On Convergence of Concurrent Systems under Regular Interactions

Pavithra Prabhakar, Sayan Mitra, Mahesh Viswanathan

University of Illinois at Urbana-Champaign

Abstract. Convergence is often the key liveness property for distributed systems that interact with physical processes. Techniques for proving convergence (asymptotic stability) have been extensively studied by control theorists. In particular, for the asynchronous model of computation Tsitsiklis [8] provides a set of necessary and sufficient conditions for proving stability and convergence under the assumption that each asynchronous operator (state transition function) is applied infinitely often. This paper generalize these results to obtain necessary and sufficient conditions for systems where the infinite sequence of operators is a member of an arbitrary omega regular language. This enables us to apply our theory to distributed systems with changing communication topology, node failures and joins. We illustrate an application of the new set of conditions in verifying the convergence of a simple (continuous) consensus protocol.

1 Introduction

Convergence or asymptotic stability is a key requirement of many concurrent and distributed systems that interact with physical processes. Roughly, a system A *converges* to a target state x^* if the state of A along infinite executions get closer and closer to x^*, with respect to some topology on the state space X, as time goes to infinity. While termination has been the defacto liveness property of interest for software systems, the more general convergence property becomes relevant for systems with both software and physical components. Examples of such systems include algorithms for mobile robots for forming a spatial pattern, synchronization of coupled oscillators, distributed control algorithms over switching networks [7](see for e.g. [4], [1], [2] and [6]). Convergence may indeed be viewed as a liveness property quantified over a (possibly infinite) sequence of shrinking predicates containing the target state.

Necessary and sufficient conditions for proving convergence of distributed systems which broadly fall under the category of *continuous consensus* have been studied extensively by control theorists for over three decades [7]. Specifically, two types of models of distributed computation have been considered. In the synchronous model, the state of the entire system $x \in X$ evolves according to some difference equation: $x_{k+1} = f(x_k)$ or differential equation $\dot{x} = f(x)$, where $f : X \to X$. Convergence conditions in this case are derived based on the eigenvalues of f. We refer the reader to [7] for a survey of the results of this type. In the asynchronous model, the evolution of the system is specified by a collection of transition functions $\{T_k\}$, where each $T_k : X \to X$, and ans execution of the system is obtained by applying an infinite sequence σ of $T_k's$ to the starting state. In [8], Tsitsiklis has identified a general set of necessary and sufficient conditions for the convergence of executions that satisfy a particular fairness assumption.

Tsitsiklis' condition, informally, is as follows. He requires one to identify a collection of shrinking neighborhoods, indexed by a totally ordered index set, that converges to x^* and satisfies the following properties. First the neighborhoods are required to be invariant, i.e., for any neighborhood U, $T_k(x) \in U$ for every $x \in U$ and every T_k. Second, for every neighborhood U, there must be a transition T_U that takes U to a strictly smaller neighborhood. Tsitsiklis shows that when such a neighborhood "system" exists, the system can be proved to converge to x^* in every execution where each transition T_k is applied *infinitely often*. Moreover, he shows that the convergence of a system also implies the existence of such a neighborhood system.

In this paper, we generalize Tsitsiklis' observations as follows. We identify necessary and sufficient conditions for convergence under executions described by an *arbitrary* ω-regular language, instead of focusing on a particular set of executions that satisfy a specific fairness condition. While this is a philosophically natural extension of Tsitsiklis' investigations, it allows us to model a variety of asynchronous behavior, such as ordered execution of certain events, communication patterns between distributed agents over a dynamically evolving or unreliable communication network, and distributed network with nodes failing and recovering, that are not captured by Tsitsiklis' original formulation.

Our necessary and sufficient condition for convergence is remarkably similar to Tsitsiklis' condition. Let us assume that $\mathcal{A}$ is a Müller automaton that describes the set of valid executions. Once again, we require a collection of shrinking neighborhoods, indexed by a totally ordered index set, that converges to x^*. We also require these neighborhoods to be "invariant". However, since every finite sequence of operations need not be the prefix of a valid execution, our definition of invariance accounts for the state of the automaton $\mathcal{A}$. Next, like Tsitsiklis, we have a condition that ensures "progress towards" x^* is eventually made. This is captured by our insight that edges crossing "cuts" in accepting cycles of $\mathcal{A}$ are traversed infinitely often, and so for every neighborhood set U, there must be some cut that ensures progress. The proof showing that these conditions are sufficient, is very similar to Tsitsiklis' proof. To demonstrate the necessity of these condition for convergence is more challenging primarily because every finite sequence of operations need not be the prefix of a valid execution.

We conclude the paper by demonstrating the application of the new set of conditions to prove the convergence of a simple continuous consensus algorithm. We consider a variety of scenarios ranging from a dynamically evolving communication network, to a situation where nodes in the distributed system can fail and recover.

Related Work. Tsitsiklis' result for the asynchronous model have been extended in several ways. For example, in [7] sufficient conditions have been given for proving convergence of distributed algorithms in which the communication graph of the participating agents is dynamic, but never permanently partitioned. More recently, in [5] sufficient conditions for convergence have been derived for partially synchronous systems where messages may be lost or delayed by some constant but unknown time. All of these constrained executions can be modelled as ω-regular languages, and therefore the results of this paper can be seen as a generalization of these observations.

2 Motivating Example

We model the behavior of a distributed system where agents starting at arbitrary positions on a line communicate based on an underlying dynamic graph to move closer to each other. In addition they can fail and join the system a finite number of times. However when they join they start at the same position in which they originally started. We show that the agents finally converge to a common point. We describe the protocol formally below.

Let $N \in \mathbb{N}$ denote the maximum number of agents that can ever be present in the system. Each agent has a unique identifier from the set $[N] = \{1, 2, \cdots, N\}$. We denote the state variable which stores the position of agent i, $i \in [N]$, by x_i, and it takes values in $\mathbb{R} \cup \{\perp\}$. For any $i \in [N]$, agent i is said to be failed if $x_i = \perp$; otherwise i is alive. We denote the collective states of all agents by vectors x, y etc.

Let $G = (V, E)$ be an undirected graph with $V = [N]$ and $E \subseteq V \times V$. Each vertex in the graph corresponds to an agent in the system. G is the underlying graph that remains fixed throughout our discussion. At a given point in the execution of the system, the actual communication graph G' is the subgraph of G restricted to the alive nodes.

Let us concentrate on a particular initial state (G, c) where $c = (c_1, \cdots, c_N)$. Let the current configuration of the system be (H, x), where $H = (V_H, E_H)$. There are three kinds of operators which can modify the configuration, namely, *join*, *fail* and *move*, and correspond to a node joining the system, a node failing and two nodes communicating to move to their average value.

We say that a configuration has reached a fixpoint if all the unfailed nodes have the same value, that is, (H, x) is a fixpoint if for every $i, j \in [N]$, $i \neq j$, $x_i \neq \perp$ and $x_j \neq \perp$ implies $x_i = x_j$. When (H, x) is a fixpoint, for all $i, j \in [N]$, $\mathit{fail}_j((H, x)) = (H, x)$, $\mathit{join}_j((H, x)) = (H, x)$ and $\mathit{move}_{i,j}((H, x)) = (H, x)$. When (H, x) is not a fixpoint,

- $\mathit{fail}_j((H, x)) = (H', x')$ where $x'_j = \perp$ and $x'_i = x_i$ for $i \neq j$ and $H' = (V'_H, E'_H)$ where $V'_H = V_H - \{j\}$ and $E'_H = E \cap (V'_H \times V'_H)$.
- $\mathit{join}_j((H, x)) = (H', x')$ where $x'_j = c_j$ and $x'_i = x_i$ for $i \neq j$ and $H' = (V'_H, E'_H)$ where $V'_H = V_H \cup \{j\}$ and $E'_H = E \cap (V'_H \times V'_H)$.
- $\mathit{move}_{i,j}((H, x)) = (H, x')$ where $x' = x$ if either $x_i = \perp$ or $x_j = \perp$, otherwise $x'_i = x'_j = (x_i + x_j)/2$ and $x'_k = x_k$ for $k \notin \{i, j\}$.

We note that $move_{i,j}$ is defined only if $(i,j) \in E$, that is, communication between (i,j) is allowed. We want to show that an infinite sequence of operations converges to a point if it contains a finite number of $fail_j$ and $join_j$ operations and the set of edges (i,j) of E such that $move_{i,j}$ occurs infinitely often forms a connected graph. We will see later that this set of sequences is an ω-regular language. We will develop sufficiency conditions for proving such properties, and apply it to this example. Several generalizations of this type of consensus protocol has been presented in the literature (see for e.g. [3]).

3 Preliminaries

3.1 Directed and Undirected graphs

A labelled directed graph (LDG) G is a triple (V, E, Σ), where V is a finite set of vertices, $E \subseteq V \times \Sigma \times V$ is a set of edges and Σ is a finite set of labels. Let $G = (V, E, \Sigma)$ be a LDG. Given $V' \subseteq V$, the restriction of G to V' is given by the LDG $G[V'] = (V', E', \Sigma)$ where $E' \subseteq E$ is the set $\{(u, a, v) \in E \mid u, v \in V'\}$. Given $E' \subseteq E$, $G - E' = (V, E - E', \Sigma)$. A *path* in G is a sequence of edges $e_1 \cdots e_n$ such that $e_i = (q_i, a, q_{i+1})$ for all i. We say that q_{n+1} is *reachable* from q_1. We say that G is *strongly connected* if for every $u, v \in V$, v is reachable from u. A set $V' \subseteq V$ is strongly connected in G if $G[V]$ is strongly connected and is *maximally strongly connected* if in addition for all V'' such that $V' \subset V''$, $G[V'']$ is not strongly connected.

An undirected graph G is a pair (V, E) where $E \subseteq V \times V$ is a symmetric relation. Whenever we refer to a set of edges of an undirected graph it is assumed to be symmetric. A path in G and reachability of a vertex is defined as before. We say that a graph G is connected if every vertex is reachable from every other vertex. A cut in a connected graph G is a non-empty set of edges E such that $G - E$ is not connected. A subgraph of $G = (V, E)$ is a graph $G' = (V', E')$ such that $V' \subseteq V$ and $E' \subseteq E \cap (V' \times V')$.

3.2 Stability and Convergence

Let $\mathcal{X}$ be a set and $T_k : \mathcal{X} \to \mathcal{X}$ for $1 \leq k \leq K$ be a finite collection of functions ("operators"). Let $X_{fp} \subseteq \mathcal{X}$ be a set of common fixpoints, that is, $T_k(x) = x$ for all $1 \leq k \leq K$ and $x \in X_{fp}$. We will denote an infinite sequence of operators by σ. Let $\sigma = a_1 a_2 a_3 \cdots$, then $\sigma(i)$ denotes the i-th element of σ namely a_i, and $Pref(\sigma, i)$ denotes the finite sequence consisting of the first i elements, namely, $a_1 \cdots a_i$. Given an $x \in \mathcal{X}$, we denote the element obtained by applying the first n operators of σ to x in order by $\sigma(x, n)$. Formally $\sigma(x, 0) = x$, $\sigma(x, n) = \sigma(n)(\sigma(x, n-1))$, for $n \geq 1$.

Next we want to define the notion of convergence. We say that starting from x a sequence σ converges to some point in X_{fp} if it moves closer and closer to X_{fp} along σ. To make this notion precise we need to define a neighborhood system around X_{fp}.

Definition 1. *An $\mathcal{X}$-neighborhood system $\mathcal{U}$ around X_{fp} is a collection of subsets of $\mathcal{X}$ such that:*

Property 1 $X_{fp} \subseteq U$, $\forall U \in \mathcal{U}$.

Property 2 *For any $y \in \mathcal{X}$ such that $y \notin X_{fp}$, there exists some $U \in \mathcal{U}$ such that $y \notin U$.*

Property 3 *$\mathcal{U}$ is closed under finite intersections.*

Property 4 *$\mathcal{U}$ is closed under unions.*

We say that $\mathcal{U}$ has a *countable base* if there exists a sequence $\{U_n\}_{n=1}^{\infty}$ of elements of $\mathcal{U}$ such that for every $U \in \mathcal{U}$ there exists some n such that $U_n \subseteq U$.

We say that a sequence $\{x_n\}_{n=1}^{\infty}$ of elements of $\mathcal{X}$ *converges* to X_{fp} with respect to $\mathcal{U}$ if for every $U \in \mathcal{U}$ there exists a positive integer N such that $x_n \in U$, $\forall n \geq N$.

We want to converge not with respect to a single sequence but a set of sequences. Let us fix a set of operators $\Sigma = \{T_1, \cdots, T_k\}$. An infinite sequence of operators from Σ will also be called an infinite word. We will denote the set of all infinite words over Σ by Σ^ω. We will call a subset L of Σ^ω a language over Σ.

Definition 2. *Stability and convergence. Given a neighborhood system $\mathcal{U}$ and $L \subseteq \Sigma^\omega$, we say that L is* stable *with respect to $\mathcal{U}$ if $\forall U \in \mathcal{U}$, $\exists V \in \mathcal{U}$ such that $\forall x \in \mathcal{X}$, $\forall \sigma \in L$, if there exists $n_0 \in \mathbb{N}$ such that $\sigma(x, n_0) \in V$ then $\forall n \geq n_0$, $\sigma(x, n) \in U$.*

We say that L converges to X_{fp} with respect to $\mathcal{U}$ if $\{\sigma(x, n)\}_{n=1}^{\infty}$ converges to X_{fp} for all $x \in \mathcal{X}$ and $\sigma \in L$.

Remark 1 *Our definition of convergence is analogous to asymptotic stability used is control theory. However our definition of stability is slightly stronger than the classical notion of Lyapunov stability in that instead of requiring that for every U there exists a V such that any trajectory starting in V remains within U, we require that if a trajectory starting anywhere enters V, then we remain within U. This stronger condition is equivalent to the weaker condition, when the L we consider is a suffix closed language.*

Example 1 *Let us now try to formalize the convergence of the Example in Section 2. Let the agents have identifiers from $[N]$ and $G = (V, E)$ be the underlying undirected graph which changes when nodes fail and join. The state space $\mathcal{X}$ is the set of all pairs (H, x) where H is a subgraph of G restricted to the nodes i which have not failed. $\mathcal{X} = \{(H, x) \,|\, H = (V_H, E_H), V_H = \{i \,|\, x_i \neq \bot\}, E_H = E \cap (V_H \times V_H)\}$. We take X_{fp} to be the set of all fixpoints, which are configurations in which all the unfailed nodes have the same value. $X_{fp} = \{(H, x) \in \mathcal{X} \,|\, \forall i, j, (x_i \neq \bot, x_j \neq \bot) \Rightarrow x_i = x_j\}$.*

Next we need to define a neighborhood system which satisfies the properties 1, 2, 3 and 4. Before defining the neighborhood, we need to set some notation.

Let $\beta(n) = (1 - 1/(2n^3)$ *when* $n > 0$ *and* $\beta(0) = 0$. *Let* $alive(x)$ *is the size of the set* $\{i \mid x_i \neq \bot\}$. *We define a function* $f : (\mathbb{R} \cup \{\bot\})^N \rightarrow \mathbb{R}_{\geq 0}$ *given by* $f(x) = \sum_{j:x_j \neq \bot}(x_j - M)^2$, *where* $M = \frac{1}{alive(x)} \sum_{j:x_j \neq \bot} x_j$ *when* $alive(x) \neq 0$, *otherwise* $f(x) = 0$. *Let* $\mathbb{I}$ *be the set of all integers. We can now define the neighborhood as:*

$$\mathcal{U} = \{U_i\}_{i \in \mathbb{I}}, \text{where } U_i = \{(H, x) \in \mathcal{X} \mid f(x) \leq \beta(alive(x))^i\}.$$

$\mathcal{U}$ *is a neighborhood system. To see the Property 1 is satisfied, observe that for all* $(H, x) \in X_{fp}$, $f(x) = 0$ *and* $\beta(alive(x))^i \geq 0$. *Hence* $X_{fp} \subseteq U_i$ *for all* i. *Given any* $(H, x) \notin X_{fp}$, $f(x) > 0$. *And* $\beta(n)^i \rightarrow 0$ *as* $i \rightarrow \infty$ *for* $n \geq 0$. *Therefore* $(H, x) \notin U_i$ *for some* i. *In particular we have* $X_{fp} \subset \cdots \subset U_3 \subset U_2 \subset U_1 \subset U_0 \subset U_{-1} \subset U_{-2} \subset U_{-3} \subset \cdots \subset \mathcal{X}$, *and* $\bigcap_{i \in I} U_i = X_{fp}$. *Clearly Properties 3 and 4 are satisfied.*

We want to show that starting from any $(H, x) \in \mathcal{X}$, *we converge to* X_{fp} *on any sequence of operations of join, fail and move with finite number of join and fail operations and such that the moves form a connected component of the alive nodes. Let* $join = \{join_j \mid j \in [N]\}$, $fail = \{fail_j \mid j \in [N]\}$ *and* $move = \{move_{i,j} \mid (i, j) \in E\}$. *Formally* $\Sigma = join \cup fail \cup move$. *We can define a function* $Nodes{-}Alive : \Sigma^* \rightarrow 2^{[N]}$ *which takes a finite sequence of operators and returns the set of nodes alive after applying the operators in the sequence. We will use . for concatenation of two finite sequences or for concatenation of an infinite sequence to the end of a finite sequence.* $Nodes{-}Alive(\epsilon) = [N]$. $Nodes{-}Alive(\sigma.T) = Nodes{-}Alive(\sigma)$ *if* $T \in move$, $= Nodes{-}Alive(\sigma) - \{i\}$ *if* $T = fail_i$ *and* $= Nodes{-}Alive(\sigma) \cup \{i\}$ *if* $T = join_i$. $L_{converge} = \{\sigma \in \Sigma^\omega \mid \exists \sigma_1 \in \Sigma^*, \sigma_2 \in move^\omega, \sigma = \sigma_1\sigma_2, G_{\sigma_1,\sigma_2}$ *is connected*$\}$, *where* $G_{\sigma_1,\sigma_2} = (Nodes{-}Alive(\sigma_1), \{(i, j) \mid move_{i,j} \in \inf(\sigma_2)$ *or* $move_{j,i} \in \inf(\sigma_2)\})$.

Remark 2 *This system is not stable even in the classical sense because starting in any configuration, executing* $join_j$, $j = 1, \cdots, N$, *will result in the initial configuration. So given a* U *which is sufficiently small, for every* V, *there exists an* $x \in V$ *and* σ *such that* σ *takes* x *out of* U. *However we will prove the convergence using our sufficiency results.*

3.3 Muller automata and ω-regular languages

A *Muller* automaton $\mathcal{A}$ over an alphabet Σ is a tuple $(Q, q_{init}, \delta, \{F_1, \cdots, F_k\})$ where:

- Q is a finite set of states.
- q_{init} is the initial state.
- $\delta \subseteq Q \times \Sigma \times Q$ is the transition relation (or the set of edges).
- $F_i \subseteq Q$ for $1 \leq i \leq k$ are accepting sets.

An automaton $\mathcal{A}$ defines a language over Σ. Given an infinite sequence sequence $\tau = \tau_1\tau_2\cdots$, we define $\inf(\tau) = \{\tau_i \mid \{j \mid \tau_j = \tau_i\}$ is an infinite set $\}$. A run of $\mathcal{A}$ on $\sigma \in \Sigma^\omega$ is an infinite sequence of states $\rho = q_1q_2\cdots$ such that $q_1 = q_{init}$ and

$(q_i, \sigma(i), q_{i+1}) \in \delta$ for $i \geq 1$. A run ρ of $\mathcal{A}$ on σ is *accepting* if $\inf(\rho) = F_i$ for some i. An infinite word $\sigma \in \Sigma^\omega$ is *accepted* by $\mathcal{A}$ if there exists a run of $\mathcal{A}$ on σ which is accepting. The language accepted by $\mathcal{A}$, denoted $Lang(\mathcal{A})$ is the set of all infinite words accepted by $\mathcal{A}$. A language $L \subseteq \Sigma^\omega$ is *ω-regular* if there exists a Muller automaton whose language is L. We associate a labelled directed graph with $\mathcal{A}$ denote $Graph(\mathcal{A})$ and is defined as $Graph(\mathcal{A}) = (Q, \delta, \Sigma)$. Henceforth when we refer to a path of an automaton or a state being reachable, we refer to the underlying graph.

We call a Muller automaton $\mathcal{A} = (Q, q_{init}, \delta, \{F_1, \cdots, F_k\})$ *simple* if:

- Every state in Q is reachable from q_{init}, and every edge $e \in \delta$ is useful, that is, there exists $\sigma \in Lang(\mathcal{A})$ and an accepting run ρ of $\mathcal{A}$ on σ such that $e = (\rho(i), \sigma(i), \rho(i+1))$ for some i.
- $Graph(\mathcal{A})[F_i]$ is maximally strongly connected in $Graph(\mathcal{A})$ for all i.
- All edges going out of F_i go into F_i for all i, that is, $(q, a, q') \in \delta$ and $q \in F_i$ implies $q' \in F_i$.

The next proposition states that the class of languages accepted by simple Muller automata is exactly the class of ω-regular languages.

Proposition 1 *For every Muller automaton $\mathcal{A}$, there exists a simple Muller automaton $\mathcal{B}$ such that $Lang(\mathcal{A}) = Lang(\mathcal{B})$. Further $\mathcal{B}$ can be constructed in time polynomial in the size of $\mathcal{A}$.*

Proposition 2 *Given a simple Muller automaton $\mathcal{A}$ and a set of edges $E \subseteq \delta$ such that $(Graph(\mathcal{A}) - E)[F_i]$ is not strongly connected for every i, an accepting run ρ of $\mathcal{A}$ on any $\sigma \in \Sigma^\omega$ has infinitely many indices i such that $(\rho(i), \sigma(i), \rho(i+1)) \in E$.*

Example 1 *The language $L_{converge}$ of the Example in section 2 is ω-regular. It is accepted by the following automaton $\mathcal{A}_{converge} = (Q, q_{init}, \delta, \mathcal{F})$. Q consists of two types of states: the first set $Q_1 = 2^{[N]}$ stores the set of alive nodes, the second set $Q_2 = \{(S, E_S, e) \mid S \subseteq [N], (S, E_S)$ is a connected subgraph of G, $e \in E_S\}$. $q_{init} = [N]$. $\mathcal{F} = \{F_{S,E_S} \mid (S, E_S, e) \in Q_2\}$, where $F_{S,E_S} = \{(S, E_S, e) \in Q\}$. All nodes in F_{S,E_S} ensure that eventually the set of alive nodes will be S and those which will communicate infinitely often will be those in E_S. δ consists of three sets of transition:*

- $\delta_1 = \{(S, T, S \cup \{j\}) \mid S \in Q_1, T \in join\} \cup \{(S, T, S - \{j\}) \mid S \in Q_1, T \in fail\} \cup \{(S, T, S) \mid S \in Q_1, T \in move\}$.
- $\delta_2 = \{((S, E_S, e), T, (S, E_S, e')) \mid e = (i,j), T = move_{i,j}, e' \in E_S - \{e\}\}$.
- $\delta_3 = \{(S, T, (S, E_S, e)) \mid T \in move\}$.

3.4 $\mathcal{A}, \mathcal{X}$-neighborhood system

For the rest of the paper, let us fix some notation. Let $\mathcal{X}$ be a set and $\Sigma = \{T_1, \cdots, T_k\}$ a set of operators on $\mathcal{X}$. Let X_{fp} be a non-empty set of common fixpoints of $\mathcal{X}$ with respect to the operators in Σ, and $\mathcal{U}$ a neighborhood system around X_{fp}. Let $\mathcal{A}$ be a Muller automaton on Σ. Let $\mathcal{Y} = Q \times \mathcal{X}$.

We will define some concepts related to $\mathcal{Y}$. Given $Y \subseteq \mathcal{Y}$ and $q \in Q$, $Proj_q(Y) = \{x \mid (q,x) \in Y\}$ and $Proj(Y) = \cup_{q \in Q} Proj_q(Y)$. Given an edge $e = (q, T_i, q') \in \delta$, $func_e((q,x)) = (q', T_i(x))$. Given $Y \subseteq \mathcal{Y}$, $func_e(Y) = \{q'\} \times T_i(Proj_q(Y))$. A set $Y \subseteq \mathcal{Y}$ is said to be $\mathcal{A}$-invariant if for all $e \in \delta$, $func_e(Y) \subseteq Y$. We say that a state $y \in \mathcal{Y}$ is *reachable* if there exists an $x \in \mathcal{X}$, a $\sigma \in Lang(\mathcal{A})$ and a run ρ of $\mathcal{A}$ on σ such that $y = (\rho(i), \sigma(x, i-1))$ for some i. We say that y is reached from x using $w = Pref(\sigma, i-1)$. We will denote the set of all reachable states of Y by $Reachable(Y)$.

Let $Y_{fp} = Reachable(Q \times X_{fp})$. Note that Y_{fp} is an $\mathcal{A}$-invariant set. A $\mathcal{A}, \mathcal{X}$-neighborhood system around X_{fp} is a $Q \times \mathcal{X}$-neighborhood system $\mathcal{W}$ around Y_{fp}. When $\mathcal{X}$ is clear from the context we will drop the $\mathcal{X}$ and call it an $\mathcal{A}$-neighborhood system. $\mathcal{W}$ is said to be *finer* than $\mathcal{U}$, if for every $U \in \mathcal{U}$, there exists a $W \in \mathcal{W}$ such that $Proj(W) \subseteq U$. $\mathcal{U}$ is said to be *finer* than an $\mathcal{W}$, if for every $W \in \mathcal{W}$, there exists a $U \in \mathcal{U}$ such that W contains $Reachable(Q \times U)$. $\mathcal{U}$ and $\mathcal{W}$ are said to be *equivalent*, if $\mathcal{U}$ is finer than $\mathcal{W}$ and $\mathcal{W}$ is finer than $\mathcal{U}$. An $\mathcal{A}, \mathcal{X}$-neighborhood system $\mathcal{W}$ is said to be $\mathcal{A}$-*invariant* if every $W \in \mathcal{W}$ is $\mathcal{A}$-invariant.

Proposition 3 *Let $\mathcal{W}$ be an $\mathcal{A}$-neighborhood system equivalent to $\mathcal{U}$ such that every $W \in \mathcal{W}$ is a subset of $Reachable(Q \times \mathcal{X})$. Then $\mathcal{U}$ has a countable base if and only if $\mathcal{W}$ has a countable base.*

Proof Let $\{U_n\}_{n=1}^{\infty}$ be a countable base for $\mathcal{U}$. Let $W \in \mathcal{W}$. Then there exists $U \in \mathcal{U}$ such that $Reachable(Q \times U) \subseteq W$ since $\mathcal{U}$ is finer than $\mathcal{W}$. There exists $U_n \subseteq U$ by the definition of countable base. Also there exists W_n such that $W_n \subseteq Q \times U_n$ since $\mathcal{W}$ is finer than $\mathcal{U}$. Since $W_n \subseteq Reachable(Q \times \mathcal{X})$ by assumption, we have $W_n \subseteq Reachable(Q \times U_n)$. Therefore $W_n \subseteq W$. $\{W_n\}_{n=1}^{\infty}$ is a countable base for $\mathcal{W}$.

Similarly let $\{W_n\}_{n=1}^{\infty}$ be a countable base for $\mathcal{W}$. Let $U \in \mathcal{U}$. Then there exists $W \in \mathcal{W}$ such that $Proj(W) \subseteq U$, or there exists W_n such that $Proj(W_n) \subseteq U$. Also there exists $U_n \in \mathcal{U}$ such that $Reachable(Q \times U_n) \subseteq W_n$. Hence $Proj(Reachable(Q \times U_n)) \subseteq U$, or $U_n \subseteq U$. $\{U_n\}_{n=1}^{\infty}$ is a countable base for $\mathcal{U}$. □

4 Stability

From now on we will assume that $\mathcal{A}$ is a simple Muller automaton.

The following result generalizes the result of Tsitsiklis [8].

Theorem 1. *The following are equivalent.*

1. *$Lang(\mathcal{A})$ is stable with respect to an $\mathcal{X}$-neighborhood system $\mathcal{U}$ around X_{fp}.*
2. *There exists an $\mathcal{A}, \mathcal{X}$-invariant neighborhood system $\mathcal{W}$ around $Reachable(Q \times X_{fp})$ which is equivalent to $\mathcal{U}$.*

Proof The proof is along the lines of that given in [8]. Let $Lang(\mathcal{A}) = L$.
$(1 \Rightarrow 2)$: We assume that L is stable with respect to $\mathcal{U}$ and we need to construct an $\mathcal{A}$-invariant neighborhood system $\mathcal{W}$ of subsets of $\mathcal{Y}$ which is equivalent to $\mathcal{U}$.

We do this as follows. Given $U \in \mathcal{U}$, we define W_U as the union of all $\mathcal{A}$-invariant subsets of $U' = Reachable(Q \times U)$. Note that $Y_{fp} = Reachable(Q \times X_{fp})$ is an $\mathcal{A}$-invariant subset of U', which shows that W_U is nonempty for all $U \in \mathcal{U}$. Also W_U is the largest $\mathcal{A}$-invariant subset of U'. Let $\mathcal{W}' = \{W_U : U \in \mathcal{U}\}$ and let $\mathcal{W}$ be the closure of $\mathcal{W}'$ under finite intersections and unions.

$\mathcal{W}$ is a neighborhood system. The first property is satisfied since $Q \times X_{fp} \subseteq W_U$ for all U implies $Q \times X_{fp} \subseteq W$ for all $W \in \mathcal{W}$. The sets in $\mathcal{W}$ are closed under finite intersections and arbitrary union by definition. Let $y = (q, x) \in \mathcal{Y} - Y_{fp}$. If y is not reachable then it does not belong to any W_U. Otherwise $x \notin X_{fp}$, and therefore $x \notin U$ for some $U \in \mathcal{U}$. Therefore $y \notin W_U$. Therefore for every $y \in \mathcal{Y} - Y_{fp}$, there exists $W \in \mathcal{W}$ such that $y \notin W$.

$\mathcal{W}$ is $\mathcal{A}$-invariant, since W_U are $\mathcal{A}$-invariant and finite intersections and arbitrary unions of $\mathcal{A}$-invariant sets are $\mathcal{A}$-invariant.

$\mathcal{W}$ is equivalent to $\mathcal{U}$. For every $U \in \mathcal{U}$, there exists $W_U \in \mathcal{W}$ such that $W_U \subseteq Q \times U$, hence $\mathcal{W}$ is finer than $\mathcal{U}$. To show that $\mathcal{U}$ is finer than $\mathcal{W}$, it is enough to show that $\mathcal{U}$ is finer than $\mathcal{W}'$, because the sets in $\mathcal{U}$ are also closed under finite intersections and arbitrary unions. Let $W \in \mathcal{W}'$, then $W_U = W$ for some $U \in \mathcal{U}$. Using the fact that L is *stable* $\exists V \in \mathcal{U}$ such that $\forall x \in \mathcal{X}$, $\forall \sigma \in L$, if there exists $n_0 \in \mathbb{N}$ such that $\sigma(x, n_0) \in V$ then $\forall n \geq n_0$, $\sigma(x, n) \in U$. Define $V' = \{y \mid y \text{ is reached from } x \in \mathcal{X} \text{ using } Pref(\sigma, j) \text{ for some } \sigma \in L \text{ and } \sigma(x, i) \in V$ for some $i \leq j\}$. In particular, V' contains $Reachable(Q \times V)$. Note that V' is an $\mathcal{A}$-invariant subset of U' (here we use that fact that every edge is useful). Hence $V' \subseteq W_U$. Therefore there exists $V \in \mathcal{U}$ such that $Reachable(Q \times V) \subseteq W$.
$(2 \Rightarrow 1)$: Given any $U \in \mathcal{U}$, there exists some $W \in \mathcal{W}$ such that $Proj(W) \subseteq U$, because $\mathcal{W}$ is finer than $\mathcal{U}$. Moreover, since $\mathcal{U}$ is finer than $\mathcal{W}$, there exists some $V \in \mathcal{U}$ such that $V' = Reachable(Q \times V)$ is contained in W. Consider an x, $n_0 \in \mathbb{N}$ and $\sigma \in L$ such that $\sigma(x, n_0) \in V$. Let ρ be an accepting run of $\mathcal{A}$ on σ. Then $(\rho(n_0 + 1), \sigma(x, n_0)) \in V' \subseteq W$. Since W is $\mathcal{A}$-invariant $(\rho(n + 1), \sigma(x, n)) \in V'$ for all $n \geq n_0$. Therefore $\sigma(x, n) \in Proj(W) \subseteq U$ for all $n \geq n_0$. Hence L is stable with respect to $\mathcal{U}$. □

Remark 3 *We note that the $\mathcal{W}$ constructed in the first part of the proof is such that every $W \in \mathcal{W}$ is a subset of $Reachable(Q \times \mathcal{X})$.*

5 Convergence

In this section we present necessary and sufficient conditions for convergence of a ω-regular language L. Let $\mathcal{X}$, X_{fp}, Σ, $\mathcal{U}$, $\mathcal{Y}$ and Y_{fp} be as above. Let L be an ω-regular language over Σ such that $L = Lang(\mathcal{A})$, where $\mathcal{A}$ is a simple Muller automaton.

First, we present a sufficient condition for convergence.

Condition 1 *There exists a totally ordered index set I and a collection $\{X_\alpha : \alpha \in I\}$ of distinct subsets of $\mathcal{Y}$ containing Y_{fp} with the following properties:*

Property 1 *$\alpha < \beta$ implies $X_\alpha \subseteq X_\beta$.*

Property 2 *For every $U \in \mathcal{U}$, there exists some $\alpha \in I$ such that $Proj(X_\alpha) \subseteq U$.*

Property 3 *$\bigcup_{\alpha \in I} Proj_{q_{init}}(X_\alpha) = \mathcal{X}$.*

Property 4 *X_α is $\mathcal{A}$-invariant for all $\alpha \in I$.*

Property 5 *For every $\alpha \in I$ such that $X_\alpha \neq Y_{fp}$, there exists $E \subseteq \delta$ such that for every i $(Graph(\mathcal{A}) - E)[F_i]$ is not strongly connected, and for every $e \in E$, $func_e(X_\alpha) \subseteq \bigcup_{\beta < \alpha} X_\beta$.*

Property 6 *Every non-empty subset of I which is bounded below has a smallest element.*

Following theorem states that the above condition is sufficient for convergence.

Theorem 2. *If Condition 1 holds, then L converges to X_{fp} with respect to $\mathcal{U}$.*

Proof Let I, $\{X_\alpha : \alpha \in I\}$ have the properties in Condition 1. Suppose that we are given some $U \in \mathcal{U}$, $x_0 \in \mathcal{X}$ and $\sigma \in L$. We must show that $\sigma(x_0, n)$ eventually enters and remains in U.

Let us fix an accepting run $\rho = q_1 q_2 q_3 \cdots$ of $\mathcal{A}$ on σ. Let
$J = \{\alpha \in I : \exists n \text{ such that } (\rho(n), \sigma(x_0, n-1)) \in X_\alpha\}$.

Lemma 1. *$J = I$.*

Proof Since from Property 3, we have $\mathcal{X} = \bigcup_{\alpha \in I} Proj_{q_{init}}(X_\alpha)$, there exists some $\alpha \in I$ such that $(q_{init}, x_0) \in X_\alpha$. Hence J is nonempty. We consider two cases: we first assume that J is not bounded below. Then, for every $\alpha \in I$, there exists a $\beta \in J$ such that $\beta < \alpha$. Hence for every $\alpha \in I$, there exists some $\beta < \alpha$ and some integer n such that $(\rho(n), \sigma(x_0, n-1)) \in X_\beta \subseteq X_\alpha$ (Property 1). So, every $\alpha \in I$ belongs to J, and $I = J$.

Let us now assume that J is bounded below. Since it is nonempty, it has a smallest element from Property 6, denoted by β. If $X_\beta = Y_{fp}$, then β is also the smallest element of I, and $I = J$ follows. So, let us assume that $X_\beta \neq Y_{fp}$. Then from Property 5 there exists $E \subseteq \delta$ such that $(Graph(\mathcal{A}) - E)[F_i]$ is not strongly connected, and for every $e \in E$, $func_e(X_\beta) \subseteq \bigcup_{\gamma < \beta} X_\gamma$. From the definition of J, there exists some n_0 such that $(\rho(n_0), \sigma(x_0, n_0 - 1)) \in X_\beta$ and by invariance of X_β (Property 4), $(\rho(n), \sigma(x_0, n-1)) \in X_\beta$ for all $n \geq n_0$. Since $\sigma \in L$ and ρ is an accepting run, we have an $m > n_0$ such that $(\rho(m), \sigma(m), \rho(m+1)) \in E$ by Proposition 2. Since $(\rho(m), \sigma(x_0, m-1)) \in X_\beta$, we have $(\rho(m+1), \sigma(x_0, m)) \in \bigcup_{\gamma < \beta} X_\gamma$, or $(\rho(m+1), \sigma(x_0, m)) \in X_\gamma$ for some $\gamma < \beta$. Hence $\gamma \in J$, which contradicts the assumption that β was the smallest element of J. This completes the proof of the lemma. □

Given $U \in \mathcal{U}$, there exists some $\alpha \in I$ such that $Proj(X_\alpha) \subseteq U$ (Property 2). Since $J = I$, there exists some n_0 such that $(\rho(n_0), \sigma(x_0, n_0 - 1)) \in X_\alpha$. Since $func_e(X_\alpha) \subseteq X_\alpha$ for all e, it follows that $(\rho(n), \sigma(x_0, n-1)) \in X_\alpha$ for all $n \geq n_0$. Or $\sigma(x_0, n-1) \in Proj(X_\alpha)$ for all $n \geq n_0$. Hence $\sigma(x_0, n-1) \in U$ for all $n \geq n_0$, which completes the proof. □

Next we show that Condition 1 is a necessary condition when the system satisfies some additional properties.

Theorem 3. *If L is stable and converges to X_{fp} with respect to $\mathcal{U}$ and if $\mathcal{U}$ has a countable base, then Condition 1 holds.*

Proof Since L is stable with respect to $\mathcal{U}$, we have from Theorem 1 that there exists an $\mathcal{A}$-invariant neighborhood system $\mathcal{W}$ which is equivalent to $\mathcal{U}$. Since $\mathcal{U}$ has a countable base, from Proposition 3 and Remark 3, we have that $\mathcal{W}$ has a countable base as well $\{W_n\}_{n=1}^{\infty}$. Without loss of generality we may assume that $W_{n+1} \subseteq W_n$ for all n. (Otherwise we could define a new countable base $W'_n = \bigcap_{k=0}^{n} W_k$.) Let W_0 be $Reachable(Q \times \mathcal{X})$.

Our proof consists of two main steps: for each $n \geq 0$ we construct a nested collection of subsets of $\mathcal{Y}$ which lie between W_n and W_{n+1}. Then we merge these collections to get a single nested collection.

Lemma 2. *Let $W', W'' \in \mathcal{W}$ such that $W' \subset W''$. Let Inv be the set of all $\mathcal{A}$-invariant subsets of W'' containing W'. Then there exists a function $f : Inv \rightarrow Inv$ and $g : Inv \rightarrow 2^\delta$ such that:*

- *For any $W \in Inv$, we have $W \subseteq f(W)$ and if $Proj_{q_{init}}(W) \neq Proj_{q_{init}}(W'')$ then $f(W) \subset W$.*
- *$func_e(f(W)) \subseteq W$ for all $e \in g(W)$.*
- *$(Graph(\mathcal{A}) - g(W))[F_i]$ is not strongly connected for all i.*

For the sake of continuity we prove continue with the proof of the theorem and prove the lemma later.

Let I_n be a well-ordered set with cardinality larger than that of $\mathcal{Y}$ and let $\alpha_{0,n}$ be its smallest element. We apply Lemma 2 with $W'' = W_n$ and $W' = W_{n+1}$ to obtain a function f_n satisfying the properties of the lemma above. We define a function $h_n : I_n \rightarrow Inv$ using the following transfinite recursion: $h_n(\alpha_{0,n}) = W_{n+1}$, and for all $\alpha > \alpha_{0,n}$, $h_n(\alpha) = f_n(\bigcup_{\beta<\alpha} h_n(\beta))$.

Notice that $W_{n+1} \subseteq h_n(\beta) \subseteq h_n(\alpha) \subseteq W_n$, for any α, β such that $\alpha > \beta$, and that if $Proj_{q_{init}}(h_n(\beta)) \neq Proj_{q_{init}}(W_n)$, then $h_n(\beta) \subset h_n(\alpha)$. Since I_n has cardinality larger than that of $\mathcal{Y}$, there exists some $\alpha \in I_n$ such that $Proj_{q_{init}}(h_n(\alpha)) = Proj_{q_{init}}(W_n)$. Let $\bar{\alpha}_n$ be the smallest such α and let $\bar{I}_n = \{\alpha \in I_n \,|\, \alpha < \bar{\alpha}_n\}$.

We now define $I = \{\top\} \cup \{(n, \alpha) \,|\, \alpha \in \bar{I}_n, n = 0, 1, \cdots\}$ with the following total order: $(n, \alpha) < (m, \beta)$ if either $n > m$ or $n = m$ and $\alpha < \beta$, and $(n, \alpha) < \top$ for all n and α. Finally, let $X_{(n,\alpha)} = h_n(\alpha)$ and $X_\top = W_0$.

We claim that the collection $\{X_\alpha \,|\, \alpha \in I\}$ satisfies all the properties of Condition 1. Property 1 is satisfied because $h_n(\beta) \subset h_n(\alpha)$ for every $\beta < \alpha$, where $\beta, \alpha \in \bar{I}_n$. Property 2 follows from the fact that our X_αs include the countable base $\{W_n\}_{n=1}^{\infty}$ we started with and $\mathcal{W}$ is equivalent to $\mathcal{U}$. Since $Proj_{q_{init}}(W_0) = \mathcal{X}$, Property 3 is true. Property 4 holds since all the new sets we introduce (basically in Lemma 2) are $\mathcal{A}$-invariant. Again Property 5 follows from Lemma 2, where the function g gives the set of edges E for every invariant set. Finally, since every non-empty subset of a well-ordered set has a least element, and countable concatenations of well-ordered sets is well ordered, we satisfy 6. □

5.1 Proof of Lemma 2

Given $Y \subseteq \mathcal{Y}$ and $e = (q, T, q')$, define $Reach_e(Y) = \{(q', x') \mid \exists (q, x) \in Y, x' = T(x)\}$. Given a path P in $Graph(\mathcal{A})$, $Reach_P(S)$ is defined inductively as follows. If $P = e \in \delta$, then $Reach_P(Y) = Reach_e(Y)$. Otherwise if $P = P'e$, and $Reach_P(Y) = Reach_e(Reach_{P'}(Y))$. Given an edge $e \in \delta$, $PathsEnd(e) = \{P \mid P = P'e\}$, is the set of all paths ending in e.

Given a set of edges $E \subseteq \delta$ and $W \in \mathcal{W}$, define $f_E(W) = \{(q, x) \mid \forall e \in E, P \in PathsEnd(e), Reach_P(\{(q, x)\}) \subseteq W\}$. $f_E(W)$ has the following properties.

- $W \subseteq f_E(W)$: Since W is invariant, for all P $Reach_P(W) \subseteq W$.
- $f_E(W)$ is $\mathcal{A}$-invariant: Let $x \in f_E(W)$, and $Y = f_E(\{x\})$. If there exists $y \in Y$ such that $y \notin f_E(W)$, then there exists $e \in E$ and $P \in PathsEnd(e)$ such that $Reach_P(\{y\}) \not\subseteq W$. Then $x \notin f_E(W)$, since there exists $e \in E$ and a path $e'P \in PathsEnd(e)$ such that $Reach_P(\{x\}) \not\subseteq W$.
- $func_e(f_E(W)) \subseteq W$ for all $e \in E$: The argument is similar to the previous.

Let $\mathcal{E} = \{E \subseteq (\delta \cap \cup_i (F_i \times \Sigma \times F_i)) \mid (Graph(\mathcal{A}) - E)[F_i]$ is not strongly connected for every $i\}$. Given $W \in Inv$, we claim that if $Proj_{q_{init}}(W) \neq Proj_{q_{init}}(W'')$, then $W \subset f_E(W)$ for some $E \in \mathcal{E}$. Suppose not. Then there exists $(q_{init}, x_0) \in W'' - W$ and $f_E(W) = W$ for all $E \in \mathcal{E}$. Therefore (q_{init}, x_0) does not belong to any $f_E(W)$.

We will construct a $\sigma \in Lang(\mathcal{A})$ and an accepting run ρ of $\mathcal{A}$ on σ such that $(\rho(i), \sigma(x_0, i-1)) \notin f_E(W)$ for all $E \in \mathcal{E}$ and $i \geq 1$. This contradicts the convergence of L as follows. There exists $U \in \mathcal{U}$ such that $Z = Reach(Q \times U) \subseteq W'$, since $\mathcal{U}$ is finer than $\mathcal{W}$. Since $Z \subseteq W' \subseteq W = f_E(W)$, $(\rho(i), \sigma(x_0, i-1)) \notin Z$ for all i. Therefore $\sigma(x_0, i) \notin U$ for all i, contradicting the convergence of L with respect to $\mathcal{U}$. Hence $W \subset f_E(W)$ for some $E \in \mathcal{E}$. We set $f(W) = f_E(W)$ for some E for which $W \subset f_E(W)$.

It remains to construct a $\sigma \in L$ and ρ which satisfy the above condition. Given a path P starting in q and a singleton set $\{(q, x)\}$, $Reach_P(\{(q, x)\})$ is a singleton. Hence we will write this as just $Reach_P((q, x))$.

Let $E \in \mathcal{E}$ be non-empty. Let $s_0 = (q_{init}, x_0)$. Since $s_0 \notin f_E(W)$ there exists some P ending in an edge in E such that $Reach_P(s_0) \notin f_E(W) = W$. Let us call this P as P_1 and $Reach_P((q_{init}, x_0))$ as s_1. Let the last of edge of P belong to $F_{i^*} \times \Sigma \times F_{i^*}$. Since the automaton is simple any path starting from F_{i^*} will remain within F_{i^*}. We will assume $|F_{i^*}| \geq 2$, (a similar procedure can be used when $|F_{i^*}| = 1$.

The following procedure generates a sequence of P_js:

1. Let P_1 and s_1 be as defined above. Initialize j to 1.
2. Let $Q' = \emptyset$.
3. While $Q' \neq Q$ do:
   - Add the last state of P_j to Q'.
   - Consider $E = \delta \cap Q' \times \Sigma \times (Q - Q')$.
   - Increment j.
   - Set P_j to be a path ending in E such that $Reach_{P_j}(s_{j-1}) \notin f_E(W) = W$.
   - Set $s_j = Reach_{P_j}(s_{j-1})$.

Note that we maintain the invariant that $s_{j-1} \notin f_E(W)$ for some E and hence $s_{j-1} \notin W$. Therefore there always exists a path P_j ending in E such that $Reach_{P_j}(s_{j-1}) \notin f_E(W) = W$. Let $P' = P_1 P_2 \cdots$ be the sequence of edges $e_1 e_2 \cdots$ with $e_i = (q_i, a_i, q_{i+1})$. Define $\sigma = a_1 a_2 \cdots$ and $\rho = q_1 q_2 \cdots$. ρ is a run of $\mathcal{A}$ on σ. It is accepting because each P_j contains every state from F_{i^*} at least once, and only contains states from F_{i^*}, because the automaton is simple. We have infinitely many i such that $(\rho(i), \sigma(x_0, i-1)) \notin f_{\mathcal{E}}(W)$ for any E or equivalently $(\rho(i), \sigma(x_0, i-1)) \notin W$ (They correspond to s_js). Since W is invariant, if $(\rho(i), \sigma(x_0, i-1)) \in W$ for some i, then $(\rho(j), \sigma(x_0, j-1)) \in W$ for all $j \geq i$, which contradicts the previous statement. Therefore $(\rho(i), \sigma(x_0, i-1)) \notin W$ for all i.

6 An Application

In this section, we illustrate the application of our results to prove convergence of the Example in Section 2.

We have already defined $\mathcal{X}, X_{fp}, \mathcal{U}$ and $L_{converge}$. We will prove convergence using the simple Muller Automaton $\mathcal{A}_{converge}$. We will then point out how one can prove convergence given any simple Muller automaton for $L_{converge}$.

6.1 Properties of the neighborhood system $\mathcal{U}$

Let us define $move_{i,j}(x) = x'$ as follows: If $x_i \neq \perp, x_j \neq \perp$, then $x'_i = x'_j = (x_i + x_j)/2$ and $x'_k = x_k$ for $k \notin \{i, j\}$, otherwise $x' = x$.

We recall the following two results from [5].

Proposition 4 *$f(move_{i,j}(x)) \leq f(x)$.*

Let $Sorted(x)$ from $[N] \to [N]$ be a one-one and onto function which satisfies for $i < j$, $x_{Sorted(x)(i)} < x_{Sorted(x)(j)}$, or $x_{Sorted(x)(i)} \leq x_{Sorted(x)(j)}$ and $Sorted(x)(i) < Sorted(x)(j)$. $Sorted(x)(i)$ will give the identifier of the agent with the i-th smallest value and when agents have same values, the value of the agent with the smaller identifier is considered smaller. Here $\perp$ is considered to have a value of ∞.

Proposition 5 *Given any x, there exists $k \in [N]$ such that $x_{Sorted(x)(k+1)} - x_{Sorted(x)(k)} > \frac{1}{alive(x)} \sqrt{\frac{f(x)}{alive(x)}}$. Given any i, j such that $1 \leq i \leq k$ and $k+1 \leq j \leq alive(x)$, $f(move_{i',j'}(x)) \leq \beta(alive(x)) f(x)$, where $i' = Sorted(x)(i)$ and $j' = Sorted(x)(j)$.*

The above property says that if we start at some (H, x) in U_i, then there is partition of the nodes in H, such that for all edges (i, j) which go between the partitions, $move_{i,j}(H, x)$ will be in U_{i+1}. But we need more, we need one such partition which will work for all elements of U_i.

Define $Cut(x, k) = (A, B)$ where $A = \{Sorted(x)(i) \mid i \leq k\}$ and $B = \{Sorted(x)(j) \mid alive(x) \geq j \geq k+1\}$. Define $Gap(x) = \{Cut(x, k) \mid 1 \leq k < alive(x), x_{Sorted(x)(k+1)} - x_{Sorted(x)(k)} > \frac{1}{alive(x)}\sqrt{\frac{f(x)}{alive(x)}}\}$.

Proposition 6 *For all $i, j \in [N]$ and x such that for all $(A, B) \in Gap(x)$, either $i, j \leq A$ or $i, j \geq B$, we have $Gap(x) \subseteq Gap(move_{i,j}(x))$.*

Let $\{Gap(x) \mid (H, x) \in U_i - U_{i+1}\} = \{C_1, \cdots, C_{i_n}\}$ such that if $C_i \subset C_j$ then $i < j$. Between U_i and U_{i+1} we define a finite number of sets as follows. $U_{i,0} = U_i$, $U_{i,j+1} = U_{i,j} - \{(H, x) \in \mathcal{X} \mid Gap(x) = C_{j+1}\}$ for $0 \leq j \leq i_n - 1$. We define the index set to be $J = \{(i, j) \mid 0 \leq j \leq i_n - 2\}$ with the ordering $(i, j) < (i', j')$ if $i > i'$ or $i = i'$ and $j > j'$. The required sets are $\{U_\alpha \mid \alpha \in J\}$. This set has the following property.

Proposition 7 *For all $\alpha \in J$, there exists $C \subseteq [N] \times [N]$ such that for all $(i, j) \in C$ and $(H, x) \in U_\alpha$, C is a cut in H and $move_{i,j}(H, x) \in U_\beta$ for some $\beta < \alpha$. Also, for all $\alpha \in J$, $i, j \in [N]$, $(H, x) \in U_\alpha$, we have $move_{i,j}((H, x)) \in U_\alpha$.*

Proof Given α, $U_\alpha = U_{m,n}$ for some m, n. Let $Gap(x) = Z$ which is the same for any $(H, x) \in U_\alpha - U_{\alpha+1}$ where $\alpha + 1 = m, n+1$ if $(m, n+1) \in J$, otherwise $\alpha + 1 = m+1, n$. The required cut $C = \{(i, j) \mid \exists (A, B) \in Z, i \in A, j \in B\}$. Then from Proposition 5, we have for all $(i, j) \in C$, $move_{i,j}((H, x)) \in U_{m+1,n}$ for all $(H, x) \in U_\alpha - U_{\alpha+1}$.

Given $(H, x) \in U_\alpha$, if $(i, j) \in C$, $move_{i,j}(H, x) \in U_\alpha$ from above. If $(i, j) \notin C$, then from Proposition 4 we have $move_{i,j}(H, x) \in U_{m',n'}$ where $m' \geq m$ and from 6, we have $n' \geq n$. Therefore $move_{i,j}((H, x)) \in U'_\alpha$ for some $\alpha' \leq \alpha$, hence also in U_α. □

6.2 Convergence proof

We are now ready to define the invariant sets required to prove convergence.

Recall $\mathcal{A}_{converge} = (Q, q_{init}, \delta, \mathcal{F})$ with $Q = Q_1 \cup Q_2$ and $\delta = \delta_1 \cup \delta_2 \cup \delta_3$ and $\mathcal{F} = \{func_{S,E_S} \mid (S, E_S, e) \in Q_2\}$. Let $\mathcal{Y} = Q \times \mathcal{X}$, $Y_{fp} = Reachable(Q \times X_{fp})$. The index set $I = J \cup \{\top\}$, with $\top > j$ for all $j \in J$. For $\alpha \in J$, $Y_\alpha = (Q_2 \times X_\alpha) \cup Y_{fp}$, and $Y_\top = \mathcal{Y}$. The index set I with the sets $\{Y_\alpha \mid \alpha \in I\}$ satisfy all the properties of Condition 1. It is easy to see that Properties 1, 2, 3 and 6 are satisfied. $Y_\top$ is clearly invariant. For $\alpha \in J$ and any edge e not in δ_2, $func_e(Y_\alpha) = \emptyset \subseteq Y_\alpha$. For $\alpha \in J$ and $e = (q, a, q') \in \delta_2$, $a = move_{i,j}$ for some i, j and $func_e(Y_\alpha) \subseteq (\{q'\} \times move_{i,j}(X_\alpha)) \cup (Q \times X_{fp})$. Since $move_{i,j}(X_\alpha) \subseteq X_\alpha$ from Proposition 7, we have that Y_α is invariant. Now we show that Property 5 also holds. For $Y_\top$, we can choose E to be δ_2. $(Graph(\mathcal{A}) - E)[F_i]$ is not strongly connected for every i. We need to show that for all $(q,(H,x)) \in Y_\top$, for all $e \in E$, $func_e((q,(H,x))) \in Y_\alpha$ for some $\alpha \in J$. We need to consider only $(q,(H,x)) \in Y_\top - \bigcup_{\alpha \in J} Y_\alpha$. But then $q \in Q_1$ and hence $func_e((q,(H,x))) = \emptyset$. For any other Y_α, we can choose $E = \{(q, move_{i,j}, q') \mid (i,j) \text{ or } (j,i) \in C\}$, where C is the cut associated with X_α given by Proposition 7. It is easy to see that $(Graph(\mathcal{A}) - E)[F_i]$ is not strongly connected since the labels of the edges in each F_i correspond to a connected subgraph on the unfailed nodes, and C is cut in the induced subgraph of G with unfailed nodes.

Since the system is not stable as mentioned before we cannot use Theorem 2 to guarantee existence of level sets to prove convergence of the system for any arbitrary automaton accepting $L_{converge}$. However we can always find such sets because of the following structure of any simple Muller automaton $\mathcal{A} = (Q, q_{init}, \delta, \Sigma, \{F_1, \cdots, F_k\})$ such that $Lang(\mathcal{A}) = L_{converge}$.

- There is no edge labelled by *join* or *fail* in any of the F_i, that is, there is no $a \in join \cup fail$ and $q, q' \in F_i$ for some i, such that $(q, a, q') \in \delta$. Because then we would have an accepting run with infinite *join* or *fail* operations.
- Let C be a cut of $G = (V, E)$, i.e, $G - C$ is not connected. Then removing the edges in $\mathcal{A}$ labelled by elements in C renders every F_i not strongly connected, i.e., for every i $(Graph(\mathcal{A}) - E')[F_i]$ is not strongly connected, where $E' = \{(q, a, q') \mid a = move_{i,j}, (i,j) \in C, q, q' \in \bigcup_j F_j\}$.

References

1. V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. Convergence in multiagent coordination consensus and flocking. In *Proceedings of the Joint forty-fourth IEEE Conference on Decision and Control and European Control Conference*, pages 2996–3000, 2005.
2. V. Borkar and P. Varaiya. Asymptotic Agreement in Distributed Estimation. *IEEE Trans. on Automatic Control*, 27(3):650–655, June 1982.
3. K. Mani Chandy, Sayan Mitra, and Concetta Pilotto. Convergence verification: From shared memory to partially synchronous systems. In *In proceedings of Formal Modeling and Analysis of Timed Systems (FORMATS'08)*, volume 5215 of *LNCS*, pages 217–231. Springer Verlag, 2008.
4. A. Jadbabaie, J. Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.
5. Sayan Mitra and K. Mani Chandy. A formalized theory for verifying stability and convergence of automata in PVS. In Otmane Aït Mohamed, César Muñoz, and Sofiène Tahar, editors, *TPHOLs*, volume 5170 of *Lecture Notes in Computer Science*, pages 230–245. Springer, 2008.
6. R. Olfati-Saber. Flocking for multi-agent dynamic systems: algorithms and theory. *Automatic Control, IEEE Transactions on*, 51(3):401–420, March 2006.
7. Reza Olfati-saber, J. Alex Fax, and Richard M. Murray. Consensus and cooperation in networked multi-agent systems. In *Proceedings of the IEEE*, page 2007, 2007.
8. John N. Tsitsiklis. On the stability of asynchronous iterative processes. *Mathematical Systems Theory*, 20(2-3):137–153, 1987.