the opponent does, the scheduler can guarantee that the resulting ω-word is fair under the given notion.

The notion of *bounded fairness is not a constructive liveness property*. If an ω-word is bounded fair, it is k-fair for some $k \geq |E| > 1$. Let $c \in E$ and take as the strategy of the opponent to choose in each of their turns $w_n = c^n$. Since $|E| > 1$, then w_{2k} cannot be a complete k-window. Therefore, the resulting $w = w_0 \cdot w_1 \cdot w_2 \dots$ is not bounded fair, regardless of the strategy of the scheduler.

It is worth noticing that the above opponent's strategy is reminiscent of our procedure to construct an ω-sequence in Cons. 10 using the unbounded growth of c's to prevent consensus.

Random Words. Consider a word $e_0 e_1 \dots$ where each edge or action $e_n = (i, j)$ is chosen from E *independently* with probability $p_{(i,j)} > 0$. Let us refer to such kinds of sequences as *random* words. We then say that a given notion of fairness is *random inclusive* if every random ω-word is *almost surely* (i.e., with probability one) fair under the given notion.

It follows from the Second Borel–Cantelli lemma[3] that every random word is *almost surely* strongly fair. Nevertheless, the notion of bounded fairness fails to be random inclusive: If a word is bounded fair, it is k-fair for some $k \geq |E|$, and thus it needs to have the form $w_0 \cdot w_1 \dots$ where each w_m is a complete k-window. Since $1 < |E|$, the probability that a random k window is complete is strictly smaller than 1. Therefore, the probability of a random word having an infinite number of *consecutive* complete k-windows is 0.

Random words are important in simulations of our model (see Fig. 1c). Furthermore, having a notion of fairness that is random inclusive and guarantees consensus will allow us to derive and generalize consensus results for randomized opinion models, such as gossip algorithms [17]. We elaborate on this in the related work. We now introduce our new notion of fairness.

4 A New Notion of Bounded Fairness

A natural way to relax bounded fairness to satisfy constructive liveness and random inclusion is to require that the complete k-windows need only appear infinitely often: i.e., an ω word w is said to be *weakly bounded* fair if there exists $k \in \mathbb{N}$ such that every suffix of w has a k-window. Nevertheless, as it will be derived later, weak bounded fairness is not sufficient to guarantee consensus.

[3] The lemma states that if the sum of the probabilities of an infinite sequence of events $E_0 E_1 \dots$ that are independent is infinite, then the probability of infinitely many of those events occurring is 1 [31]. Here, each event E_k expresses that the edge e_k occurs at time k and these events are independent because each edge (i, j) in a random word is chosen independently with probability $p_{(i,j)} > 0$.

It turns out that, to guarantee consensus, it suffices to require that a large enough number m of *consecutive* complete k-windows appear infinitely often. These consecutive windows are referred to as multi-windows.

▶ **Definition 15** ((m, κ) multi-window)**.** *Let w be an ω-word over E, $m \in \mathbb{N}^+$ and $\kappa \in \omega + 1$. We say that w* has an (m, κ) multi-window *if there exists a subword u of w of the form $u = w_1 \cdot w_2 \cdot \ldots \cdot w_m$ where each w_i is a κ-window of w. Furthermore, if each w_n in u is complete, we say that w* has a complete (m, κ) *multi-window. If it exists, the word u is called an* (m, κ) multi-window of w*.*

Notice that because of the concatenation of windows in Def. 15, by construction, no ω-word has a (m, ω) multi-window with $m > 1$: If $\kappa = \omega$ then $m = 1$. In this case, the multi-window is just a window of infinite length of w, i.e., a suffix of w.

▶ **Definition 16** ((m, κ)-fairness)**.** *Let w be an ω-word over E, $m \in \mathbb{N}^+$ and $\kappa \in \omega + 1$. We say that w is* (m, κ)-fair *if every suffix of w has a complete (m, κ) multi-window. We say that w is* m-consecutive bounded fair*, or* m-bounded fair*, if it is (m, k)-fair for some $k \in \mathbb{N}$.*

Clearly, w is ω-fair iff it is $(1, \omega)$-fair, and w is weakly bounded fair iff it is 1-bounded ω-fair. Let $F(m, \kappa)$ and $F(\kappa)$ be the sets of ω-words that are (m, κ)-fair and κ-fair, respectively. We have the following sequence of strict inclusions (assume $k, m \in \mathbb{N}^+$):

▶ **Proposition 17.** $F(k) \subset F(m+1, k) \subset F(m, k) \subset (\bigcup_{\kappa \in \mathbb{N}} F(m, \kappa)) \subset F(1, \omega) = F(\omega)$.

Compliance with Fairness Criteria. Let us consider the criteria for fairness in the previous section. The notion of m-bounded fairness is machine closed since bounded fairness is stronger than m-bounded fairness (Prop. 13 and Prop. 17) and bounded fairness is machine closed.

It is also a constructive liveness property since (m, k) fairness, for $k \geq |E|$, is stronger than m-bounded fairness (Prop. 17), and it is also a constructive liveness property: A winning strategy for the scheduler is to choose a complete (m, k)-window at each one of its turns.

Similarly, m-Bounded Fairness is random inclusive since the stronger notion (m, k)-Fairness is random inclusive for $k \geq |E|$. In a random ω-word $w = w_0 \cdot w_1 \ldots$ where each w_n is a $(m \times k)$-window, the probability that w_n is a complete (m, k)-multi-window is non-zero and independent. Thus again, by the Second Borel–Cantelli lemma, almost-surely w has infinitely many complete (m, k) multi-windows, i.e., it is almost-surely (m, k)-fair.

4.1 Consensus Theorem

We can now state one of our main theorems: m-bounded fairness guarantees consensus in strongly-connected, puppet-free graphs.

▶ **Theorem 18** (Consensus under m-bounded fairness)**.** *Let $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$ be an OTS where G is a strongly-connected, puppet-free influence graph. For every run π of M, if w_π is m-bounded fair and $m \geq |A| - 1$, then π converges to consensus.*

▶ Remark 19. A noteworthy corollary of Th. 18 is that, under the same assumptions of the theorem, if w_π is a bounded fair (a random ω-word), then π converges to consensus (π almost surely converges to consensus). This follows from the above theorem, Prop. 13, Prop. 17 and the fact that m-bounded fairness is random inclusive.

A proof of Th. 18 is given in the technical report [7]. Let us give the main intuitions here.

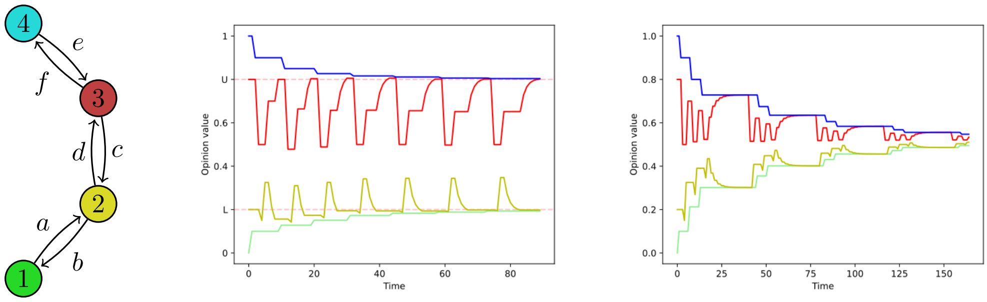

(a) OTS with $I_e = 1/2$ for every edge $e \in E$ and $\mathbf{B}_{\text{init}} = (0.0, 0.2, 0.8, 1.0)$.

(b) Opinion evolution of the 1-bounded fair ω-word w in Cons. 21 with $U = 0.8$ and $L = 0.2$.

(c) Opinion evolution of the 3-bounded fair ω-word $((bfdace)^3 a^{10} e^{10})^\omega$.

Figure 4 Examples of an m-bounded fair runs. In Fig. 4b and 4c, each plot corresponds to the opinion of the agent with the same color in Fig. 4a.

Proof sketch. The proof focuses on the evolution of maximum and minimum opinion values. The sequences of maximum and minimum opinion values in a run, $\{\max \mathbf{B}_t\}_{t\in\mathbb{N}}$ and $\{\min \mathbf{B}_t\}_{t\in\mathbb{N}}$, can be shown to be (bounded) monotonically non-increasing and non-decreasing, respectively, so they must converge to some opinion values, say U and L with $L \le U$.

We must then argue that $L = U$ (this implies convergence to consensus of π by the Squeeze Theorem [32]). Since w_π is m-bounded fair with $m \ge |A| - 1$, after performing all the actions of an (m, k) multi-window of w_π, for some $k \ge |E|$, all the agents of A would have influenced each other. In particular, the agents holding the maximum and minimum opinion values, say agents i and j. To see this, notice that since G is strongly connected, there is a path from i to j, $a_1 \dots a_l$ with length $l \le |A| - 1$. Thus, after performing the first complete k-window of the (m, k)-multi-window, a_1 must be performed, after performing the second complete k-window, a_2 must be performed and so on. Hence, after performing all the actions of the multi-window, i would have influenced j. It can be shown that their mutual influence causes them to decrease their distance by a positive constant factor (here, the puppet freedom assumption is needed). Since the w_π is m-fair, there are infinitely many (m, k)-windows to be performed, and thus the sequences of maximum and minimum opinion values converge to each other, i.e., $U = L$. ◀

It is worth pointing out that without the condition $m \ge |A| - 1$ in Th. 18, we cannot guarantee consensus. Fig. 4c illustrates an m-bounded fair run, for $m = |A| - 1$, of an OTS with 4 agents that converges to consensus. Nevertheless, the following run construction shows that for $m = |A| - 3$, we can construct an m-bounded fair run that fails to converge to consensus (the run is illustrated in Fig. 4b). It also shows that weak bounded fairness, i.e., 1-bounded fairness, is not sufficient to guarantee convergence to consensus. We do not have a counter-example or a proof for $m = |A| - 2$.

▶ **Proposition 20.** *There exists $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$, where $G = (A, E, I)$ is a strongly connected, puppet-free graph, with an m-bounded fair ω-word w, $m = |A| - 3$, such that π_w does not converge to consensus.*

The proof of the above proposition is given in the following construction.

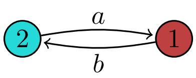

(a) $\mathbf{B}_{\text{init}} = (0.0, 1.0)$ and if $\mathbf{B}[1] = \mathbf{B}[2]$ then $I_a^{\mathbf{B}} = I_b^{\mathbf{B}} = 0.5$, otherwise $I_a^{\mathbf{B}} = \left[\frac{U-\mathbf{B}[2]}{2(\mathbf{B}[1]-\mathbf{B}[2])}\right]_0^1, I_b^{\mathbf{B}} = \left[\frac{L-\mathbf{B}[1]}{2(\mathbf{B}[2]-\mathbf{B}[1])}\right]_0^1$.

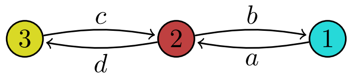

(b) $\mathbf{B}_{\text{init}} = (0.0, 0.5, 1.0)$, $I_d^{\mathbf{B}} = I_b^{\mathbf{B}} = 0.5$, if $\mathbf{B}[1] = \mathbf{B}[2]$ then $I_a^{\mathbf{B}} = 0.5$, if $\mathbf{B}[2] = \mathbf{B}[3]$ then $I_c^{\mathbf{B}} = 0.5$, otherwise $I_a^{\mathbf{B}} = \left[\frac{\frac{1}{2}(\mathbf{B}[1]+L)-\mathbf{B}[2]}{\mathbf{B}[1]-\mathbf{B}[2]}\right]_0^1$, $I_c^{\mathbf{B}} = \left[\frac{\frac{1}{2}(\mathbf{B}[3]+U)-\mathbf{B}[2]}{\mathbf{B}[3]-\mathbf{B}[2]}\right]_0^1$.

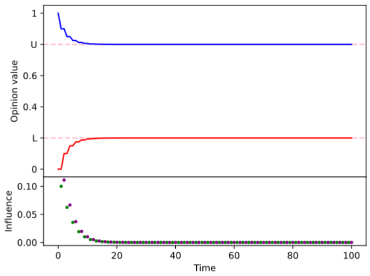

(c) Opinion and influence evolution of the ω-word $(ab)^\omega$. Each plot corresponds to the opinion of the agent with the same color in Fig. 5a. The influences $I_a^{\mathbf{B}}$ and $I_b^{\mathbf{B}}$ are plotted in green and purple.

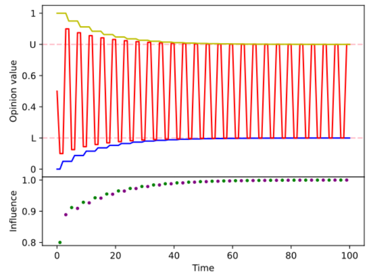

(d) Opinion and influence evolution of the ω-word $(a\,b\,c\,d)^\omega$. Each plot corresponds to the opinion of the agent with the same color in Fig. 5b. The influences $I_a^{\mathbf{B}}$ and $I_c^{\mathbf{B}}$ are plotted in green and purple.

Figure 5 Plots for DOTS in Fig. 5a and Fig. 5b with $U = 0.8$ and $L = 0.2$.[4]

▶ **Construction 21** (Counter-Example to Consensus for m-bounded fairness with $m \leq |A| - 3$)**.** *Suppose that $M = (G, \mathbf{B}_{\text{init}}, \rightarrow)$ where G is the strongly-connected, puppet-free, influence graph in Fig. 4a and $\mathbf{B}_{\text{init}}$ is any state of opinion such that $\mathbf{B}_{\text{init}}[1] < \mathbf{B}_{\text{init}}[2] < \mathbf{B}_{\text{init}}[3] < \mathbf{B}_{\text{init}}[4]$. We have $A = \{1, 2, 3, 4\}$ and $E = \{a, b, c, d, e, f\}$. We construct an ω-word w such that π_w does not converge to consensus with the following infinite iterative process. Let U and L be such that $\mathbf{B}_{\text{init}}[2] \leq L < U \leq \mathbf{B}_{\text{init}}[3]$.*

Process: *(1) Perform the sequence of actions bfdace. (2) Perform a sequence of a actions with as many a's as needed until the opinion of Agent 2 becomes smaller than L. (3) Perform a sequence of e's with as many e's as needed until the opinion of Agent 3 becomes greater than U. The result of this iteration is a sequence of the form $bfdace \cdot a^* e^*$. Repeat steps 1-3 indefinitely.*

The above process produces the ω-sequence $w = v \cdot w_1 \cdot v \cdot w_2 \cdot \ldots$ of the form $(bfdace\, a^* e^*)^\omega$ where $v = bfdace$ and $w_i = a^{m_i} e^{n_i}$ are results of the i-th iteration of the process, and $n_i \geq 0$ and $m_i \geq 0$ are the number of a's and e's in each w_i. (The opinion evolution of run π_w, with $L = 0.2$, $U = 0.8$ and $\mathbf{B}_{\text{init}} = (0.0, 0.2, 0.8, 1.0)$ is illustrated Fig. 4b)

Since the subword v is a complete (1,6)-multi-window and appears infinitely often in w, w is m-bounded fair *for $m = |A| - 3 = 1$. Furthermore, right after each execution of edge f in step 1, the opinion of Agent 1 gets closer to L, but it is still smaller than L since the opinion of Agent 2 at that point is smaller than L. For symmetric reasons, after action b, the opinion of Agent 4 gets closer to U, but it is still greater than U since the opinion of Agent 3 at that point is greater than U. Consequently, the opinion of Agent 1 is always below L, while the opinion of Agent 4 is always above U with $L < U$. Therefore, they cannot converge to the same opinion.*

5 Dynamic Influence

The static weights of the influence graph of an OTS imply that the influence that each individual has on others remains constant throughout opinion evolution. However, in real-life scenarios, the influence of individuals can vary depending on many factors, in particular the state of opinion (or opinion climate). Indeed, individuals may gain or lose influence based on the current opinion trend or for expressing dissenting and extreme opinions, among others.

To account for the above form of dynamic influence, we extend the weight function $I : E \to (0,1]$ of the influence graph $G = (A, E, I)$ as a function $I : E \times [0,1]^{|A|} \to [0,1]$ on edges and the state of opinion. The resulting graph is said to have *dynamic influence*.

▶ **Definition 22** (Dynamic OTS). *A Dynamic OTS (DOTS) is a tuple* $(G, \mathbf{B}_{\text{init}}, \to)$ *where* $G = (A, E, I)$ *has dynamic influence* $I : E \times [0,1]^{|A|} \to [0,1]$*. We write* $I^{\mathbf{B}}_{ij}$ *for* $I((i,j), \mathbf{B})$*. The labeled transition* $\to$ *is defined as in Def. 2 but replacing* I_{ij} *with* $I^{\mathbf{B}}_{ij}$ *in Eq. 1.*

The notions of runs, words, e-paths, and related notions for DOTS remain the same as those for OTS (Def. 4). Let us consider some examples of dynamic influence.

Confirmation Bias. Under *confirmation bias* [8], an agent j is more influenced by those whose opinion is closer to theirs. The function $I^{\mathbf{B}}_{ij} = 1 - |\, \mathbf{B}[j] - \mathbf{B}[i] \,|$ captures a form of confirmation bias; the closer the opinions of i and j, the stronger the influence of i over j.

Bounded Influence. Nevertheless, if we allow dynamic influence that can converge to 0 in a given run $\mathbf{B}_{\text{init}} \xrightarrow{e_0} \mathbf{B}_1 \xrightarrow{e_1} \ldots$, i.e, if $\lim_{t\to\infty} I^{\mathbf{B}_t}_{i,j} = 0$, we may reduce indefinitely influence and end up in a situation similar to non-strong connectivity of the graph, thus preventing consensus as in Section 3.1 (Fig. 2). Analogously, if $\lim_{t\to\infty} I^{\mathbf{B}_t}_{i,j} = 1$, we may end up in puppet situations preventing consensus like in Section 3.2 (Fig. 2c). Both situations are illustrated in the DOTS in Fig. 5. To prevent them, we bound the dynamic influences.

▶ **Definition 23** (Bounded Influence). *A DOTS* $(G, \mathbf{B}_{\text{init}}, \to)$ *with* $G = (A, E, I)$ has bounded influence *if there are constants* $I_L, I_U \in (0,1)$ *such that for each* $\mathbf{B} \in [0,1]^{|A|}$*,* $(i,j) \in E$*, we have* $I^{\mathbf{B}}_{i,j} \in [I_L, I_U]$*.*

The previous form of confirmation bias influence $I^{\mathbf{B}}_{ij} = 1 - |\mathbf{B}[j] - \mathbf{B}[i]|$ is not bounded. Nevertheless, the linear transformation $I_L + (I_U - I_L)I^{\mathbf{B}}_{ij}$ can be used to scale any unbounded influence $I^{\mathbf{B}}_{ij}$ into a bounded one in $[I_L, I_U]$ while preserving its shape.

We conclude with our other main theorem, whose proof is given in the technical report [7].

▶ **Theorem 24** (Consensus with bounded influence). *Let* $M = (G, \mathbf{B}_{\text{init}}, \to)$ *be a DOTS where* G *is a strongly-connected, influence graph. Suppose that* M *has bounded influence. For every run* π *of* M*, if* w_π *is* m*-bounded fair with* $m \geq |A| - 1$*, then* π *converges to consensus.*

The result generalizes Th. 18 to dynamic bounded influence. Therefore, in strongly-connected and dynamic bounded influence graphs, convergence to consensus is guaranteed for all runs that are m-bounded fair, which include each random run almost surely.

---

[4] We use a clamp function for $[0,1]$ defined as $[r]^1_0 = \min(\max(r, 0), 1)$ for every $r \in \mathbb{R}$.

6 Conclusions and Related Work

We introduced a DeGroot-based model with asynchronous opinion updates and dynamic influence using labelled transition systems. The model captures opinion dynamics in social networks more faithfully than the original DeGroot model. The fairness notions studied and the consensus results in this paper show that the model is also tractable and brings new insights into opinion formation in social networks. To our knowledge, this is the first work that uses fairness notions from concurrent systems in the context of DeGroot-based models.

There is a great deal of work on DeGroot-based models for social learning (e.g., [4, 13, 12, 38, 37, 15, 11]). We discuss work with asynchronous updates and dynamic influence, which is the focus of this paper. The work [15] introduces a version of the DeGroot model in which self-influence changes over time, while the influence on others remains the same. The works [11, 12] explore convergence and stability, respectively, in models where influences change over time. The works mentioned above do not take into account asynchronous communication, whereas this paper demonstrates how asynchronous communication, when combined with dynamic influence, can prevent consensus.

Recent works on gossip algorithms [17, 30, 2, 35] study consensus with asynchronous communications for distributed averaging and opinion dynamics. The work in [30] studies *reaching* consensus (in finite time) rather than *converging* to consensus. The works [2, 35] consider undirected cliques rather than directed graphs as influence graphs. The closest work is [17], which states consensus for random runs in directed strongly connected graphs but unlike our case all edges have the same fixed weight $q \in (0, 1)$ (i.e., they assume static influence with the same influence value for all edges). The dynamics of asymmetric gossip updates in [17] can indeed be captured as OTS, and their random runs are almost-surely m-bounded fair. Consequently, our work generalizes the consensus result in [17] by extending it to graphs with (bounded) dynamic influence and whose edges may have different weights. Furthermore, the framework in [17] does not address fairness notions which are the focus and the main novelty of our work.

The work [19] discusses probabilistic fairness as a method equally strong as strong fairness to prove liveness properties, where a liveness property is characterized by a set of states such that a run holds this property iff the run reaches a state of this set. However, the property of *(convergence to) consensus* (Def. 6) does not correspond to this notion of liveness since it is not about *reaching* a specific set of states but about *converging* to a consensual state. In fact, unless there are puppets or the initial state of a run is already a consensual state, consensus is never reached in finite time in our model.

Bounded fair ω-words can be characterized by Prompt Buchi Automata (PBW) [3]. Indeed, the set of bounded-fair words of an OTS can be characterized as the language of PBW. Hence, the closure properties of these automata may prove valuable for future developments of our work. It would also be interesting to see in future work whether or not the $m-$bounded fair words of an OTS can be characterized as the language of a PBW (or of an elegant variant of it).

In future work, we plan to study the actual value of consensus in a given system. This may provide information about the most influential agents. We also plan to study how actions can be scheduled (or manipulated), while preserving the fairness assumptions, to converge more quickly or slowly to a consensus, or to a given consensus value. For example, giving priority to edges whose agents have a greater opinion disagreement, while respecting fairness assumptions. We may build on previous work on priorities in concurrent communications [6] for this purpose.

Finally, we plan to extend our model with agents that can learn by exchanging beliefs, lies, and information, by building upon our work in concurrent constraint programming (e.g. [24, 23, 18]).

References

1 Martín Abadi and Leslie Lamport. An old-fashioned recipe for real time. In *Real-Time: Theory in Practice*. Springer Berlin Heidelberg, 1992.

2 Emerico Aguilar and Yasumasa Fujisaki. Opinion dynamics via a gossip algorithm with asynchronous group interactions. *Proceedings of the ISCIE International Symposium on Stochastic Systems Theory and its Applications*, 2019:99–102, 2019. doi:10.5687/sss.2019.99.

3 Shaull Almagor, Yoram Hirshfeld, and Orna Kupferman. Promptness in ω-regular automata. In *Automated Technology for Verification and Analysis*. Springer Berlin Heidelberg, 2010.

4 Mário S. Alvim, Bernardo Amorim, Sophia Knight, Santiago Quintero, and Frank Valencia. A Multi-agent Model for Polarization Under Confirmation Bias in Social Networks. In *41th International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE)*, 2021. URL: https://inria.hal.science/hal-03740263.

5 Krzysztof R. Apt, Nissim Francez, and Shmuel Katz. Appraising fairness in languages for distributed programming. In *14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, POPL 1987*, 1987. doi:10.1145/41625.41642.

6 Jesús Aranda, Frank D. Valencia, and Cristian Versari. On the expressive power of restriction and priorities in CCS with replication. In *Foundations of Software Science and Computational Structures, FoSSaCS 2009*, volume 5504 of *Lecture Notes in Computer Science*, 2009. doi:10.1007/978-3-642-00596-1_18.

7 Jesús Aranda, Sebastián Betancourt, Juan Fco. Díaz, and Frank Valencia. Fairness and consensus in opinion models (technical report), 2024. arXiv:2312.12251.

8 Elliot Aronson, Timothy Wilson, and Robin Akert. *Social Psychology*. Upper Saddle River, NJ : Prentice Hall, 7 edition, 2010.

9 Eike Best. Fairness and conspiracies. *Information Processing Letters*, 18(4):215–220, 1984. doi:10.1016/0020-0190(84)90114-5.

10 Arun G Chandrasekhar, Horacio Larreguy, and Juan Pablo Xandri. Testing models of social learning on networks: Evidence from a lab experiment in the field. Working Paper 21468, National Bureau of Economic Research, August 2015. doi:10.3386/w21468.

11 S. Chatterjee and E. Seneta. Towards consensus: Some convergence theorems on repeated averaging. *Journal of Applied Probability*, 14(1):89–97, 1977. doi:10.2307/3213262.

12 Zihan Chen, Jiahu Qin, Bo Li, Hongsheng Qi, Peter Buchhorn, and Guodong Shi. Dynamics of opinions with social biases. *Automatica*, 106:374–383, 2019. doi:10.1016/j.automatica.2019.04.035.

13 Pranav Dandekar, Ashish Goel, and David Lee. Biased assimilation, homophily and the dynamics of polarization. *Proceedings of the National Academy of Sciences of the United States of America*, 110, March 2013. doi:10.1073/pnas.1217220110.

14 Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 1974. URL: http://www.jstor.org/stable/2285509.

15 Peter M. DeMarzo et al. Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, 118(3):909–968, 2003. URL: http://www.jstor.org/stable/25053927.

16 Nachum Dershowitz, D. N. Jayasimha, and Seungjoon Park. *Bounded Fairness*, pages 304–317. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. doi:10.1007/978-3-540-39910-0_14.

17 F. Fagnani and S. Zampieri. Asymmetric randomized gossip algorithms for consensus. *IFAC Proceedings Volumes*, 2008. doi:10.3182/20080706-5-KR-1001.01528.

18 Moreno Falaschi, Carlos Olarte, Catuscia Palamidessi, and Frank D. Valencia. Declarative diagnosis of temporal concurrent constraint programs. In *Logic Programming. ICLP 2007*, 2007. doi:10.1007/978-3-540-74610-2_19.

19 Rob Van Glabbeek and Peter Höfner. Progress, justness, and fairness. *ACM Computing Surveys*, 52(4):1–38, August 2019. doi:10.1145/3329125.

20 Elizabeth B. Goldsmith. *Introduction to Social Influence: Why It Matters*, pages 3–22. Springer International Publishing, Cham, 2015. doi:10.1007/978-3-319-20738-4_1.

21 Benjamin Golub and Evan Sadler. Learning in social networks. *Available at SSRN 2919146*, 2017.

22 Orna Grumberg, Nissim Francez, Johann A. Makowsky, and Willem P. de Roever. A proof rule for fair termination of guarded commands. *Information and Control*, 66(1):83–102, 1985. doi:10.1016/S0019-9958(85)80014-0.

23 Michell Guzmán, Stefan Haar, Salim Perchy, Camilo Rueda, and Frank Valencia. Belief, Knowledge, Lies and Other Utterances in an Algebra for Space and Extrusion. *Journal of Logical and Algebraic Methods in Programming*, September 2016. doi:10.1016/j.jlamp.2016.09.001.

24 Stefan Haar, Salim Perchy, Camilo Rueda, and Frank Valencia. An Algebraic View of Space/Belief and Extrusion/Utterance for Concurrency/Epistemic Logic. In *17th International Symposium on Principles and Practice of Declarative Programming (PPDP 2015)*, 2015. doi:10.1145/2790449.2790520.

25 M.Z. Kwiatkowska. Survey of fairness notions. *Information and Software Technology*, 31(7):371–386, 1989. doi:10.1016/0950-5849(89)90159-6.

26 Leslie Lamport. Fairness and hyperfairness. *Distributed Computing*, 13(4):239–245, November 2000. doi:10.1007/PL00008921.

27 D. Lehmann, A. Pnueli, and J. Stavi. Impartiality, justice and fairness: The ethics of concurrent termination. In Shimon Even and Oded Kariv, editors, *Automata, Languages and Programming*, pages 264–277, Berlin, Heidelberg, 1981. Springer Berlin Heidelberg.

28 R.D. Mauldin. *The Scottish Book: Mathematics from the Scottish Café*. Birkhäuser, 1981. URL: https://books.google.com.co/books?id=gaqEAAAAIAAJ.

29 Hossein Noorazar. Recent advances in opinion propagation dynamics: a 2020 survey. *The European Physical Journal Plus*, 135(6):521, June 2020. doi:10.1140/epjp/s13360-020-00541-2.

30 Guodong Shi, Bo Li, Mikael Johansson, and Karl Henrik Johansson. Finite-time convergent gossiping. *IEEE/ACM Transactions on Networking*, 24(5):2782–2794, 2016. doi:10.1109/TNET.2015.2484345.

31 Albert N. Shiryaev. *Probability-1: Volume 1*. Springer New York, 2016. doi:10.1007/978-0-387-72206-1.

32 Houshang H. Sohrab. *Basic Real Analysis*. Birkhauser Basel, 2nd edition, 2014. doi:10.1007/0-8176-4441-5.

33 Hagen Völzer and Daniele Varacca. Defining fairness in reactive and concurrent systems. *J. ACM*, 59(3), June 2012. doi:10.1145/2220357.2220360.

34 Hagen Völzer, Daniele Varacca, and Ekkart Kindler. Defining fairness. In Martín Abadi and Luca de Alfaro, editors, *CONCUR 2005 – Concurrency Theory*, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

35 Xing Wang, Bingjue Jiang, and Bo Li. Opinion dynamics on social networks. *Acta Mathematica Scientia*, 42(6):2459–2477, November 2022. doi:10.1007/s10473-022-0616-8.

36 Stanley Wasserman and Katherine Faust. Social network analysis in the social and behavioral sciences. In *Social Network Analysis: Methods and Applications*, pages 1–27. Cambridge University Press, 1994.

37 Chen X, Tsaparas P, Lijffijt J, and De Bie T. Opinion dynamics with backfire effect and biased assimilation. *PLoS ONE*, 16(9), 2021. doi:10.1371/journal.pone.0256922.

38 Weiguo Xia, Mengbin Ye, Ji Liu, Ming Cao, and Xi-Ming Sun. Analysis of a nonlinear opinion dynamics model with biased assimilation. *Automatica*, 120:109113, 2020. doi:10.1016/j.automatica.2020.109113.