Université catholique de Louvain
Ecole Polytechnique de Louvain
Département d'Ingénierie Mathématique

Center for Systems Engineering and Applied Mechanics

Graphs and Networks for the Analysis of Autonomous Agent Systems

Julien M. Hendrickx

Thesis submitted in partial fulfillment
of the requirements for the degree of
Docteur en sciences appliquées

Dissertation committee:
Prof. Brian D.O. Anderson, ANU (Australia)
Prof. Vincent D. Blondel, Advisor, UCL
Prof. Karl Henrik Johansson, KTH (Sweden)
Prof. Rodolphe Sepulchre, ULg
Prof. Paul Van Dooren, Chair, UCL
Prof. Sandro Zampieri, University of Padova (Italy)

February 2008

Acknowledgments

First of all, I would like to thank my advisor Vincent Blondel for his invaluable support and help, for providing me so much useful advice, many ideas and opportunities, and giving me a great liberty in my research at the same time. I am also very grateful for his careful reading of the thesis, especially when time was running short for me.

I would also like to thank all the members of my thesis committee for their questions and feedback, which have helped me greatly in improving this thesis.

My gratitude goes to Brian Anderson and his team, and to John Tsitsiklis, with whom I had the pleasure to collaborate during this PhD, and thanks to whom I could spend some months at A.N.U. and at M.I.T.

Visiting other labs and teams is always an enjoyable and interesting experience from which much can be learned. I am therefore also thankful to Rodolphe Sepulchre (ULg), Stephen Morse (Yale), Magnus Egerstedt (Georgia Tech), Pierre-Alexandre Bliman (INRIA) and Karl Henrik Johansson (KTH) who did me the honor of inviting me in their universities or research centers.

Working at Cesame has been a pleasure, and I would like to thank all the people who made these years enjoyable. In particular, I am grateful to Christel, Sergey, Emmanuel and Tzvetan with whom I have shared my office, and to Brieux and Thomas with whom I have often shared the building at some improbable times. Michèle, Isabelle and Lydia of the secretary staff also deserve a special attention for their great work.

Writing a thesis takes several iterations. I would like to thank Raphaël Jungers for his careful reading of the first draft, and for his many suggestions. I am also grateful to him for the different collaborations and the many interesting discussions we had. I also really appreciated the detailed remarks and suggestions of Paul Van Dooren and Sandro Zampieri concerning both the text and the mathematical notations and proofs. The advices of Laure Ninove and Jean-Charles Delvenne during the writing and the whole PhD were very useful too. Besides, Brainstorming sessions with Jean-Charles were very rewarding experiences.

It would have been very difficult to accomplish this work without the support of my parents, family and friends, who I would like to thank all.

Last but certainly not least, I am extremely grateful to Sylvie for her constant understanding, help and support throughout these last years, and for the many hours spent rereading this thesis.

Pecunia nervus belli: This thesis has been financially supported by a F.R.S.-FNRS (Fund for Scientific Research) fellowship. On various occasions, the support of the Concerted Research Action "Large Graphs and Networks" of the Communauté française de Belgique, and of the Belgian Programme on Interuniversity Attraction Poles (initiated by the Belgian Federal Science Policy Office) has been appreciated.

Contents

Chapter 1

General Introduction

A natural trend when designing a system is either to make it monolithic or hierarchic. A monolithic system is one designed as a whole to accomplish a certain task. A hierarchic system is one in which blocks are designed to accomplish particular tasks, and those blocks are articulated and controlled according to a certain hierarchy. Many of the complex devices that we use are hierarchical. Think for example of a computer where the controller leads the CPU, computing unit of the processor, and where the CPU itself controls all other devices such has the hard-disk or the screen. Large computer programs are also designed hierarchically, with large functions implemented by many smaller ones. Similarly, the main decisions about an airplane in flight are made by the pilot that transmits them to the appropriate actuators via piloting devices. This thesis is also written in a hierarchical way, with a main file containing different parts that contains chapters etc. Finally many industrial plants also work in a hierarchical (and often centralized) way with a control room where all measures outputs are shown and where levers and buttons (or computers) can be used to control the plant.

Nature provides however many examples of sophisticated group behaviors that would never be achievable by individual members of the group, and that are produced in the absence of clear or fixed hierarchy. Fish and birds move sometimes as a group with such an efficient and smooth coordination that they can give the impression of being one single large animal (see Figure 1.1). There is no obvious leader or commanding chain in such groups, and the member interactions depend on the opportunity of the moment, and not according to some pre-specified order. Group of ants are also known to accomplish impressive tasks without clear organization. In addition to what they are able to build and move, they can also find as a group acceptable solutions to shortest paths problems, something that no single ant could do. As an other example, think of the kind of problems a group of persons can tackle as opposed to what a single person can do.

Such systems have a large robustness with respect to the addition or removal of agents. A school of fish remains approximately the same when the number of fish is increased or decreased by 25%, but removing (or adding) one single piece of a computer may lead to dramatic consequences. Whoever had to fight with insects also knows that removing 75% of them may decrease but certainly not suppress the nuisance.

Finally, the complementarity between different agent may also improve the efficiency. Characteristics offering an advantage for one task may be a drawback for other tasks. The collaboration between different agents having different characteristics may thus lead to the accomplishment of different tasks that no single agent would be able to do. This specialization may however lead to a decrease in robustness if some sort of agent has too few representatives.

The use of multi-agent systems and of its increased robustness properties comes however at a cost: that of complexity. Describing and predicting the behavior of a system of equal entities interacting with each other without clear hierarchy may be extremely difficult, even if the systems observed in nature are sometimes simple. In mechanics for example, describing the behavior of a single point mass body is usually a simple task, but describing the evolution of a system with three or more bodies interacting with each other may be delicate, as it can present chaotic behavior. This inherent complexity might actually be the reason why hierarchical structures are often *designed*, while non-hierarchical structures are often rather *observed*. The former allows indeed focussing on one single task almost without taking the rest of the system into account. Note that in parallel with the development of the multi-agent system design ideas, there was recently an increasing interest in the analysis of “natural” systems that can be observed, probably in the hope of learning how to cope with their inherent complexity.

Finally, providing a clear definition of multi-agent systems is not an easy task, as any multi-agent system can be abstracted as a monolithic one using appropriate state variables, and many monolithic systems can be decomposed into subsystems interacting with each other. Besides, multi-agent systems can sometimes present a relative hierarchy during their evolution. Let us however take the risk of proposing a non-mathematical definition. In our opinion, a system is multi-agent or totally decentralized if it can be decomposed into identical entities called agents, that interact with each other, and whose number can be modified without affecting the nature of the system. More generally, there might be different classes of agent having different properties, but the number of agent in each class should be modifiable without affecting the nature of the system. As simple example of application of this definition, a car is not a multi-

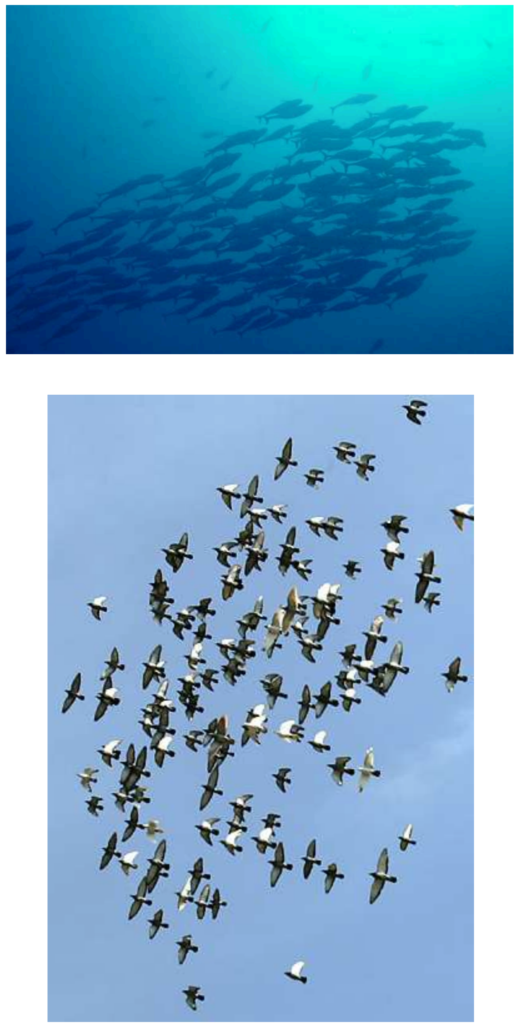

Figure 1.1: Example of school of fish and flock of birds spontaneously agreeing on a common directions. The second image is used under the CC-BY-2.0 licence (http://creativecommons.org/licenses/by/2.0/), it was originally posted to Flickr by kumon at http://flickr.com/photos/93965446@N00/11740533.

agent system as the addition or removal of a wheel makes a big difference. Also, no system with one fixed central controller is a multi-agent system as it does not allow the removal of this controller or the addition of a new one. A school of fish is on the other hand a multi-agent system, as it remains a school of fish when fish are added or removed. The essence of a multi-agent is thus in our opinion not the large number of agents, but the fact that their exact number is of little importance.

We have worked on different issues related to multi-agent systems and the use of graph theory in their analysis. Graphs arise indeed naturally when representing the interactions between agents. The research presented in this thesis contains two main parts.

In the first part, we consider the ability of a group of agents to remain in formation, that is, to ensure that the agent positions describe a certain constant shape, as represented in Figure 1.2. We describe and analyze the theory of persistent graphs, which we have introduced in [64]. Persistence is a notion related to rigidity. It characterizes the ability of agents to remain in formation by the preservation of distances between agents, where each such distance constraint is the responsibility of one agent.

The second part concerns consensus issues. The domain of consensus contains all multi-agent systems in which agents have to or tend to agree on some value. Examples provided by nature include schools of fish or flocks of birds agreeing on a common direction, as in Figure 1.1. We analyze more particularly two models of opinion dynamics, in which agents are influenced by other agent opinions, provided that their opinions are not too different. Such systems are among the simplest ones for which the interactions depend on the agent values, and we take explicitly this dependence into account.

Let us finally mention that for the sake of conciseness and coherence we do not present in this thesis other results which we have also obtained during the last three years. They include the derivation of an approximative but accurate formulas to compute the distance distribution in Erdős-Rényi graphs [16]. This allowed us to reproduce theoretically and analyze experimental results obtained by Guillaume and Latapy [58] on networks exploration. We have also taken part in a study of local degree leaders in random graphs [15]. A local degree leader is a vertex having a degree larger than all its neighbors. Finally, we have proposed in [17] a simple linear time algorithm for a combinatorial game, based on a representation of the game as an optimization problem over a set of words.

Figure 1.2: Jet-fighters flying in formations, (c) Marcel Grand, www.swissjet.ch

Key contributions

Part I

Brian D.O. Anderson, Vincent D. Blondel, Jean-Charles Delvenne, Baris Fidan and Brad (Changbin) Yu have all collaborated to the research presented in part I.

Our contributions in the first part of this thesis can be summarized as the definition and the analysis of graph-theoretical notions characterizing the properties of structure of unilateral distance constraints in multi-agent formations. In particular, we have defined persistence and constraint consistence, characterizing respectively the ability to maintain a formation shape and the ability to have all constraints satisfied once the formation reaches an equilibrium. We have also defined structural constraint consistence and structural persistence, partially characterizing the ability of the formation to reach an equilibrium. We have provided necessary and sufficient conditions for a directed graph to be constraint consistent, persistent and structural persistent. This lead to explicit algorithms to check persistence for agents lying in a two-dimensional space. For higher dimensions, the absence of known algorithm to check rigidity made it impossible for us to design an algorithm to check persistence in the general case. We have nevertheless obtained a polynomial-time algorithm to check the structural persistence of graphs that are known to be persistent. We have further analyzed the properties of persistent acyclic persistent graphs and of minimally persistent graphs (in two dimensions), that are persistent graphs with minimal sets of edges. For both sorts of graphs, we have provided sequential methods to transform one graph into another, while keeping persistence at each step of the process. Note that we also present a polynomial time algorithm to check persistence for a certain class of graphs that was obtained by Jørgen Bang-Jensen and Tibor Jordán, and that is based on our characterization of persistence. Finally, we have proposed several extensions of these notions, and obtained partial characterizations for some of them.

Part II

The results presented ini part II were obtained in collaboration with Vincent D. Blondel and John N. Tsitsiklis

Our main contribution in part II is a theoretical analysis of two paradigm multi-agent systems, in which we have explicitly used the dependence of the interaction topology in the system state. The two paradigm systems are Krause's opinion dynamics model and a continuous-time variation of it. We have (re)-proved their convergence properties, and proposed a notion of stability with respect to the addition of an agent to explain why some equilibria are almost never reached. We have made some conjectures about their asymptotic behavior for a growing number of agents. We have introduced and analyzed extensions of these systems defined on agent continuum. For those models, we have proved stronger results on the equilibria to which they can converge, and partial convergence results. For Krause's opinion dynamics model, we have also formally established that the system on an agent continuum represents the asymptotic behavior of the system on discrete agents for a growing number of agents, on any finite time interval. Finally, we have analyzed several extensions of these systems, to which we have generalized some of our results.

As a side contribution, we have also proposed a classification of multi-agent systems involving consensus based on their important characteristics.

Publications

The symbol "" marks the preprints that have not been accepted for publication yet.*

On the thesis's first part

Journal articles

Julien M. Hendrickx, Brian D.O. Anderson, Jean-Charles Delvenne and Vincent D. Blondel, **Directed graphs for the analysis of rigidity and persistence in autonomous agents systems.** International journal of robust and nonlinear control, volume 17, issue 10, pages 960-981, July 2007.

Changbin Yu, Julien M. Hendrickx, Baris Fidan, Brian D. O. Anderson, Vincent D. Blondel, **Three and higher dimensional autonomous formations: rigidity, persistence and structural persistence**. Automatica, Volume 43, Issue 3, pages 387-402, March 2007.

Julien M. Hendrickx, Baris Fidan, Changbin Yu, Brian D.O. Anderson and Vincent D. Blondel, **Formation Reorganization by Primitive Operations on Directed Graphs**. To appear in IEEE Transactions on Automatic Control, 2008.

Julien M. Hendrickx, Changbin Yu, Baris Fidan, and Brian D.O. Anderson, **Rigidity and persistence for ensuring shape maintenance in multiagent meta-formations**. To appear in Asian Journal of control's special issue on Collective Behavior and Control of Multi-Agent Systems, 2008.

* Brian D.O. Anderson, Changbin Yu, Baris Fidan and Julien M. Hendrickx, **Control and Information Architectures for Autonomous Formations**. In revision for IEEE Control System Magazine.

Book chapter

Baris Fidan, Brian D.O. Anderson, Changbin Yu and Julien M. Hendrickx, **Persistent autonomous formations and cohesive motion control**. To appear in Modeling and Control of Complex Systems, P. Ioannou and A. Pitsillides (ed.), CRC Press, 2007.

Proceedings of peer-reviewed conferences

Julien M. Hendrickx, Baris Fidan, Changbin Yu, Brian D.O. Anderson and Vincent D. Blondel, **Rigidity and Persistence of Three and Higher Dimensional Formations**. Proceedings of the 1st International workshop on Multi-Agent Robotic Systems MARS 2005, Barcelona (Spain), 39-46, September 2005.

Changbin Yu, Julien M. Hendrickx, Baris Fidan and Brian D.O. Anderson, **Structural Persistence of Three Dimensional Autonomous Formations**. Proceedings of the 1st International workshop on Multi-Agent Robotic Systems MARS 2005, Barcelona (Spain), 47-55, September 2005.

Julien M. Hendrickx, Brian D. O. Anderson and Vincent D. Blondel, **Rigidity and persistence of directed graphs**. Proceedings of the 44th IEEE Conference on Decision and Control, Sevilla (Spain), December 2005.

Julien M. Hendrickx, Baris Fidan, Changbin Yu, Brian D. O. Anderson and Vincent Blondel, **Elementary operations for the reorganization of minimally persistent formations**. Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS2006), Kyoto (Japan), 859-873, July 2006.

Brian D. O. Anderson, Changbin Yu, Baris Fidan and Julien M. Hendrickx, **Use of Meta-Formations for Cooperative Control**. Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS2006), Kyoto (Japan), 2381-2387, July 2006.

Brian D. O. Anderson, Changbin Yu, Baris Fidan and Julien M. Hendrickx, **Control and Information Architectures for Formations**. Proceedings of the IEEE International Conference on Control Applications (CCA), München (Germany), 1127-1138, October 2006.

Changbin Yu, Baris Fidan, Julien M. Hendrickx, Brian D.O. Anderson, **Merging Multiple Formations: A Meta-Formation Prospective**. Proceedings of the 45th IEEE Conference on Decision and Control, San Diego (CA, USA), 4657-4663, December 2006.

Julien M. Hendrickx, Changbin Yu, Baris Fidan, Brian D.O. Anderson, **Rigidity and Persistence of Meta-Formations**. Proceedings of the 45th IEEE Conference on Decision and Control, San Diego (CA, USA), 5980-5985, December 2006.

On the thesis's second part

Journal articles

* Vincent D. Blondel, Julien M. Hendrickx and John N. Tsitsiklis, **Opinion dynamics systems with position dependent communication topologies**, preprint.

* Julien M. Hendrickx, **Order preservation in a generalized version of Krause's opinion dynamics model**, in revision for Physica A.

Proceedings of peer-reviewed conferences

Vincent D. Blondel, Julien M. Hendrickx, Alex Olshevsky, and John N. Tsitsiklis, **Convergence in Multiagent Coordination, Consensus, and Flocking**. Proceedings of the 44th IEEE Conference on Decision and Control, Sevilla (Spain), December 2005.

Julien M. Hendrickx and Vincent D. Blondel, **Convergence of different linear and non-linear Vicsek models**. Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS2006), Kyoto (Japan), 1229-1240, July 2006.

Vincent D. Blondel, Julien M. Hendrickx and John N. Tsitsiklis, **On the 2R conjecture for multi-agent systems**. Proceedings of the European control conference 2007 (ECC 2007), Kos (Greece), 874-881, July 2007.

Other publications

Journal articles

Julien M. Hendrickx, Philippe H. Geubelle and Nancy R. Sottos, **A spectral scheme for the simulation of dynamic mode 3 delamination of thin films**. Engineering Fracture Mechanics, Volume 72, Issue 12, Pages 1866-1891, August 2005.

* Vincent D. Blondel, Julien M. Hendrickx and Raphaël Jungers, **Linear time algorithms for clobber**. submitted.

Vincent D. Blondel, Jean-Loup Gullaume, Julien M. Hendrickx and Raphaël Jungers, **Distance distribution in random graphs and application to complex networks exploration**. Physical Review E, volume 76, 066101, 2007.

Vincent D. Blondel, Jean-Loup Guillaume, Julien M. Hendrickx, Cristobald de Kerchove and Renaud Lambiotte, **Local Leaders in Random Networks**, To appear in Physical Review E.

Proceedings of peer-reviewed conferences

Philippe. H. Geubelle, Julien M. Hendrickx and Nancy R. Sottos, **Spectral Scheme for the analysis of Dynamic Delamination of a Thin Film**. Proceedings of ICF XI International Symposium on Fracture Mechanics, Turin (Italy), March 2005.

Preliminary

Undirected graphs

An *undirected graph* $G(V, E)$ consists of a set V whose elements are called vertices (or nodes), and a set E of unordered pairs of vertices, called edges. In this thesis, we label vertices with integer $1, 2, \ldots, n$, and denote by n their number. We have thus $|V| = n$, where $|V|$ denotes the cardinality of V. The same convention is used for all sets. Vertices are usually represented by points in the plane, and edges as lines connecting those points, as in the example in Figure 1.3. We say that two elements i and j of V are *connected* or *neighbors* if the unordered pair (i, j) or equivalently (j, i) is in E. Alternatively, we may say that i is connected to j or j to i. Both i and j are then said to be *incident* to the edge (i, j). In Figure 1.3 for example, 5 and 6 are connected. A graph is complete if every pair of its vertices is connected by an edge. The complete graph on n vertices is denoted by K_n, and contains $\frac{1}{2}n(n-1)$ edges.

The *degree* d_i of a vertex i is the number of edges to which it is incident, or equivalently its number of neighbors. We note it $d_{i,G}$ when there is a risk of ambiguity on the graph being considered. The vertex 3 in Figure 1.3 has a

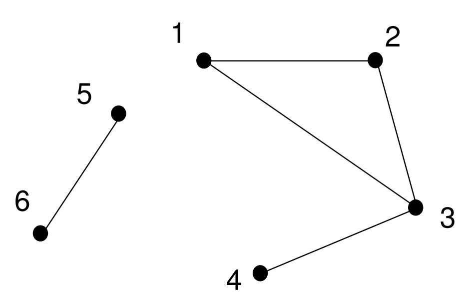

Figure 1.3: Representation of an undirected graph.

degree 3. The sum over all vertices of their degree is twice the number of edges: $\sum_{i \in V} d_i = 2|E|$.

A *path* is a sequence of vertices $v_1, v_2, \ldots, v_p$ such that there is an edge incident to v_i and v_{i+1} for every $i = 1, \ldots, p-1$. In Figure 1.3, $(1,3,4)$ is a path because $(1,3)$ and $(3,4)$ are edges, but $(1,5,6)$ is not a path. The path can also be described by its sequence of edges. We say that the first and the last vertices of a path are *connected by a path*. A graph is *connected* if every two vertices are connected by a path. The graph in Figure 1.3 is for example not connected because there is no path from 5 to 1. A graph is $k-$connected if it is connected, and if it is impossible to disconnect it by removing less than k vertices. In particular, a connected graph is $1-$ *connected*.

A *subgraph* of a graph $G(V,E)$ is a graph $G'(V',E')$ with $V' \subseteq V$ and $E' \subseteq E$. A subgraph G' of G is a *proper subgraph* if it is different from G. A subgraph is *induced* by a set of vertices V' if it is obtained from $G(V,E)$ by removing all vertices of $V \setminus V'$ and the edges that are incident to them. In an induced subgraph, two vertices are connected if and only if they are connected in the initial graph. A *connected component* of a graph is a maximal connected subgraph, that is, a subgraph that is connected and that is the proper subgraph of no larger connected subgraph. In Figure 1.3, the vertices 5, 6 and the edge $(5,6)$ constitute a connected component. On the other hand, the graph containing 1, 2, 3, and the edges connecting them is not a connected component of the initial graph, because it is a proper subgraph of the graph containing 1, 2, 3, 4 and all the edges between them, which is connected. Every graph can be decomposed in a unique way into its connected components. The graph in Figure 1.3 can for example be decomposed into two connected components.

Depending on the applications, one can also define an undirected graph as a set of vertices V and a set of edges E, where each edge is *incident* to two (possibly identical) vertices. This definition allows then self-loops and multiple edges. A *self-loop* is an edge connecting a vertex to itself. It increases the degree of the vertex by 2. *Multiple edges* are different edges connecting the same vertices, i.e. repetitions of a an edge. Graphs without self-loops and multiple edges are sometimes referred to as *simple graphs*. Unless otherwise specified, all graphs in this thesis are simple graphs.

We often consider representation in real vector space, obtained by assigning a position to each vertex. We denote by D the dimension of the space in which the graph is represented.

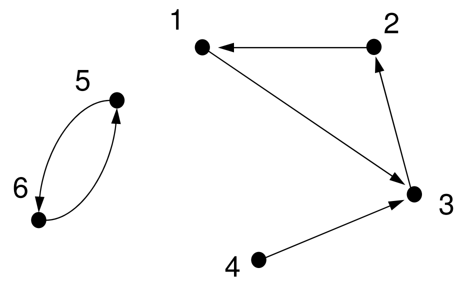

Figure 1.4: Representation of a directed graph.

Directed graphs

In simple words, a directed graph is a graph in which edges have a direction as in Figure 1.4. Two vertices can thus be connected by two edges in opposite directions. More formally, a *directed graph* consists of a set V of vertices and a set E of *ordered* pairs of vertices called (directed) edges. It is denoted by $G(V, E)$, as for the undirected case. We mention explicitly the directed or undirected character of the graph when there is some ambiguity.

We say that a vertex i is connected to j by a directed edge, or that j is a neighbor of i, if $(i, j) \in E$. This does however not imply that i is a neighbor of j, as the edge (j, i) is not necessarily present in the graph. The edge (i, j) is then an *outgoing edge* for i and an *ingoing edge* for j. The direction of the edges is usually represented by an arrow as in Figure 1.4. To every directed graph corresponds an *underlying undirected graph* on the same vertices, obtained by replacing directed edges by undirected ones. The graph in Figure 1.3 is for example the undirected graph underlying the graph in Figure 1.4. As in the undirected case, some applications require the use of a different definition allowing self-loops and multiple edges in directed graphs.

In a directed graph, the *out-degree* d_i^+ of a vertex i is its number of outgoing edges, or number of vertices to which it is connected. The *in-degree* d_i^- is its number of ingoing edges, or number of vertices that are connected to it. The total degree is the sum of the in-degree and the out-degree. We note them $d_{i,G}^+$ and $d_{i,G}^-$ when there is a risk of ambiguity on the graph being considered. In Figure 1.4, the vertex 3 has an out-degree 1 and an in-degree 2, so that its total degree is 3. In every directed graph, there holds $\sum_{i \in V} d_i^+ = \sum_{i \in V} d_i^- = |E|$.

A *directed path* is a sequence of vertices $v_1, v_2, \ldots, v_p$ such that v_i is connected to v_{i+1} by a directed edge for each $i = 1, \ldots, p-1$. We then say that v_i is connected to v_p by a directed path. In Figure 1.4, there is a directed path from 4 to 2 but not from 2 to 4. The path can also be described by the corresponding sequence of directed edges. A *cycle* is a path starting and finishing at the same vertex (and containing at least another vertex), as the one containing 1, 3 and 2.

A directed graph is said to be *weakly connected* if its underlying undirected graph is connected. The graph in Figure 1.4 is not weakly connected, but the subgraph containing $1, 2, 3, 4$ and the edge connecting them is weakly connected. A directed graph is *strongly connected* if for every two vertices, there is a directed path connecting the first vertex to the second. The subgraph containing $1, 2, 3, 4$ is not strongly connected as there is no path from 2 to 4. On the other hand, the subgraph containing $5, 6$ and the two edges connecting them is strongly connected.

Finally, we say that a directed graph is complete if its undirected underlying graph is complete. Since edges can have different directions and since two vertices can also be connected by two edges with opposite directions, there is no unique complete directed graph on n vertices.

Part I

Rigidity and Persistence of Directed Graphs

Prologue

I became aware of the issue of formation shape maintenance on October the 7th in 2004, during the second week of my PhD. On this particular day there was at Cesame a seminar by Brian Anderson on the use of distance constraints to maintain a formation shape. I felt a great interest in this topic, among other reasons because it makes an extensive use of graph theory to analyze some real dynamical systems. So I went to see him after his seminar, and asked if there were some interesting open issue that could be relevant for me to study. He immediately mentioned the possible extension of those results to directed graphs, representing unilateral distance constraints, and gave me a few papers to read in which this idea was explored, [11, 41, 45].

A few days later, we met to talk about ideas I had to solve some conjectures in those papers. A few weeks later, after several discussions with Brian Anderson, with Jean-Charles Delvenne and with my advisor Vincent Blondel, we had a first definition of persistence, extending rigidity to directed graphs, and which is described in our first paper on the subject [64]. After three years and two visits in Canberra to work with Brian and his team, I am now happy to see that, not only we have obtained many results and solved many problems about persistence, but several other people have too. For example, Jørgen Bang-Jensen and Tibor Jordán have provided a polynomial time algorithm to check the persistence of a wide class of graphs representing two-dimensional formations [12], and thus partly closed one of the main open questions in our first paper [64]. Brian Smith, Magnus Egerstedt and Ayanna Howard have analyzed the automatic generation of persistent graphs taking an additional restriction into account, namely that the agents have bounded sensing and communication ranges [119]. Also, in the continuation of earlier works [11] (preexisting ours), John Baillieul and Lester McCoy have investigated the enumeration of graphs that are equivalent to two-dimensional acyclic minimally persistent graphs [10]. Finally, Laura Krick has analyzed in a part of her thesis [79]Section 6.2 some stability issues for formations governed by unilateral distance constraints and modelled by persistent graphs.

In the meantime, and more particularly during this last year, I have often thought about the basic definitions of persistence and related notions, and tried to better understand what they exactly mean, and how they relate to each other. So instead of compiling the different papers published on persistence, I have taken this thesis as an opportunity to present things in a different (but equivalent) way, resulting from this reprocessing work. In particular, the definition of persistence is re-obtained from basic hypotheses on the agents' behavior. Structural persistence is not presented as a condition defined for avoiding a particular problem, but as a part of a more general convergence issue. Simpler proofs are also presented for many results, partly thanks to the new formalism used.

Chapter 2

Introduction to Rigidity and Persistence

2.1 Shape maintenance in formation control

We consider formations of *autonomous agents* evolving in a D-dimensional space. By autonomous agent, we mean here any human controlled or unmanned vehicle that can move by itself and has a local intelligence or computing capacity, such as ground robots, air vehicles or underwater vehicles. We suppose here that the agents have no physical extension, that is, that their positions are single points. A *formation* is a group of autonomous agents with communication capacities, in which the agents collaborate to achieve a common goal. In many applications, the shape of an autonomous agent formation needs to be preserved, that is, the positions of all the agents need to remain constant up to a same translation, rotation, and/or reflection. For example, target localization by a group of unmanned airborne vehicles (UAVs) appears to be best achieved (in the sense of minimizing a localization error) when the UAVs are located at the vertices of a regular polygon [40]. Other examples of optimal placements for groups of moving sensors can be found in [97]. The displacement of the agents can be decomposed into two components, the general formation trajectory, and the local "internal" displacements of the agent seeking to maintain or re-obtain the formation shape. In accordance with the literature, we do not consider here the general displacement of the formation. The sole goal of the agents is thus to maintain the shape of the formation, independently of their position.

Obviously, to maintain a formation shape, agents have to sense some aspect of the formation geometry, i.e., an agent will need to measure some geometrically-relevant variable involving some of the other agents in the formation, in order to apply a control to correct any error in formation shape. For example, if three agents i, j and k have to maintain a triangle, this could be done by sensing and correcting distances between agents, or by sensing and maintaining one distance and two angles. Other sensed measurements in other situations could include bearing relative to north, inclination/declination relative to the horizon, time-difference of arrival at two agents of a pulse transmitted from a third agent, and so on.

2.2 Rigidity and distances preservation

In this thesis, we confine our attention to maintaining the formation shape by explicitly keeping some inter-agent distances constant. Obviously, if the distance between any pair of agent is kept constant, the shape of the formation is maintained. Maintaining all inter-agent distances may however be costly, and not necessary, as it suffices to maintain a certain number of them. In other words, *some* inter-agent distances can be explicitly maintained constant so that all the inter-agent distances remain constant. Constraining the distances is however not the only way to maintain a formation shape, one could for example imagine imposing constraints on the agent bearings (see Section 7.1.6).

The information structure arising from such a system can be efficiently modelled by a graph, where agents are abstracted by vertices and actively constrained inter-agent distances by edges. In Figure 2.1(a) for example, the distances between the agents 1 and 2, 2 and 3, 3 and 4, and 4 and 1 are maintained constant, which is not sufficient to maintain the shape if the formation evolves in a two-dimensional space. On the other hand, constraints in Figure 2.1(c) are sufficient to maintain the formation shape. Note that we always assume the existence of a reference configuration in which all distance constraints are satisfied.

Analyzing the ability of a structure of distance constraint to maintain the shape of a formation can be done using the notion of *rigidity*, which has already been the object of many studies in the past [6, 7, 11, 30, 37–39, 41–45, 47, 56, 57, 72, 84, 85, 95, 104, 105, 109, 110, 123, 124, 129]. A formation is rigid if the fact that all distance constraints are satisfied is sufficient to maintain its shape constant. This notion is reviewed in Chapter 3, in which it is shown that it almost only depends on the graph of distance constraints. It is important not to confuse a formation, a graph abstracting its distance constraints, and a representation or drawing of this graph. At the intuitive level however, one can say that a graph is rigid if its representation cannot be deformed without modifying the length

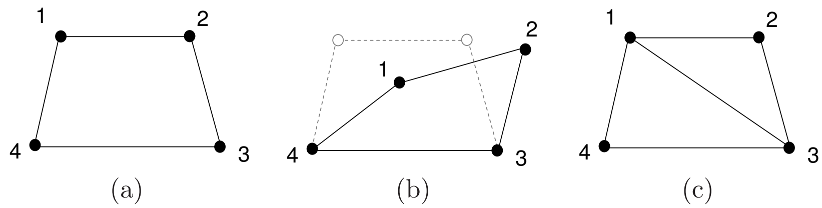

Figure 2.1: For a two-dimensional space, the graph (a) is not rigid because its representation can be deformed without changing the distance along the edges as represented in (b). Such a deformation could not occur to the formation (c), which is therefore rigid.

of its edges. The notion depends thus of course on the dimension in which the graph is drawn, or in which the agents lie. The graphs represented in Figure 2.1 (a) and (c) are thus respectively non-rigid and rigid for two-dimensional spaces, but they would both be rigid in a one-dimensional space, and both nonrigid in a three-dimensional space. Note that this notion of rigidity also corresponds to the undeformability of a structure of joints and beams. This was actually the motivation behind its first formalization [30, 72, 124].

For large graphs, determining whether a given graph is rigid or not is not as immediate as on the examples discussed above. Intuitively, a rigid body in the plane has 3 degrees of freedom, while a set of n unconstrained point-agents has $2n$ degrees of freedom. Each edge of the graph corresponds to one constraint, and removes thus up to one degree of freedom. Rigidity in the plane (respectively in the 3D-space) requires thus $2n-3$ edges (respectively $3n-6$ edges). But this condition is not sufficient as some of these edges may be redundant and may not remove any degree of freedom. Examples of such a graph are presented in Figure 2.2. The issue of recognizing rigid graphs is further reviewed in Chapter 3, and in particular in Section 3.4.

2.3 Persistence and unilateral distance constraints

Unlike in the case of frameworks where distance constraints are guaranteed by the presence of bars between joints, constraints on inter-agent distances in formations have to be maintained by means of measurements and control actions. A distance between two agents can be cooperatively maintained by the two agents, in which case the rigidity theory can directly be applied. But one can also give full responsibility of maintaining the constraint to one agent, which has to maintain its distance from the other constant, this latter agent being unaware of that fact and taking therefore no specific action helping to satisfy

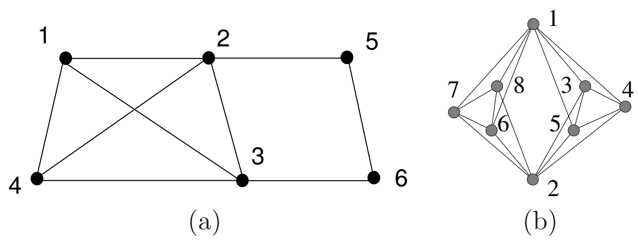

Figure 2.2: Examples of non-rigid graphs having enough edges to be rigid ($2n-3$ and $3n-6$), but that are not rigid due to the redundance of some edges. In the representation (a) in a two-dimensional space, 5 and 6 are not fixed with respect to the other points. In the representation (b) in a three-dimensional space, $(3,4,5)$ and $(6,7,8)$ can rotate around the axis defined by 1 and 2.

the distance constraint.

This unilateral character can be a consequence of the technological limitations of the autonomous agents. Some UAV's can for example not efficiently sense objects that are behind them or have an angular sensing range smaller than 360° [21,48,106]. Some of our collaborators are for example working with agents in which optical sensors have blind three dimensional cones. Unilateral distance constraints can also be desired to ease the trajectory control of the formation, as it allows so-called leader-follower formations [11,44,121]. In such a formation, one agent (leader) is free of inter-agent constraints and is only constrained by the desired trajectory of the formation, and a second agent (first follower) is responsible for only one distance constraint and can set the relative orientation of the formation. The other agents have no decision power and are forced by their distance constraints to follow the first two agents. An example of such a formation is shown in Figure 2.3(c). Finally, it has been argued [11] that for some classes of control law, having the distance constraints maintained by both agents can lead to unstable behaviors in the presence of measurement errors if the agents are not allowed to communicate. Such behaviors can however be avoided by introducing dead-zones at the cost of limited inaccuracy in the preservation of formation shape [52].

A structure of unilateral distance constraints can be represented using a directed graph, where an edge from i to j means that the agent abstracted by i has to maintain its distance to the agent abstracted by j constant. In Figure 2.3(a) for example, the agent 2 is responsible for maintaining its distance to 1 constant, but 1 is not aware of that and cannot be assumed to take any action helping 2 in its task.

Rigidity can still be applied to formations represented by directed graphs as this notion does not depend on the presence or absence of directions for the edges. A directed graph is thus rigid if the undirected graph obtained by forgetting all edge directions is rigid. It still means that the formation shape is maintained provided that the inter-agents distance along all edges are preserved. But, the use of control laws, and particularly of their unilateral character, brings a new issue: *Nothing guarantees a priori that the constraints will be satisfied*. Consider for example the directed graph represented in Figure 2.3 (a), which we suppose to model a formation evolving in a two-dimensional space. It is clearly rigid. Observe however that the agent 4 is responsible for the maintenance of three constraints, while agents 2 and 3 are responsible for only one constraint: each of them has to remain at a constant distance from agent 1, which has no constraint. Suppose that, starting from a reference position in which all constraints are satisfied, 3 moves in the plane, while keeping its distance to 1 constant as represented in Figure 2.3(b). Agent 4 is then unable to simultaneously satisfy the three constraints for which it is responsible. On the other hand, the three other agents satisfy all their constraints, and are thus a priori not forced to change their positions. They are indeed not aware of the constraints for which 4 is responsible, and have no reason for moving to positions where these constraints can be satisfied. So, although the graph is rigid, the formation shape is not always maintained, because some constraints are not necessarily satisfied. Moreover, this is a consequence of the topology of unilateral distance constraints, independently of the control laws used.

The characterization of the directed information structures which can efficiently maintain the formation shape has begun to be studied under the name of "directed rigidity" or "rigidity of a directed graph" [8, 11, 41]. These works included several conjectures about minimal directed rigidity, i.e., directed rigidity with a minimal number of edges for a fixed number of vertices. We have proposed a theoretical framework analyzing these issues in [64], where the name of "persistence" was used in preference to "directed rigidity". Rigidity can indeed be indifferently applied to undirected and directed graphs, and we have argued above that the rigidity of a directed graph is not sufficient for the shape of its corresponding formation to be preserved. This part of the thesis presents this framework and its further extensions [12,63,67–69,133,134]. Note that our goal here is not to design particular control laws achieving shape maintenance for certain types of constraint topologies, but to characterize the topologies for which this goal can be achieved, and this as independently as possible from the particular control laws used.

A simple way of avoiding the problem of agents having sets of constraints that cannot be satisfied is to prevent agents evolving in a D-dimensional space from being responsible for more than D constraints. It is indeed generally pos-

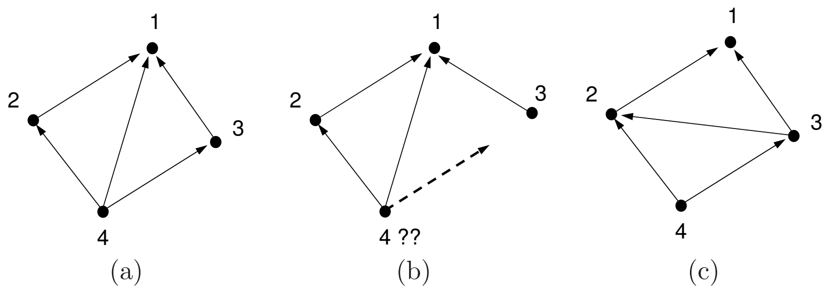

Figure 2.3: The graph (a) is rigid, but the shape of the corresponding 2D-formation is not necessarily maintained as all constraints are not necessarily always satisfied: Agent 3 has one single distance constraint to satisfy, and can freely move on a circle centered on agent 1. If it does so, it becomes impossible for 4 to simultaneously satisfy the three constraints for which it is responsible, while the other agents satisfy all their constraints and have thus nor reason to move, as represented in (b). Such a situation never happens with the formation represented in (c). Formally, the graph (a) is rigid but not persistent because it is not constraint consistent, while (c) is both rigid and constraint consistent, making it persistent. The graph (c) represents also a so-called “leader-follower” formation.

sible to satisfy as many constraints as the dimension of the space in which the position is. This solution is for example preconized in [11]. Remember however that our goal is not to design one particular system but to characterize the directed graphs for which the shape of the corresponding formation can be maintained. And, bounding the number of constraints for each agent unnecessarily excludes some formations whose shape can efficiently be maintained even though some agents are responsible for more than D constraints. Consider for example the graph in Figure 2.4(a) representing the constraint topology of a two-dimensional formation. The agent 4 is responsible for three distance constraints. Its task is thus impossible unless the positions of 1, 2 and 3 are such that these three constraints are compatible, which is for example the case in the reference configuration in which all constraints are satisfied. Observe now in Figure 2.4(b) that 1, 2 and 3 can always satisfy their constraints, and that when they do so, the shape of the sub-formation they constitute is preserved, i.e. is the same as in the reference configuration. As a result, the three constraints of 4 are compatible, and the shape of the whole formation can be maintained. Another way of looking at this is to see that if 4 satisfies two of its constraints, it automatically satisfies the third one. Note that the use of redundant constraints such as in Figure 2.4 can be desired for robustness purpose. It allows maintaining the formation shape even if some distances are temporarily not sensed or controlled due for example to noise or technical problems.

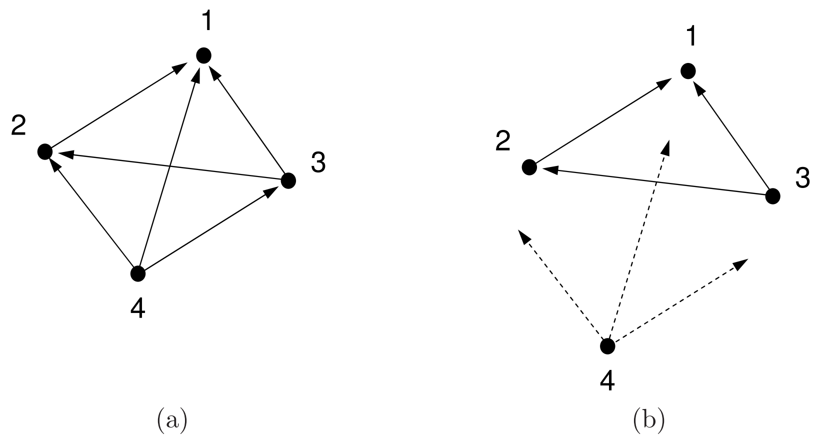

Figure 2.4: Example of 2D-formation whose shape can be maintained although one agent is responsible for three constraints (a). The three constraints of 4 are indeed made compatible by the fact that the constraints of 1, 2 and 3 are sufficient to maintain the shape the triangle $(1, 2, 3)$. The corresponding graph is thus persistent.

In view of the non persistent formation represented in Figure 2.3(a) and the persistent one in Figure 2.4(a), an agent can be responsible for more than D constraints, provided that these constraints are made consistent with each other by the rest of the constraint structure. We introduce in Chapter 4 the notion of *constraint consistence* characterizing this property. Constraint consistence is a notion defined independently of rigidity, and for which the direction of the edges plays a major role. Intuitively, a graph is constraint consistent if in the corresponding formation, the fact that all agents are trying[1] to satisfy the constraints for which they are responsible is sufficient to guarantee that all constraints will be satisfied. As rigidity, this notion depends of course on the dimension of the space in which the formation lies. Figure 2.5 represents a constraint consistent graph and a non constraint consistent graph for a two-dimensional space. In Figure 2.5(a), 1 can freely choose its position, and 2, 3, and 4 have to remain at a constant distance from 1, which is always feasible for them. In Figure 2.5(b), the agents 2, 3 and 4 can freely choose their positions, and 1 is supposed to remain at a constant distance from the three of them, which is generally impossible. The first formation is thus constraint consistent while the second is not. But the analysis of constraint consistence can be more

[1]We do not consider here malevolent agents. The idea of an agent "trying" to satisfy constraints is formalized in Chapter 4 by the notion of equilibrium position, the convergence of the formation to which is further considered in Chapter 5.

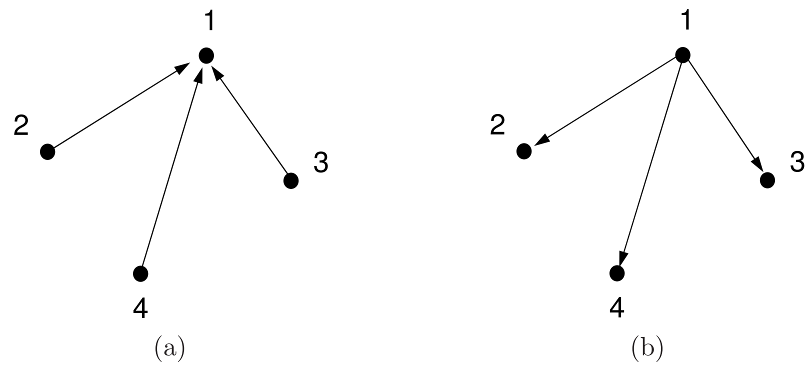

Figure 2.5: Examples of constraint consistent (a) and non-constraint consistent (b) planar formations, sharing the same undirected underlying graph. In (a), 2, 3, and 4 can always remain at a constant distance from 1, which can freely choose its position. In (b), it is generally impossible for 1 to maintain its distance to 2, 3 and 4 constant while they freely move in the plane.

complicated. For example, the graph in Figure 2.4 is constraint consistent, as we have already argued that all its corresponding distance constraints could be satisfied. Also, it is not evident at first sight to determine if the graphs represented in Figure 2.6 are persistent. It will be seen that the persistence of (a) follows from Corollary 4.2. The persistence of (b) and (c) follow from Theorem 4.5, the general criterion that we will obtain to determine if a graph is persistent or not. The persistence of (c) can also be obtained in an easier way, taking into account the fact that the graph can be made acyclic by removing three edges leaving the same vertex. Finally, it will be seen in Figure 4.5 that (d) is not persistent.

Persistence, introduced in Chapter 4 to characterize the ability of the structure of unilateral constraint to maintain the formation shape, is the combination of rigidity and constraint consistence. A persistent formation is one for which the fact that all agents are trying to satisfy their constraints is sufficient for the shape to be maintained. So, not only the shape is maintained when all constraints are satisfied (rigidity), but all constraints are necessarily satisfied when all agents are trying to satisfy the constraints for which they are responsible (constraint consistence). This simple equivalence, formally proved in Section 4.2, is summarized below.

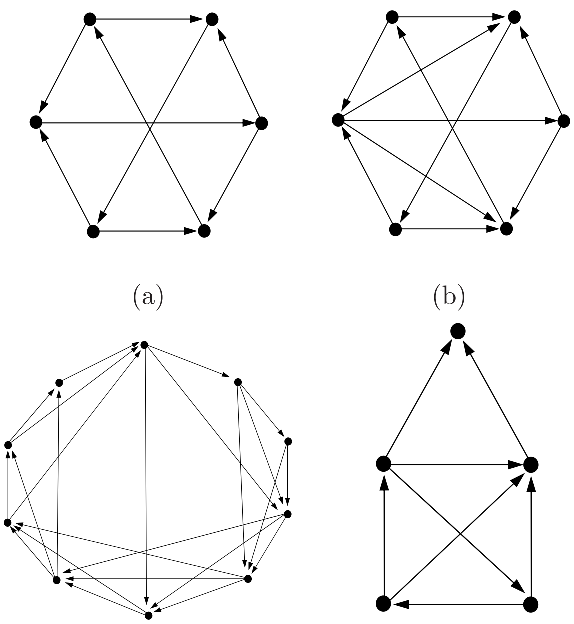

(a) (b) (c) (d)

Figure 2.6: Examples of graph for which it is uneasy to determine a priori if they are persistent. Using results obtained in the next chapters, one can prove that they are actually all persistent except (d).

Constraint Consistence	:	*Agents trying to satisfy constraints*	$\Rightarrow$ *All constraints satisfied*
$+$			
Rigidity	:	*All constraints satisfied*	$\Rightarrow$ *Shape preserved*
$=$			
Persistence	:	*Agents trying to satisfy constraints*	$\Rightarrow$ *Shape preserved*

2.4 Outline of part I

We begin by reviewing rigidity and some related concepts in Chapter 3. We then introduce and give a full characterization of persistence and constraint consistence in Chapter 4. These notions are based on an analysis of equilibrium situation for formations, but do not take into account the issue of reaching this equilibrium. The latter issue is considered in Chapter 5, leading to the notion of structural persistence, for which we also provide a full characterization. We then focus in Chapter 6 on particular classes of graphs, for which the analysis of the notions introduced can be pushed further. We consider acyclic graphs in Section 6.1, minimally persistent graphs in Section 6.2, and graphs with three degrees of freedom in Section 6.3, the two latter sections considering only two-dimensional spaces. Minimally persistent graphs are persistent graphs with a minimal set of edges. Graphs with three degrees of freedom in a two-dimensional space are those whose corresponding formation trajectory can fully be controlled using only local agent decisions, i.e., for which no decision needs to be taken collectively. We present in Chapter 7 further directions for research that could be undertaken. We mention some track that could be followed, and partial results when available. We close then this first part of the thesis by the concluding remarks in Chapter 8. To ease the reading, the main definitions related to this part of the thesis are listed in Appendix A.

Since this part of the thesis is focussed on understanding the main notions related to persistence, we do not present all the results that we have obtained during these three years. In particular, we omit some recent results on merging of persistent formations [70, 71]. We also only present the simplest method to build all two-dimensional minimally persistent graphs in Section 6.2, and not some other methods to build them, and results on the inexistence of certain classes of methods [68].

Chapter 3

Rigidity of Graphs

In this chapter we review the definition and main characteristics of rigidity. Intuitively, a graph representation is rigid if it cannot be continuously deformed without modifying the edge lengths. We formalize this notion in Section 3.1. In Section 3.2 we introduce the notion of infinitesimal rigidity, which, roughly, is a first-order analysis of rigidity. It allows a simpler analysis using linear-algebraic tools. We describe the close link between infinitesimal rigidity and rigidity, and see that the latter is actually a stronger condition than the former as every infinitesimally rigid representation is rigid. We see in Section 3.3 that rigidity actually only depends on the graph and not on the particular representation that is considered (up to a zero-measure set of representations forming particular ill-conditioned cases). Finally, we introduce in Section 3.4 the notion of minimal rigidity corresponding to rigidity with a minimal set of edges, and show how this notion can be used to give a necessary condition for rigidity. We also present a force based approach of rigidity in Section 3.5 and a generalization of rigidity called tensegrity in Section 3.6. Note that in this chapter, we do note take into account the possible directions of the edges, nor the fact that they are directed or undirected. All definitions can thus be applied to both directed and undirected graphs. For the sake of clarity, all graphs represented are however undirected, as the possible direction of the edges has no influence on the rigidity properties of the graph.

Before beginning, we believe it is worth emphasizing that rigidity is a notion defined for representations of graphs and not for graphs. It is then a non-trivial result presented in Section 3.3 that rigidity is a *generic property of a graph*. A property of graph representations is generic for graphs if for each graph, either almost all its representations have the property or almost all of them do not have the property, where "almost all representations" refer to all representations excluding a set of measure zero.

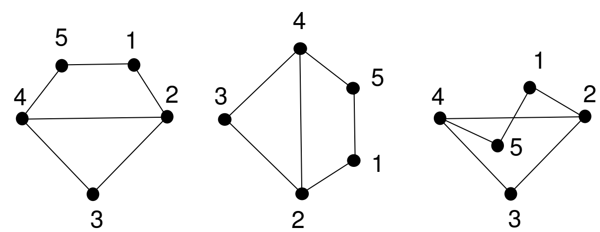

Figure 3.1: Three representations of the same graph in $\Re^2$, all realizations of the same distance set. (a) and (b) are congruent as they can be obtained one from each other by rotations and translations, but (c) is not congruent to (b) nor to (c).

3.1 Rigidity

A representation of a graph $G(V, E)$ in $\Re^D$ is a function $p : V \rightarrow \Re^D : i \rightarrow p_i$. We say that $p_i \in \Re^D$ is the *position* of the vertex i, and define the distance between two representations p and q of the same graph by

$$d(p, q) = \max_{i \in V} ||p_i - q_i|| ,$$

endowing the set of representations with a structure of metric space. Two representations p and q are *congruent* if they can be obtained one from each other by a composition of Euclidean transformations (i.e., translations, rotations, and reflections), which is actually the case if and only if the distance between the positions of every pair of vertices (connected by an edge or not) is the same in both of them: $||p_i - p_j|| = ||q_i - q_j||$ for all $i, j \in V$.

A *distance set* $\bar{d}$ for G is a set of distances $d_{ij} \geq 0$, defined for all edges $(i, j) \in E$. A distance set is *realizable* if there exists a representation p of the graph for which $||p_i - p_j|| = d_{ij}$ for all $(i, j) \in E$. Such a representation is then called a *realization*. Intuitively, a distance set $\bar{d}$ is realizable if it is possible to draw the graph such that the distance between the positions of any pair of vertices i, j connected by an edge is d_{ij}. Note that each representation p of a graph induces a realizable distance set (defined by $d_{ij} = ||p_i - p_j||$ for all $(i, j) \in E$). Figure 3.1 shows three realizations of the same distance set for a graph, two of which are congruent.

Definition 3.1. *A representation p of a graph G is* rigid *if there is a neighborhood of p in which all realizations of the distance set induced by p and G are congruent to p.*

As an example of the application of this definition, Figure 3.2(c) shows a graph representation p and a realization p' of the induced distance set - the

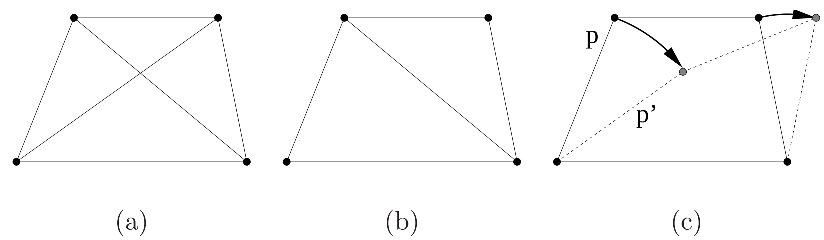

Figure 3.2: Example of rigid (a),(b) and non-rigid (c) representations. (b) is also minimally rigid, as the removal of any edge would lead to a loss of rigidity.

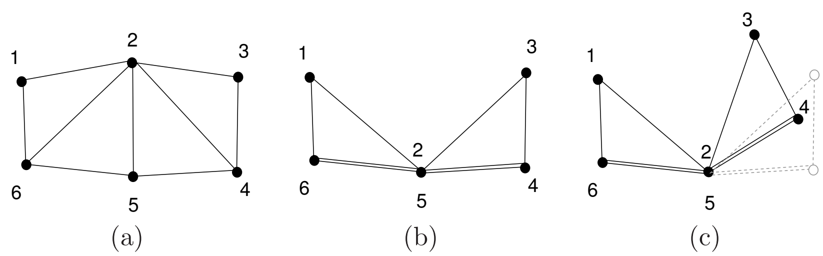

Figure 3.3: Rigid (a) and not rigid (b) representations of the same graph. The representation (b) is not rigid because there exists (arbitrarily close) realizations of its induced distance set to which it is not congruent, as in (c).

lengths of all edges are indeed the same in p and p' - which is not congruent to p. Since such realizations p' can be found arbitrarily close to p, this latter is not rigid. On the other hand, it is possible to prove that the representations in Figures 3.2(a) and (b) are rigid. Figure 3.3 shows two representations of a same graph, one of which is rigid and the other is not. This shows that rigidity does depend on the representation considered, and not only on the graph.

Before going further in our analysis, note that the definition of rigidity only requires the congruence of all realizations in a certain neighborhood of p, so that this notion is related to local variations. If large variations are considered, a realization p' of a distance set induced by a rigid graph representation p is not necessarily congruent to p, as shown on the example in Figure 3.4. A graph representation whose distance set defines a unique realization (up to congruence) is called a *globally rigid representation*. We do not consider global rigidity in this thesis, and refer the reader for example to [61] for more information on this topic.

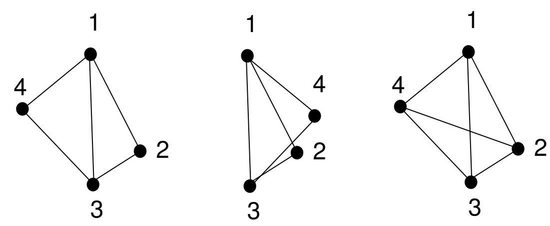

Figure 3.4: The two-dimensional graph representation in (a) is provably rigid but is not globally rigid, as (b) shows a realization of its induced distance set that is not congruent with the representation in (a). Since it is rigid there exists however a neighborhood of the representation (a) in which no such realization can be found. The graph representation in (c) is on the other hand globally rigid.

Note also that rigidity can equivalently be defined in the following more compact way. For a graph $G(V, E)$,let us define the morphisms

$$\Phi_{V\times V} : (\Re^D)^n \to (\Re^+)^{n^2} : p \to \{d_{ij}\}_{(i,j)\in V\times V} = \{||p_i - p_j||\}_{(i,j)\in V\times V}$$

and

$$\Phi_E : (\Re^D)^n \to (\Re^+)^{|E|} : p \to \{d_{ij}\}_{(i,j)\in E} = \{||p_i - p_j||\}_{(i,j)\in E}$$

mapping the space of sets of n $D-$dimensional points (representations) on the distances between respectively all the pair of points, and all the pair of points whose corresponding vertices are connected by an edge in the graph G. We say that p *is rigid as a representation of* G if there exits a $\delta > 0$ such that for all $p' \in B(p, \delta)$, $\Phi_E(p') = \Phi_E(p)$ implies that $\Phi_{V\times V}(p) = \Phi_{V\times V}(p')$.

3.2 Infinitesimal rigidity

Rigidity is conveniently analyzed by linearizing the constraints on the positions. We only present here an intuitive version of this linearization, and refer the reader to [6, 7] or [124] for a more formal analysis.

Let p be a graph representation, which we consider here not as a function defined on V but as a vector of $\Re^{nD}$ (with $n = |V|$) containing all vertices positions:

$$p = \begin{pmatrix} p_1^T & p_2^T & \dots & p_n^T \end{pmatrix}^T .$$

Let δp be an infinitesimal displacement, that is a difference between two representations p and $p + \delta p$ so small that $||\delta p||_2^2$ is negligible[1] δp is *a Euclidean displacement* if it is a combination of infinitesimal translations and rotations. More formally, an infinitesimal displacement δp is Euclidean if there exists a time-continuous length-preserving transformation $E : \Re^D \times \Re^+ \to \Re^D : (x,t) \to E(x,t)$ such that $\delta p_i = K \frac{dE(x,t)}{dt}|_{p_i,0}$ holds with the same K for all $i \in V$. We denote by Eu_p the set of Euclidean infinitesimal displacements for a representation p. If the convex hull of p has a dimension at least $D-1$, the dimension of Eu_p is $f_D = \frac{1}{2}D(D+1)$, corresponding to D independent translations along all independent directions in $\Re^D$ and $\frac{1}{2}D(D-1)$ independent rotations for each pair of independent directions in $\Re^D$, a rotation being uniquely defined by a plane and an intensity. If the convex hull of p has a dimension smaller than $D-1$, then the dimension of Eu_p is smaller than f_D, as some rotations corresponding to planes orthogonal to the convex hull of p can be expressed as combinations of translations. Applying for example a rotation to a zero-dimensional object either has no effect or is equivalent to a translation. For the sake of conciseness, we do not explicitly treat this case here. Unless otherwise specified, we suppose thus that the convex hull of the positions in p has a dimension $D-1$ or D, which is almost always the case if the number of vertices is at least D. The case of graphs with a smaller number of vertices is considered at the end of Section 3.4 and in Proposition 3.2. Note that reflections are not contained in infinitesimal Euclidean displacements as they are by essence transformations of finite size.

The distance between two vertices i, j in the representation $p + \delta p$ is the same as in the representation p if there holds

$$\begin{aligned} ||p_i - p_j||^2 &= ||p_i + \delta p_i - p_j - \delta p_j||^2 \\ &= ||p_i - p_j||^2 + 2(p_i - p_j)^T(\delta p_i - \delta p_j) + ||\delta p_i - \delta p_j||^2 . \end{aligned}$$

Since $||\delta p||_2^2$ is neglected, this is equivalent to

$$(p_i - p_j)^T(\delta p_i - \delta p_j) = 0. \tag{3.1}$$

We say that an infinitesimal displacement δp is *admissible*[2] by a representation p of a graph $G(V,E)$ if (3.1) holds for all $(i,j) \in E$. Intuitively, δp is admissible for p and G if it preserves the distance between the positions of

[1] Some authors [124] define more formally an infinitesimal displacement as an assignment of velocities to the vertices. We prefer here using the idea of small position variation as it provides a better intuitive insight in the sequel.

[2] Some authors simply call infinitesimal displacements [124] what we call admissible infinitesimal displacements. Admissible infinitesimal displacements are indeed sufficient to analyze rigidity so that inadmissible ones do not need to be considered. When analyzing persistence as we do in the next sections however, it is essential to take inadmissible infinitesimal displacements into account.

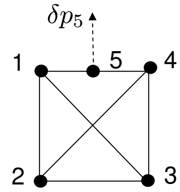

Figure 3.5: Example of rigid representation that is not infinitesimally rigid. All sufficiently close realizations of the distance set induced are congruent to this representation. However, one can verify that the infinitesimal displacement represented (with $\delta p_1 = \delta p_2 = \delta p_3 = \delta p_4 = 0$) is admissible but does not correspond to a Euclidean displacement of the whole representation.

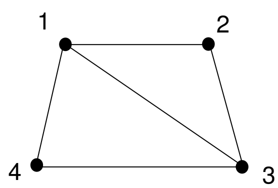

(a)

$$R_{G,p} = \begin{pmatrix} (p_1-p_2)^T & (p_2-p_1)^T & 0 & 0 \\ 0 & (p_2-p_3)^T & (p_3-p_2)^T & 0 \\ (p_1-p_3)^T & 0 & (p_3-p_1)^T & 0 \\ (p_1-p_4)^T & 0 & 0 & (p_4-p_1)^T \\ 0 & 0 & (p_3-p_4)^T & (p_4-p_3)^T \end{pmatrix}$$

(b)

Figure 3.6: Rigidity matrix (b) associated to a graph representation (a).

any two vertices that are connected by an edge in G, and thus if $p + \delta p$ is a realization of the distance set induced by p. It can be easily proved that all Euclidean infinitesimal displacements are admissible by p and G. Following the intuitive meaning of rigidity, a representation p should be infinitesimally rigid if for any admissible δp, $p + \delta p$ and p are "congruent", that is, can be obtained by Euclidean transformations one from each other. We have thus the following definition:

Definition 3.2. *A representation of a graph is* infinitesimally rigid *if all its admissible infinitesimal displacements are Euclidean.*

It can be proved that every infinitesimally rigid representation is rigid [6,7]. The converse only holds for almost all representations, i.e., there is a zero-measure set of representations that are rigid but not infinitesimally rigid, such as the one shown in Figure 3.5.

Since all constraints defining the admissibility of the infinitesimal displacements are linear, they can be written under the more compact form

$$R_{G,p}\delta p = 0, \tag{3.2}$$

where the $R_{G,p} \in \Re^{|E| \times Dn}$ is the *rigidity matrix* associated to G and p. Vertices of the graphs correspond to D columns of the matrix, and each line of the latter corresponds to an edge (i, j) in G and is equal to

$$\begin{pmatrix} \dots & 0 & (p_i - p_j)^T & 0 & \dots & 0 & (p_j - p_i)^T & 0 & \dots \end{pmatrix},$$

the non-zero column being the $(i-1)D+1^{st}$ to the iD^{th} and the $(j-1)D+1^{st}$ to the jD^{th}. An example of correspondence between a graph representation and its rigidity matrix is shown in Figure 3.6. A representation is then infinitesimally rigid if each infinitesimal displacement satisfying (3.2) is Euclidean, that is if $\mathrm{Ker} R_{G,p} \subseteq Eu_p$. Since $Eu_p \subseteq \mathrm{Ker} R_{G,p}$ always trivially holds because every Euclidean displacement is admissible by any representation of any graph, this condition is then equivalent to $\mathrm{Ker} R_{G,p} = Eu_p$ and can be expressed in terms of the matrix rank as

$$rank R_{G,p} = Dn - \dim Eu_p = Dn - f_D = Dn - \frac{1}{2}D(D-1).$$

The infinitesimal rigidity of p is thus equivalent to the presence of $Dn - f_D$ independent lines in $R_{G,p}$. Note that since $Eu_p \subseteq \mathrm{Ker} R_{G,p}$, a rigidity matrix never has a rank larger than $Dn - f_D$.

3.3 Generic rigidity of graphs

Infinitesimal rigidity is a notion defined for representations of graphs, and not for graphs. We now show however that they (almost) only depend on graphs.

Definition 3.3. *Let P be a property defined for graph representations. A graph is* generically P *if the set of its representations not having the property P has zero measure. A graph is* generically not P *if the set of its representations having the property P has zero measure. The property is* a generic property *if every graph is either* generically P *or* generically not P.
For the sake of conciseness, we omit the word "generically" in the sequel, unless when the context could allow ambiguities.

Intuitively, a generic notion is one that, although defined for representations of graphs, only depends on the graph to the exception of a zero-measure set of particular cases. Having a convex hull of dimension D is for example a property defined for a representation (in $\Re^D$) that is a generic notion.

For a representation p of a graph G, we say that a set of edges is *independent* if the corresponding lines in $R_{G,p}$ are linearly independent. It can be proved [6, 7] that the independence of edges is a generic notion. As a consequence, so is infinitesimal rigidity. We call *generic representations* those representations for which every generically independent set of edges is independent.

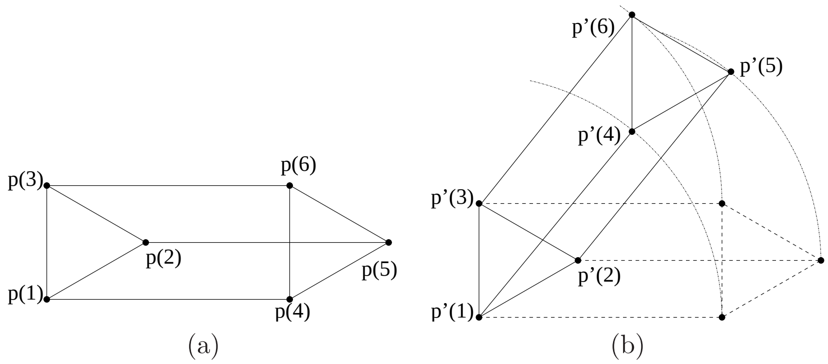

Figure 3.7: p is a representation of a graph that can be proved (by Theorem 3.2) to be rigid in $\Re^2$. However, (b) shows a realization p' of the distance set induced by p, but which is not congruent to p. p is thus a non-generic representation.

Intuitively, any representation that is not a particular case (in term of independence of rigidity matrix lines) is a generic representation. Clearly, almost all representations are generic representations. Trivial examples of non-generic representations involve collinear positions or vertices with same positions as in Figure 3.3, but more surprising examples exist as well; see e.g. Figure 3.7. We can define the *generic rank of the rigidity matrix of a graph* as the rank of the rigidity matrix associated to a generic representation of a graph. It follows from the definition of generic representation that this is a well-defined notion which only depends on the graph.

It can be proved that rigidity is also a generic notion, and that a graph is rigid if and only if it is infinitesimally rigid[3] [6, 7]. For the sake of conciseness, we simply call such a graph a *rigid graph* in the sequel. The rigidity of the graph depends of course on the dimension in which its representations are considered. We generally designate by D the dimension in which the graph is represented, without mentioning it explicitly when the context allows no ambiguity.

Let us also mention that when a graph is not rigid, none of its representations are infinitesimally rigid, but some of it non-generic representations forming a zero-measure set are rigid. One can for example verify that any representation for which all vertices have the same positions are rigid provided that the graph is connected. On the other hand, a rigid graph may have some non trivial non-generic representations that are not rigid, see for example Figure 3.7.

[3]Formally, generic rigidity is equivalent to generic infinitesimal rigidity

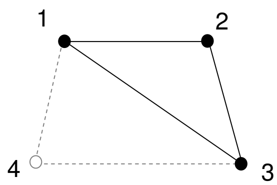

(a)

$$R_{G',p'}$$

$$R_{G,p} = \begin{pmatrix} (p_1-p_2)^T & (p_2-p_1)^T & 0 & 0 \\ 0 & (p_2-p_3)^T & (p_3-p_2)^T & 0 \\ (p_1-p_3)^T & 0 & (p_3-p_1)^T & 0 \\ (p_1-p_4)^T & 0 & 0 & (p_4-p_1)^T \\ 0 & 0 & (p_3-p_4)^T & (p_4-p_3)^T \end{pmatrix}$$

(b)

Figure 3.8: The rigidity matrix of the subgraph representation (a) is a submatrix (b) of the graph rigidity matrix.

3.4 Minimal rigidity for a characterization of rigid graphs

The analysis of rigidity is made easier by using the notion of minimal rigidity. A a graph is *minimally rigid* if it is rigid and if the removal of one or several of its edges automatically destroys its rigidity. For example, the graph represented in Figure 3.2(b) is minimally rigid while the one in Figure 3.2(a) is not. It is indeed possible to remove an edge to the graph represented in (a) without losing rigidity. Clearly, every rigid graph contains a minimally rigid subgraph on all its vertices. The converse is also true, so that a graph is rigid if and only if it contains a minimally rigid graph on all its vertices. Indeed, adding edge to a graph correspond to adding lines to its rigidity matrix, which never decreases its rank. Rigidity is thus preserved by addition of edges, and a graph containing a (minimally) rigid subgraph on all its vertices is thus rigid.

Consider a generic representation p of a minimally rigid graph. All lines of the associated rigidity matrix $R_{G,p}$ are linearly independent. If it was not the case one edge could indeed be removed without affecting the rank of $R_{G,p}$, preserving therefore rigidity and preventing the initial graph from being minimally rigid. The number of edges in a minimally rigid graph is thus equal to the generic rigidity matrix rank, that is $Dn - f_D$ since the representation is rigid. Moreover, if a graph is rigid and contains $Dn - f_D$ edges, the removal of any edge decreases its generic rank and destroys its rigidity. A rigid graph is thus minimally rigid if and only if its number of edges is the smaller one allowing rigidity for its number of vertices, that is $Dn - f_D$.

The independence of the lines in a minimally rigid graph also has consequences on the number of edges in the subgraphs of G. The rigidity matrix $R_{G',p}$ of a subgraph G' of G for the restriction of p to its vertices is the restriction of $R_{G,p}$ to the columns and lines corresponding respectively to the vertices and edges that are present in G', as can be seen on the example in Figure 3.8.

One can verify that its lines are thus obtained from lines of $R_{G,p}$ by removing zero-elements. As a consequence, they are linearly independent if and only if the corresponding lines in $R_{G,p}$ are independent. The rigidity matrix of any subgraph of G contains thus only independent lines, which again implies that any subgraph $G''(V'', E'')$ of G (on more than D vertices) contains at most $D\,|V''| - f_D$ edges. The above discussion is summarized by the following result:

Theorem 3.1. *A graph $G(V, E)$ with at least D vertices is rigid in $\Re^D$ if and only if it contains a minimally rigid (in $\Re^d$) subgraph $G'(V, E')$ on all its vertices. Moreover, if a graph $G'(V, E')$ is minimally rigid, then*

- *There holds $|E'| = D\,|V| - f_D$*
- *For every subgraph $G''(V'', E'')$ of G there holds $|E''| \leq D\,|V''| - f_D$*

As a consequence, a rigid graph $G(V, E)$ is minimally rigid in $\Re^d$ if and only if $|E| = D\,|V| - f_D$ holds.

For two-dimensional spaces, the necessary condition for independence of the lines, i.e. absence of subgraph $G(V'', E'')$ with $|E''| > 2\,|V''| - 3$, is actually also sufficient, leading to necessary and sufficient condition for rigidity:

Theorem 3.2 (Laman [84])**.** *A graph $G(V, E)$ (with $|V| \geq 2$) is rigid in $\Re^2$ if and only if there exists a subgraph $G'(V, E)$ on the same vertices such that*

- *There holds $|E'| = 2\,|V| - 3$*
- *For all subgraphs $G''(V'', E'')$ of G' with $V'' \geq 2$, there holds $|E''| \leq 2\,|V''| - 3$.*

As a consequence, a rigid graph is minimally rigid if it is rigid and if $|E| = 2\,|V| - 3$ holds.

Note that several algorithms have been proposed to check the criterion of Theorem 3.2 in a polynomial time, see [74,101] for example or [124] for a survey on minimal rigidity. The equivalence of Theorem 3.2 does not hold for larger dimensions, as shown by the so-called "double-banana" graph in Figure 3.9, and no result similar to Theorem 3.2 is known for these dimensions. However, the following result allows in some situation to guarantee the rigidity of some subgraphs.

Proposition 3.1. *Let $G(V, E)$ be a minimally rigid graph (in $\Re^D$). A subgraph $G'(V', E')$ of G on at least D vertices is rigid if and only if there holds $|E'| \geq D\,|V'| - f_D$. In that case, there necessarily holds $|E'| = D\,|V'| - f_D$ and the subgraph G' is also minimally rigid.*

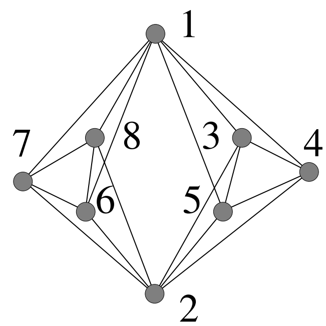

Figure 3.9: Example of a graph satisfying the necessary condition for rigidity of Theorem 3.1 with $D = 3$ but that is not rigid in $\Re^3$. The parts of the graphs defined by $\{3, 4, 5\}$ and $\{6, 7, 8\}$ can rotate around the axis defined by the positions of vertices 1 and 2.

Proof. Consider a generic representation p of G and its restriction p' to G', which clearly is also generic. Since G is minimally rigid, all lines of $R_{G,p}$ are linearly independent, and so are therefore those of $R_{G',p'}$. The rank of $R_{G',p'}$ is thus equal to the number of its lines, that is, to $|E'|$. The result follows then from the fact that a representation of a graph on at least D vertices is infinitesimally rigid if and only if the rank of its rigidity matrix is $D\,|V| - f_D$. $\square$

We have assumed until here that the number of vertices was at least as large as the dimension of the space in which the graph is represented. If this assumption is not satisfied, then it can be proved that a graph is rigid if and only if it is the complete graph. The sufficiency of this condition is trivial as, for generic representations, a complete graph implies constraints on the distance between all vertices positions. To understand its necessity, consider the fact that n vertices define generically a $n-1$ dimensional affine variety. If the graph is rigid in $\Re^D$, it must be rigid in this affine variety, and contains therefore at least $(n-1)n - f_{n-1} = \frac{1}{2}(n-1)n$ edges, which is only the case if the graph is complete. As a result, one can verify that a rigid graph $G(V, E)$ always satisfies $|E| \geq nD - f_D$.

Proposition 3.2. *A graph G on less than D vertices is rigid in $\Re^D$ if and only if it is a complete graph.*

Finally, we present the following corollary with an intuitive meaning that in a rigid graph, each vertex must have sufficiently many constraints on its position.

Corollary 3.1. *If a graph $G(V, E)$ contains a vertex with a degree smaller than $\min(D, |V| - 1)$, it is not rigid.*

Proof. If $|V| < D$, the result follows from Proposition 3.2. If $|V| \geq D$, the result can for example be obtained by applying Theorem 3.1, taking for G'' the subgraph of G' obtained by removing the vertex with degree smaller than D. □

3.5 Force based approach to rigidity matrix

Our obtention of the rigidity matrix and of the related results is based on distance constraints, consistently with the motivations of this work. The same results can actually be obtained using a different approach based on force equilibrium, which we briefly describe here. For more information we refer the reader to [124].

Suppose that the vertices represent joints and the edges beams between joints. Each joint i has a position $p_i \in \Re^3$. Beams can transmit forces parallel to their axis. The beam corresponding to the edge (i,j) can thus apply simultaneously two opposite forces $\lambda_{ij}(p_i - p_j)$ and $\lambda_{ij}(p_j - p_i)$ on the joint i and j, where $\lambda_{ij} = \lambda_{ji} \in \Re$ is proportional to the intensity of the force and inversely proportional to the length of the beam. If we suppose moreover than a force $f_i \in \Re^3$ can be applied on each joint, the equilibrium condition of the joint is described by the following three-dimensional relation

$$-f_i = \sum_{j:(i,j)\in E} \lambda_{ij}(p_i - p_j).$$

Let now $f = (f_1^T, \ldots, f_n^T)^T$, one can verify that the equilibrium conditions of all joints can be written under the compact form

$$-f = R_{G,p}^T \lambda,$$

where $R_{G,p}$ is the rigidity matrix, and λ a vector containing all λ_{ij}, in an order consistent with that of the edges in the rigidity matrix. As an example, the system of equations describing the equilibrium conditions of all joints in a framework modelled by the graph in Figure 3.6 is

$$-f = \begin{pmatrix} (p_1 - p_2) & 0 & (p_1 - p_3) & (p_1 - p_4) & 0 \\ (p_2 - p_1) & (p_2 - p_3) & 0 & 0 & 0 \\ 0 & (p_3 - p_2) & (p_3 - p_1) & 0 & (p_3 - p_4) \\ 0 & 0 & 0 & (p_4 - p_1) & (p_4 - p_3) \end{pmatrix} \begin{pmatrix} \lambda_{12} \\ \lambda_{23} \\ \lambda_{13} \\ \lambda_{14} \\ \lambda_{34} \end{pmatrix}.$$

Intuitively, such a framework is rigid if the application of forces to the joint does not result in internal deformations, but only in rotations and translations of the framework as a whole, i.e., Euclidean displacements. By the superposition principle, a rigid framework is thus one that undergoes no deformation when loaded by a set of forces $f \in Eu_p^{\perp}$. Note that in a three-dimensional space, the orthogonality to Eu_p conveniently expressed by $\sum_i f_i = 0$ and $\sum_i p_i \times f_i = 0$. Let us now define formally static rigidity.

Definition 3.4. *A representation p of a graph G is statically rigid if for any load $f \in Eu_p^{\perp}$, there exists a $\lambda \in \Re^{|E|}$ such that all joints are at equilibrium, that is, a λ solution of $-f = R_{G,p}^T \lambda$.*

More compactly, p is statically rigid if $Eu_p^{\perp} \subseteq \mathrm{Im}R_{G,p}^T$, or equivalently, if $\left(\mathrm{Im}R_{G,p}^T\right)^{\perp} \subseteq Eu_p$. Using the well known relation $\mathrm{Ker}A^T = (\mathrm{Im}A)^{\perp}$, we finally obtain that p is statically rigid if and only if $\mathrm{Ker}R_{G,p} \subseteq Eu_p$. Since this condition is also necessary and sufficient for infinitesimal rigidity, the two notions are equivalent.

Note finally that a more comprehensive notion of rigidity exists in civil and mechanical engineering, allowing beams to also transmit orthogonal forces and torque. This is equivalent to consider structures of more complex geometric constraints, such that constraints on the angle between edges. Such extended notions of rigidity are however out of the scope of this thesis.

3.6 Tensegrity

The initial motivation of rigidity was to characterize frameworks of bars, that impose a certain distance between pairs of joints and can transmit any force parallel to their axes. Tensegrity has been introduced [54] to characterize frameworks that also contain cables and struts. A cable imposes a *maximal distance* between two joints, and can be loaded in tension. A strut imposes a *minimal distance* between two joints, and can be loaded in compression. Simple examples of such structures are shown in Figure 3.10, and more complex structures that cannot be deformed are shown in Figure 3.11. In this section, we give a formal definition of tensegrity, or more formally of "rigidity of a tensegrity framework", and mention some of its properties. For more information, we refer the reader to [29, 114].

A *tensegrity framework* $G(V, B, C, S)$ is a set V of vertices, together with pairwise disjoint sets of pairs of vertices, referred to as bars, cables, and struts respectively. We denote by E the union $B \cup C \cup S$ of these three sets. The notion of representations and distance set can immediately be extended. We say that a representation p of G is a *tensegrity-realization* of a distance set d

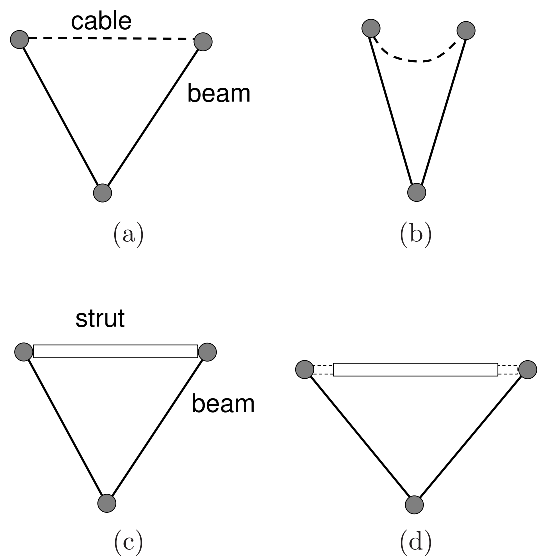

Figure 3.10: Examples of structures containing two bars and one cable (a), or two bars and one strut (c). Similar structures containing only bars could not be deformed. However, the structure (a) can be deformed by reducing the distance between the extremities of the cable as in (b), and the structure (c) by increasing the length of the strut as in (d). These structures are thus not rigid.

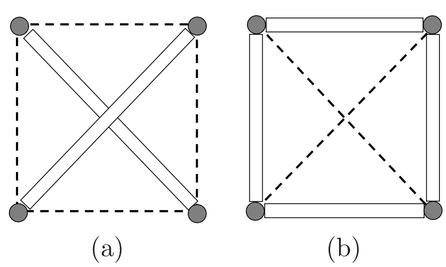

Figure 3.11: Rigid representations of two tensegrity frameworks. Neither of the two represented structures can thus be deformed. The two frameworks are complementary, i.e., can be obtained one from the other by replacing the struts by cables and the cables by struts.

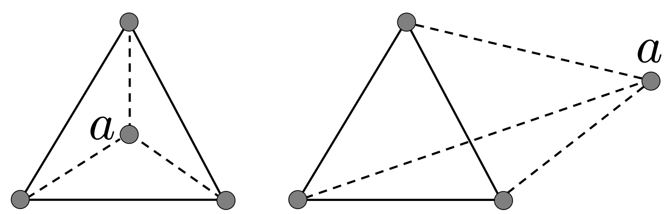

(a) (b)

Figure 3.12: Two representations of a same tensegrity framework. The representation in (a) is rigid, but the one in (b) is not. The position of a can indeed be moved to the left hand side without breaking any of the constraints implied by the three cables. Moreover, any representation "close" to (a) is rigid, and any representation "close" to (b) is not, which proves that rigidity is not a generic notion for tensegrity frameworks.

associated to a tensegrity framework $G(V,B,C,S)$ if there hold

$$\begin{array}{rclr} ||p_i - p_j|| & = & d_{ij} & \forall (i,j) \in B, \\ ||p_i - p_j|| & \leq & d_{ij} & \forall (i,j) \in C, \\ ||p_i - p_j|| & \geq & d_{ij} & \forall (i,j) \in S. \end{array} \tag{3.3}$$

Definition 3.5. *A representation p of a* tensegrity framework $G(V,B,C,S)$ is rigid *if there is a neighborhood of p in which all tensegrity-realizations of the distance set induced by p and G are congruent to p.*

A consequence of the introduction of cable and struts is that rigidity is not a generic notion for tensegrity frameworks. Figure 3.12 shows for example two representations of a same tensegrity framework, one of which is rigid and the other not, and both belonging to positive measure sets of representations having the same rigidity properties.

A notion of infinitesimal rigidity can also be defined for tensegrity frameworks similarly as in Section 3.2, but with a different definition of admissible displacements that takes the new types of constraints into account. It follows from the linearization of (3.3) that an infinitesimal displacement δp is admissible by a representation p of G if there hold

$$\begin{array}{rclr} (p_i - p_j)^t(\delta p_i - \delta p_j) & = & 0 & \forall (i,j) \in B, \\ (p_i - p_j)^t(\delta p_i - \delta p_j) & \leq & 0 & \forall (i,j) \in C, \\ (p_i - p_j)^t(\delta p_i - \delta p_j) & \geq & 0 & \forall (i,j) \in S. \end{array}$$

A representation of a tensegrity framework is then *infinitesimally rigid* if all its admissible infinitesimal displacements are Euclidean. Again, every infinitesimally rigid representation of tensegrity framework is rigid, and almost all rigid representations of a tensegrity framework are infinitesimally rigid. A notion equivalent to infinitesimal rigidity can also be defined using a force transmission approach, as in Section 3.5. A representation p of a tensegrity framework $G(V, B, S, C)$ is statically rigid if for any load $f \in Eu_p^{\perp}$, there exists a $\lambda \in \Re^{|E|}$ such that $-f = R_{G,p}^T \lambda$, with $\lambda_{ij} \geq 0$ for all $(i, j) \in S$ and $\lambda_{ij} \leq 0$ for all $(i, j) \in C$. Using algebra on cones, one can prove that static rigidity and infinitesimal rigidity are equivalent notions for tensegrity frameworks.

We finish this section by mentioning an interesting complementarity property. Let G be a tensegrity framework. We call *complement of* G the tensegrity framework $\bar{G}$ obtained by replacing all struts of G by cables, and all cables of G by struts. The two frameworks represented in Figure 3.11 are for example complement one of each other.

Proposition 3.3. *A representation p of a tensegrity framework G is infinitesimally rigid if and only if it is infinitesimally rigid as a representation of the complement framework $\bar{G}$.*

Proof. Observe that an infinitesimal δp displacement is admissible by p as a representation of G if and only if the opposite infinitesimal displacement $-\delta p$ is admissible by $\bar{G}$. Moreover, since Eu_p is a vector space, $\delta p \in Eu_p \Leftrightarrow -\delta p \in Eu_p$. Therefore, every infinitesimal displacement admissible by p and G is Euclidean if and only if every infinitesimal displacement admissible by p and $\bar{G}$ is Euclidean. $\square$

3.7 History and literature

Putting a precise date on the first apparition of rigidity is of course a hard, if not impossible, task. Early human beings leaving caves must indeed have had at least some intuitive idea of rigidity to build huts. Mentions of rigidity in scientific works go back to at least Euler who has conjectured in 1766 that "A closed spatial figure allows no changes, as long as it is not ripped apart" [47].

In the first part of the twentieth century, several rules for building rigid and isostatic frameworks have been developed, for example by Henneberg in 1911 [72] and Cox in 1936 [30]. According to Tay and Whiteley [124], some of the rules proposed were however vague or unproven, and even uncorrect [30].

A more rigorous mathematical study of rigidity has then been undertaken since the years 1970', following the famous theorem of Laman [84], which fully characterizes rigidity in two-dimensional space. These works were conducted by authors such as Asimow and Roth [6, 7], Tay [123, 124], Whiteley [129], Gluck [56], Graver et al. [57], Lovasz and Yemini [95], and Recski [109, 110].

The first appearances of formation shape maintenance by enforcement of distance constraints seems to go back to Desai et al. [37–39] and Lewis and Tan [85]. Further studies of the use of rigidity for autonomous agent formation were then proposed notably by Eren et al. [41–45], by Olfati-Saber and Murray [104, 105], and by Baillieul an Suri [11].

Tensegrity, presented in Section 3.6, appears to have been defined by Richard Buckminster Fuller in is second "Synergetics" book [54]. It was latter studied by Calladine [24], and by various authors such as Roth, Whiteley and Connely [29, 114].

Chapter 4

Persistence of Directed Graphs

We now consider directed graphs abstracting formations governed by unilateral distance constraints. Each vertex represents an agent, and a vertex i is connected to a vertex j by a directed edge (i, j) if the agent i has the responsibility of maintaining the distance between i and j constant. Our goal is to analyze the ability of a set of unilateral distance constraints to maintain the shape of the formation, as independently as possible from the particular control laws used to govern the agents. Remember that when control laws are used to satisfy the constraints, nothing guarantees a priori that all constraints will be satisfied. To analyze this, we need to make some basic assumptions on the agents' behavior, which we detail in Section 4.1. We then define persistence of a graph representation in Section 4.2 in the same way as in [64, 134]. We introduce infinitesimal persistence in Section 4.3. Both persistence and infinitesimal persistence are notions defined for representations of graphs and not for graphs, but the characterization we give in Section 4.4 shows that they are both generic notions for graphs. Finally, we introduce in Sections 4.5 and 4.6 the related notions of degrees of freedom and minimal persistence.

We would like to emphasize that that the concepts introduced in this chapter are *formally defined* for graphs or for their representations, and are *motivated* by formation issues and by the assumptions on the agents' behavior in Section 4.1. In particular, they can be used to characterize the situation of such formations at equilibrium. Whenever possible, we provide thus an interpretation of our graph-theoretical notions in terms of autonomous agent formations.

Before going further, it is also important to notice that two opposite directed edges (i, j) and (j, i) are not equivalent to one undirected edge. The latter represents a distance constraint whose responsibility is shared by two agents, which can thus collaborate in order to maintain the prescribed distance. Two opposite directed edges represent on the other hand two unilateral constraints on the same distance, and the agents are not assumed to collaborate, nor even to be aware of the fact that they are both responsible for the maintenance of the same distance. We do not consider in this thesis hybrid formations with both unilateral and bilateral constraints, and are not aware of any work in which this would have been done, although it could be an interesting research topic.

4.1 Agent behavior

Let i be an agent responsible for maintaining its distance to the agents $j_1, \dots, j_d$ constant. We suppose that the behavior of i is dictated by a control law [1]

$$\dot{p}_i = u_i \left(p_1, p_2, \dots, p_d\right),$$

where u should be invariant under rotation and translation of all agent positions, as we assume that the agents have no absolute or common notion of positions and directions. We say that an agent is *at equilibrium* when $\dot{p}_i = 0$. The position of an agent at equilibrium remains thus constant unless one of its neighbors moves. To represent the fact that the control laws are designed with the objective of satisfying the agent's distance constraints, we need to make basic assumptions on the conditions under which an agent is at equilibrium according to these control laws, and on the fact that they drive the agent to an equilibrium.

Assumption 4.1.
(a) All other agent positions being fixed, an agent reaches or asymptotically tends to equilibrium.
(b) An agent satisfying all its distance constraints is at equilibrium.
(c) An agent not satisfying all its constraints but for which there is a position at which it would satisfy all its constraints is not at equilibrium.

This assumption is a natural one as it just means that the control laws' goal is for the agent to satisfy all its constraints when it is possible. When the constraints assigned to an agents are compatible (i.e. simultaneously satisfiable), Assumptions 4.1(b) and 4.1(c) reduce thus to "*An agent is at equilibrium if and only if it satisfies all its constraints*". This is for example almost always the case

[1]For the sake of clarity, we use here a first order control law depending only on the present position of the neighbors. More elaborate control laws are likely to be use in practice, using for example second order models, and the history or the positions. Our reasoning remains valid for such control laws provided that the notion of equilibrium is appropriately redefined. One could for example say that i is at equilibrium is $u_i = 0$ and would remain so if all other agent positions were fixed.

when an agent has no more than D constraints. An agent in a D-dimensional space can indeed satisfy at least D constraints except in some degenerated cases (see Section 4.4 for more details). More problematic is the case of an agent for which the constraints assigned are incompatible, about the equilibrium condition of which Assumption 4.1 does not state anything. Incompatible constraints can appear in some degenerate cases and, more importantly, when more than D constraints are assigned to an agent. The latter case may not be ruled out as an unrealistic one, as redundances in structure are often used for robustness purposes. Moreover it may be possible for an agent to simultaneously satisfy more than D constraints, as represented for example in Figure 2.4. Following the idea initially proposed in [63, 64], we further assume that an agent unable to satisfy all its constraints satisfies then a maximal subset of compatible constraints[2]. Intuitively, this may be seen as the result of a greedy process, in which the agent tries to satisfy additional constraints as long as it is possible to do so without breaking those already satisfied. It is then at equilibrium if and only if it cannot satisfy any additional constraint without breaking one that it already satisfies. This behavior together with Assumption 4.1 fully characterize the agent equilibrium conditions, and can be reformulated as

Assumption 4.2.
(a) All other agent positions being fixed, an agent reaches or asymptotically tends to equilibrium.
(b) An agent is at equilibrium if and only if it satisfies a maximal subset of its constraints, i.e., there is no position at which it would satisfy more distance constraints while still satisfying those that it already satisfies.

Unlike Assumption 4.1, Assumption 4.2(b) is partly arbitrary and therefore debatable by definition. One could have also assumed for example that the agent minimizes some continuous cost function[3]. We partly explore the use of such different hypotheses in Section 7.1.2. Assumption 4.2 represents however well the fact that redundant constraints are added for the purpose of robustness with respect to the "loss" of a constraint, due for example to problems of communications or of measures. Note that it does not forbid the agent to choose the constraints that it ignores when its set of constraints becomes incompatible.

Naturally, a formation is at an equilibrium when all agents are at an equilibrium. In view of Assumption 4.2, a formation is thus at an equilibrium if no agent can satisfy any additional constraint by modifying its position without breaking a constraint that it already satisfies, considering the positions of the

[2] A maximal set of compatible constraints is a set of compatible constraints that is strictly contained in no other set of compatible constraints. It should not be confused with a maximum set of compatible constraints, which is a set containing the largest possible number of compatible constraints

[3] Such functions would however necessarily be non-convex if considered globally.

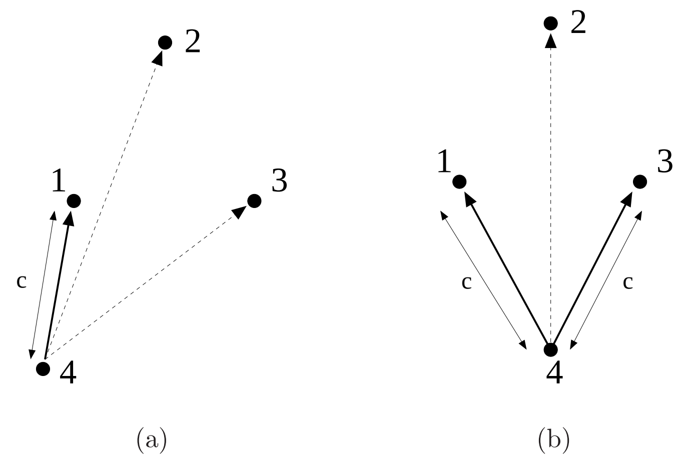

Figure 4.1: Suppose that the desired distances are $d_{41} = d_{42} = d_{43} = c$. Agent 4 is not at equilibrium in (a) because it only satisfies the constraint on the distance d_{41} while there exists a position where it could satisfy the distance constraints on both d_{41} and d_{43}. On the other hand, its position in (b) is an equilibrium one because no point can be at a distance $d_{42} = c$ of 2 in addition to being at a distance $d_{41} = d_{43} = c$ of 1 and 3.

other as fixed. The equilibrium of a formation can be interpreted as a Nash equilibrium where the agents are players trying to maximize the set of satisfied constraints. Figure 4.1 shows for example a formation at equilibrium and another one that is not at equilibrium.

Constraint consistence intuitively introduced in Chapter 2 characterizes the fact that all distance constraints are satisfied when each agent is *trying* to satisfy its constraints. Intuitively, the fact that a formation is at equilibrium means that all agents are "trying" to satisfy their constraints. A constraint consistent formation should thus be one for which all constraints are satisfied at equilibrium. Similarly, persistence characterize the fact that the formation shape is maintained when each agent is *trying* to satisfy its constraints. A persistent formation should thus be one for which the shape at equilibrium is (locally) unique. We formally define and analyze persistence and constraint consistence in the next sections. Remember however that these notions characterize the situation at equilibrium and not the dynamics of the agents. They are thus independent of the particular control laws used. On the other hand, they do not guarantee the possibility of reaching an equilibrium, which is a non trivial issue as there exist formations for which an equilibrium may never be reached or approached. We analyze this question further in Chapter 5.

4.2 Persistence and constraint consistence

Remember that the notions of *graph representation, distance set realization, congruence* introduced in Section 3.1 can also be applied to directed graphs, and so is thus rigidity. Let us now fix a directed graph G, desired distances $d_{ij} > 0$ for all $(i, j) \in E$, and a representation p. Consistently with Assumption 4.2, we have the following definition of equilibrium position and representation.

Definition 4.1. *Given a representation p, a vertex i is at an* equilibrium position[4] *p_i for a distance set $\bar{d}$ and a graph G if there is no $p^* \in \Re^D$ for which the following strict inclusion holds:*

$$\{(i, j) \in E : ||p_i - p_j|| = d_{ij}\} \subset \{(i, j) \in E : ||p^* - p_j|| = d_{ij}\}, \tag{4.1}$$

A representation is an equilibrium representation *(or a representation at equilibrium) for a certain distance set $\bar{d}$ and graph G if all the vertices are at equilibrium positions for $\bar{d}$ and for G.*

Note that any realization of a distance set is always an equilibrium representation of this distance set, as it satisfies by definition all distance constraints. As a consequence, any realizable distance set trivially admits equilibrium representations. But when the distance set is not realizable, the existence of an equilibrium representation is possible but not guaranteed. As a simple example, consider a cycle graph of length 3, and a distance set that does not satisfy the triangular inequality. The distance set is not realizable, so there is no representation simultaneously satisfying the three constraints. On the other hand, the vertices have an out-degree 1, so that there is always a p^* for which the inclusion (4.1) holds if the constraint corresponding to their out-going edge is not satisfied. There is thus no equilibrium representation. All distance sets considered in the sequel are however realizable as they are induced by graph representations, so that they always admit equilibrium representations. We can now formally define the concepts of constraint consistence and persistence.

Definition 4.2. *A representation p of a graph G is* constraint consistent *if there is a neighborhood of p in which every representation that is at equilibrium for the distance set induced by p and G is a realization of this distance set.*
A representation p of a graph G is persistent *if there is a neighborhood of p in which every representation that is at equilibrium for the distance set induced by p and G is congruent to p.*

The following theorem shows that we could have alternatively defined persistence as the intersection of rigidity and constraint consistence, where rigidity was introduced in Section 3.1 for both directed and undirected graphs.

[4] "equilibrium position" is preferred to the term "fitting position" used in [63] as it corresponds more to the interpretation of the notion given in Section 4.1

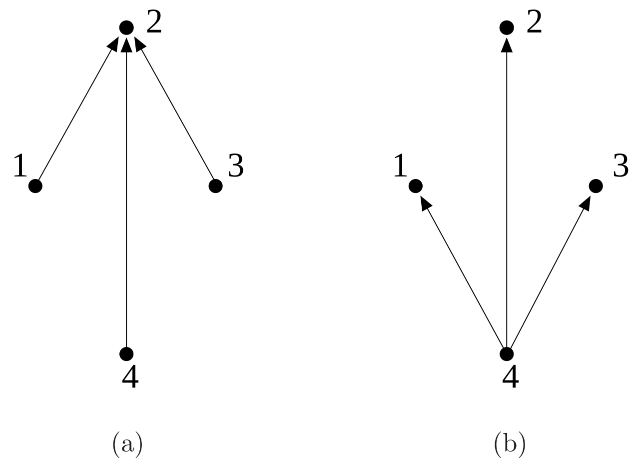

Figure 4.2: The formation represented in (a) is constraint consistent. Each of 1, 3 and 4 can indeed always satisfy its unique distance constraint. On the other hand, the representation (b) is not constraint consistent because there exists a configuration of positions of 1, 2 and 3 such that 4 is unable to satisfy its three distance constraints.

Theorem 4.1. *A representation of a directed graph is* persistent *if and only if it is* rigid *and* constraint consistent*.*

Proof. Let p be a rigid and constraint consistent representation, and call N_r and N_c the neighborhood of p coming from the application of the definitions of rigidity and constraint consistence to p. Take now a representation $p' \in N_r \cap N_c$ that is at equilibrium for the distance set induced by p. The constraint consistence of p implies that p' is a realization of this distance set, and the rigidity of p implies then that p and p' are congruent.

Conversely, consider now a persistent representation p of a graph G, its induced distance set $\bar{d}$, and the neighborhood N coming from the application of the definition of persistence to p. We show that this neighborhood is appropriate for both constraint consistence and rigidity. Since p is persistent, any representation $p' \in N$ at equilibrium for G and $\bar{d}$ is congruent to p, and is thus a realization of $\bar{d}$. The representation p is therefore rigid. Moreover, any realization $p' \in N$ of $\bar{d}$ by definition also an equilibrium representation for $\bar{d}$. The persistence of p implies then that it is congruent to p, which is therefore also rigid. $\square$

Note that this result formally holds for any definition of equilibrium representation, provided that every realization of a distance set is at equilibrium for this distance set. Before continuing, we believe important to insist on the fact that as rigidity, the notions of persistence and constraint consistence are defined with respect to local variations of positions. They can thus a priori not be applied when large deformations are considered. We see in Section 7.1.3 that this can cause practical problems for formations that are constraint consistent but not rigid, due to the possibility of having large deformations while satisfying all constraints. Persistent formations are however immune to those phenomena.

4.3 Infinitesimal persistence

As with rigidity and infinitesimal rigidity, it is convenient to use the notion of infinitesimal persistence, which corresponds to a first order analysis of persistence. We do not analyze formally the link between persistence and infinitesimal persistence, but we arrive in Section 4.4 at a characterization of infinitesimal persistence equivalent to the characterization of persistence in [64, 134], proving the equivalence of the two notions (up to a zero-measure set as will be seen).

Let us again consider a representation p of a graph G, and an infinitesimal displacement δp, which as in Section 3.2 is assumed to be sufficiently small so that $||\delta p||^2$ is negligible. Suppose that i is connected by directed edges to $j_1, j_2, \dots, j_{d_i^+}$. The distance constraints corresponding to the outgoing edges of i are

$$||(p_i + \delta p_i) - (p_{j_k} + \delta p_{j_k})||^2 = ||p_i - p_{j_k}||^2$$

for $k = j_1, j_2, \dots, j_{d_i^+}$. Due to our linearization hypothesis $||\delta p||^2 = 0$, this is equivalent to

$$\begin{array}{rcl} (p_i - p_{j_1})^T(\delta p_i - \delta p_{j_1}) & = & 0 \\ (p_i - p_{j_2})^T(\delta p_i - \delta p_{j_2}) & = & 0 \\ & \vdots & \\ (p_i - p_{j_{d_i^+}})^T(\delta p_i - \delta p_{j_{d_i^+}}) & = & 0. \end{array} \tag{4.2}$$

This system is actually the restriction of the system (3.2) $R_{G,p}\delta p = 0$ to the lines of the rigidity matrix corresponding the edges leaving i. It is therefore satisfied by any admissible δp.

From an autonomous agent point of view, the agent i has the responsibility for the satisfaction of the constraints (4.2), but can only act on its own displacement[5] δp_i. According to Assumption 4.2, it is thus at equilibrium if and only if it satisfies a maximal set of those constraints, considering all other agent displacements δp_j as fixed. The formation is then at equilibrium if no agent

[5]We assume here that i is not taking the potential reaction of other agents to its action into account.

could satisfy any additional constraint without breaking one that it already satisfies, considering the displacement of the other agents as fixed.

To formalize this idea, let us rewrite the system (4.2) in a way that emphasizes the role of δp_i,

$$\begin{pmatrix} (p_i - p_{j_1})^T \\ (p_i - p_{j_2})^T \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \end{pmatrix} \delta p_i = \begin{pmatrix} (p_i - p_{j_1})^T \delta p_{j_1} \\ (p_i - p_{j_2})^T \delta p_{j_2} \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \delta p_{j_{d_i^+}} \end{pmatrix}. \tag{4.3}$$

Observe that the matrix of the system (4.3) contains d_i^+ lines and D columns. For almost all representations, any subset of at most D lines is thus linearly independent. So if the right hand-side term is considered as fixed, then for almost all representations this system admits a unique solution when $d_i^+ = D$ and a non-unique one when $d_i^+ < D$. We say that *i is at equilibrium with δp_i* (with respect to p, G, and δp), if δp_i satisfies a maximal subsystem of (4.3), that is, if there is no $\delta p_i^* \in \Re^D$ satisfying all equations of (4.3) satisfied by δp_i and at least another one, *all other δp_j being fixed*. We say then that the infinitesimal displacement δp is an *equilibrium infinitesimal displacement* if every vertex $i \in V$ is at equilibrium with δp_i, and denote by $\mathrm{Equil}_{G,p}$ the set of equilibrium infinitesimal displacement for a representation p of a directed graph G. We can now formally define infinitesimal persistence and constraint consistence.

Definition 4.3. *A representation p of a directed graph G is infinitesimally constraint consistent if all its equilibrium infinitesimal displacements are admissible, that is if* $\mathrm{Equil}_{G,p} \subseteq \mathrm{Ker} R_{G,p}$ *holds.*
A representation p of a directed graph G is infinitesimally persistent if f all its equilibrium infinitesimal displacements are Euclidean, that is if $\mathrm{Equil}_{G,p} \subseteq Eu_p$ *holds.*

The structure of $\mathrm{Equil}_{G,p}$ is a priori unknown, but clearly every admissible infinitesimal displacement is in $\mathrm{Equil}_{G,p}$ as for such displacements all constraints are satisfied. We have thus the inclusion relation $Eu_p \subseteq \mathrm{Ker} R_{G,p} \subseteq \mathrm{Equil}_{G,p}$. The following theorem analogous to Theorem 4.1 follows then directly from the definitions of infinitesimal rigidity, constraint consistence, and persistence.

Theorem 4.2. *A representation of a directed graph is infinitesimally persistent if and only if it is infinitesimally rigid and infinitesimally constraint consistent.*

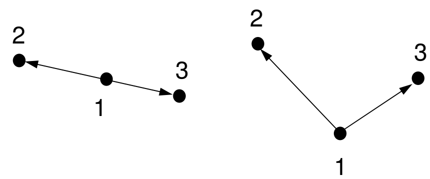

(a) (b)

Figure 4.3: The representation (a) is degenerate (in $\Re^2$) as the vectors $(p_2 - p_1)$ and $(p_3 - p_1)$ define a one-dimensional space. On the other hand (b) is non-degenerate as the same two vectors define a two dimensional-space.

4.4 Characterization and generic character

Infinitesimal constraint consistence and persistence are notions defined for representations of graphs and not for graphs. We now show that they are also generic notions and give a characterization of constraint consistent and persistent graphs. To establish this generic character, we need a notion of non-degenerate representation. A representation p of a directed graph G is *non-degenerate* if for any i and any subset $\{j_1, j_2, \ldots, j_{n'_i}\}$ of at most D of the vertices to which it is connected by directed edges, the collection of vectors $\left\{(p_{j_1} - p_i), (p_{j_2} - p_i), \ldots, (p_{j_{n'_i}} - p_i)\right\}$ spans a n'_i-dimensional space. Examples of degenerate and non-degenerate representations are shown in Figure 4.3. Clearly, the set of degenerate representations has zero measure.

Lemma 4.1. *Let p be a non-degenerate representation of a directed graph G and δp be an infinitesimal displacement. The vertex i is at equilibrium with δp_i if and only if the constraints corresponding to at least* $\min\left(D, d^+_{i,G}\right)$ *of its outgoing edges are satisfied. As a consequence, δp is an equilibrium infinitesimal displacement if and only if, for each i, at least* $\min\left(D, d^+_{i,G}\right)$ *constraints corresponding to edges leaving i are satisfied by δp.*

Proof. Consider a δp, one particular vertex i and its associated system (4.3). Since p is non-degenerate, every collection of at most D lines of the system matrix is linearly independent. As a consequence, every subsystem of at most D equations admits a solution, and this solution is unique if the subsystem contains exactly D equations.
Suppose first that $d^+_i \leq D$, that is, that the system associated to i contains no more than D equations. Then it admits a solution, so that i is at equilibrium if and only if it satisfies all equations. Suppose now that the system contains more than D equations and that δp_i satisfies S of them. If $S < D$, then there is a subsystem of D equations containing those already satisfied and admitting a solution, and i is not at equilibrium. On the other hand, if $S \geq D$,

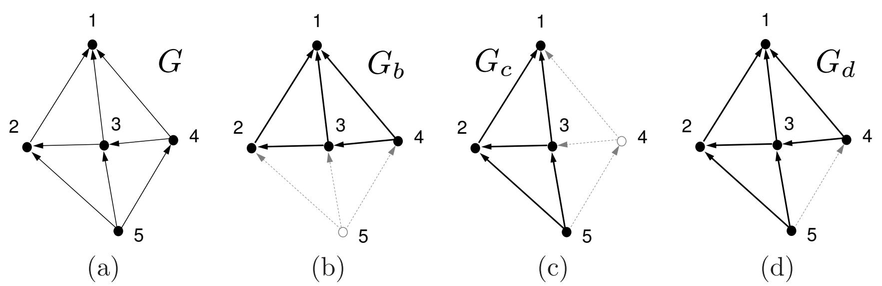

Figure 4.4: Examples of closed subgraphs, $\min D$-subgraphs and strict $\min D$-subgraphs (with $D = 2$). G_b is a closed subgraph of G as all vertices of G_a have the same out-degree as in G. G_c and G_d are strict $\min D$- subgraphs of G as all their vertices have an out-degree 2 or the same out-degree as in G when it is smaller than 2. Note that G_d is a proper subgraph although it contains all the vertices of G.

then the subsystem of satisfied equations admits a unique solution, so that no different δp_i^* could satisfy another equation in addition to the S already satisfied. i is thus at equilibrium (for p and δp) if and only if δp_i satisfies at least $\min\left(D, d^+_{i,G}\right)$ equations of (4.3). The first part of the result follows then from the fact that the equations of the system (4.3) for each i are equivalent to those in $R_{G,p}\delta p = 0$ restricted to the lines of the rigidity matrix corresponding to the outgoing edges of i. The second part of the result follows from the fact that an equilibrium infinitesimal displacement is one for which every i is at equilibrium. $\square$

The notion of closed subgraph of a directed graph is well known. A subgraph $G'(V', E')$ of $G(V, E)$ is a *closed subgraph* of G if all vertices of V' have the same outgoing edges in G' as in G. Equivalently, G' is a closed subgraph of G if all its vertices have the same out-degree in G' as in G. Motivated by Lemma 4.1, we introduce a more general class of subgraphs. A subgraph G' is a *a* $\min D$*-subgraph of* G if every vertex of G' has in G' an out-degree at least D or equal to its out-degree in G if the latter is smaller than D. In other words, if there holds $d^+_{i,G'} \geq \min\left(D, d^+_{i,G}\right)$. We say that such a subgraph is a *strict* $\min D$*-subgraph* if the condition is tied, that is if it contains no vertex with an out-degree larger than D. In a strict $\min D$-subgraph, there holds thus $d^+_{i,G'} = \min\left(D, d^+_{i,G}\right)$. Examples of such subgraphs are shown in Figure 4.4. It trivially follows from these definitions that every closed subgraph is a $\min D$-subgraph, although not necessarily a strict one.

Intuitively, observe that the information about other agent's position in a formation travels in the directions opposite to the edges. Agents in a closed subgraph $G' \subseteq G$ receive thus no information about agents in $G \setminus G'$ because no edge leaves G'. Besides, when an agent has more than D constraints, it may select D or more of them and ignore the others if they become incompatible. It then ignores the information travelling on the corresponding edges. So, in a $\min D$-subgraph $G' \subseteq G$, the agents may receive information about agents in $G \setminus G'$ but are not necessarily influenced by it, because all constraints corresponding to edges leaving G' can simultaneously be ignored. Even if the subgraph is not closed, the agents may thus behave independently of the rest of the formation. For this reason, we called such subgraphs "practically closed subgraphs" in [133, 134]. The term $\min D$-subgraph is preferred here for conciseness and readability reasons. Note that the selection of constraints described above is of course not necessarily unique, and could change with time.

Lemma 4.2. *Let p be a non-degenerate representation of a directed graph G. δp is an equilibrium infinitesimal displacement for p and G if and only if δp is admissible by p and at least one strict $\min D$-subgraph S of G containing all the vertices of G, that is, if there is such an S for which $\delta p \in Ker R_{S,p}$.*

Proof. Consider a graph G and a non-degenerate representation p of G in $\Re^D$. Suppose that δp is an equilibrium infinitesimal displacement. It follows from Lemma 4.1 that for each i, at least $\min\left(D, d^+_{i,G}\right)$ of the constraints corresponding to the outgoing edges of i are satisfied. Let us build a subgraph S of G by taking all the vertices of G and, for each vertex i, $\min\left(D, d^+_{i,G}\right)$ outgoing edges whose corresponding constraints are satisfied. S is by construction a strict $\min D$-subgraph G. Moreover δp is admissible by S since it satisfies the constraints corresponding to all edges of S.
To prove the reverse implication, suppose now that δp is admissible by p and consider a strict $\min D$-subgraph S on all vertices of G. For each i, δp satisfies at least $\min\left(D, d^+_{i,G}\right)$ constraints corresponding to edges leaving i, which by Lemma 4.1 implies that δp is an equilibrium infinitesimal displacement for p and G. □

We can now give a characterization of (infinitesimal) constraint consistence and persistence[6].

Theorem 4.3. *Let G be a directed graph. A non-degenerate representation p of G in $\Re^D$ is constraint consistent if and only if $rank R_{G,p} = rank R_{S,p}$ holds for every subgraph $\min -D$-subgraph S of G.*
As a consequence, p is infinitesimally persistent as representation of G if it is infinitesimally rigid as a representation of all $\min D$-subgraphs of G subgraphs in $\Sigma(G)$.

[6] In [64, 134] we only gave a characterization of persistence, which was a particular case of the characterization of constraint consistence. Different characterizations of constraint consistence have also been obtained independently by Jia Fang [50].

Proof. Let $\Sigma(G)$ be the set of $\min -D$-subgraphs of G. By Lemma 4.2, there holds $\text{Equil}_{G,p} = \bigcup_{S\in\Sigma(G)} KerR_{S,p}$. And by Definition 4.3, p is infinitesimally constraint consistent as a representation of G if and only if $\text{Equil}_{G,p} \subseteq KerR_{G,p}$ holds. It is thus infinitesimally constraint consistent if and only if $KerR_{S,p} \subseteq KerR_{G,p}$ holds for every $S \in \Sigma(G)$. Since $KerR_{G,p} \subseteq KerR_{S,p}$ holds trivially for every $S \subseteq G$, this condition is equivalent to $rankR_{G,p} = rankR_{S,p}$.
The result for infinitesimal persistence follows then from Theorem 4.2 and from the fact the infinitesimal rigidity only depends on the rank of the representation rigidity matrix. □

A consequence of this theorem is that a representation p of a directed graph G is infinitesimally constraint consistent if and only if $\text{Equil}_{G,p}$ is a vectorial space. More importantly, since the rank of the rigidity matrix is a generic notion, Theorem 4.3 shows that both infinitesimal constraint consistence and persistence are actually generic notions. We have then the following two theorems:

Theorem 4.4. *A graph is (infinitesimally) constraint consistent in $\Re^D$ if all its strict* min *D-subgraphs on all its vertices have the same generic rigidity matrix rank as itself.*

Theorem 4.5. *A graph is (infinitesimally) persistent in $\Re^D$ if all its strict* min *D-subgraphs on all its vertices are rigid.*

This last condition is exactly the one obtained in [64,134] for persistence, so that persistence and infinitesimal persistence are equivalent notions. Moreover, one could also prove that infinitesimal constraint consistence is equivalent to constraint consistence[7]. For the sake of conciseness, we therefore say in the sequel that a graph is *persistent* in $\Re^D$ if almost all its representations in $\Re^D$ are infinitesimally persistent and that a graph is *constraint consistent* if almost all its representations in $\Re^D$ are infinitesimally constraint consistent. These notions depend thus again on the particular dimension in which the graph is suppose to be represented. Moreover, we omit in the sequel the word "infinitesimal" except when it is essential.

Figure 4.5 shows an example of application of Theorem 4.5. The graph in Figure 4.5(a) is rigid but not persistent for $\Re^2$. Observe that only the vertex 3 has an out-degree larger than 2. The strict $\min D$-subgraphs on all vertices are thus obtained by taking two edges leaving 3 and all edges leaving other vertices. Since at least one such graph, represented in Figure 4.5(b), is not rigid, the graph in Figure 4.5(a) is not persistent. From an autonomous agent point of view, 3 has an out-degree 3 and can thus be led to "ignore" one of them.

[7] This is not theoretically difficult, but long. Moreover, it does not provide more insight on the notion than the result for infinitesimal constraint consistence.

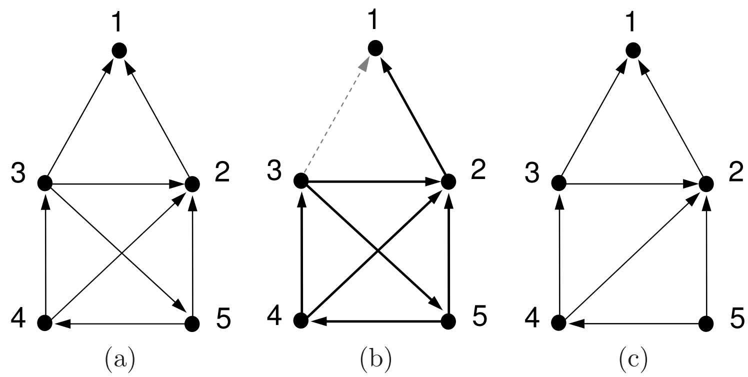

Figure 4.5: The graph (a) is rigid but not persistent for $\Re^2$ because one of its strict min D-subgraph on all its vertices is (b), which is not rigid. Besides, the graph (a) can be obtained by adding an edge to the provably persistent graph (c), which proves that persistence is not necessarily preserved by addition of edge.

It if ignores $(3, 1)$, the remaining graph is not rigid and thus not sufficient to maintain the formation shape. Observe also that the graph in Figure 4.5(a) is obtained from the provably persistent graph in Figure 4.5(c) by addition of only one edge. This shows that persistence is not necessarily preserved by addition of an edge, a major difference between rigidity and persistence. From an autonomous agent point of view, the addition of an edge introduces a redundant constraint. In the presence of too many redundant constraints, a combination of unfortunate selections among the various possible constraints architecture available to the agents may lead to one that does not preserve the shape, as the one in Figure 4.5(b).

The criterion of Theorems 4.3 and 4.5 does not immediately lead to a polynomial-time algorithm to check the persistence of a directed graph, as it requires checking the rigidity of a number of subgraphs that can grow exponentially with the graph size. The existence of a polynomial-time algorithm remains actually an open question (see Section 7.2.3) in the general case. Results can however be obtained for particular classes of directed graphs, as will be seen in Chapter 6.

The following corollaries are direct consequences of Theorems 4.4 and 4.5:

Corollary 4.1. *A graph is persistent if and only if it is rigid and constraint consistent.*

Corollary 4.2. *Every graph containing no vertex whose out-degree is larger than D is constraint consistent in $\Re^D$. As a consequence, such a graph is persistent in $\Re^D$ if and only if it is rigid.*

Corollary 4.3. *Let G be a graph that is persistent in $\Re^D$. A graph obtained from G by removing an edge leaving a vertex with an out-degree larger than D is also persistent in $\Re^D$.*

Proof. Let $G'(V, E')$ be a graph obtained from $G(V, E)$ by removing an edge leaving a vertex with an out-degree larger than D. For every $i \in V$, there holds $\min\left(D, d^+_{i,G'}\right) = \min\left(D, d^+_{i,G}\right)$. Therefore, every strict $\min D$-subgraph of G' is also a strict $\min D$-subgraph of G. The result follows then from Theorem 4.5, stating that persistence of a graph is equivalent to the rigidity of all its $\min D$-subgraphs. $\square$

4.5 Degrees of freedom

In a non-degenerate representation, for a vertex i to be at equilibrium, δp_i must either a solution of the system (4.3) or a solution of one of its subsystems containing at least D equations. Considering the other displacements δp_j as fixed, there is thus a unique δp_i for which it is at equilibrium if the out-degree d^+_i is D or more. But if $d^+_i < D$, the set of δp_i for which i would be at equilibrium is a $(D - d^+_i)$-affine variety. In the corresponding autonomous agent formation, this intuitively means that the agent i has in the latter situation some degrees of freedom or of decision in the choice of its displacement.

Definition 4.4. *The* number of degrees of freedom *(in $\Re^D$) of a vertex in a directed graph is the generic dimension of the set of its possible equilibrium displacement (in $\Re^D$), considering the other displacements as fixed.*

The following proposition allowing the computation of the number of degrees of freedom follows directly from the discussion above.

Proposition 4.1. *Let $G(V, E)$ be a directed graph. The number $dof_{i,G}$ of degrees of freedom of a vertex $i \in V$ in $\Re^D$ is $D - \min\left(D, d^+_i\right)$.*

We now show that the total number of degrees of freedom in a persistent graph is bounded by the number of independent rotations and translations in a D-dimensional space.

Proposition 4.2. *The sum of the degrees of freedom over all vertices of a persistent graph is at most $f_D = \frac{1}{2}D(D+1)$.*

Proof. Let $G(V, E)$ be a directed graph and $S(V, E_S)$ be a strict $\min D$-subgraph subgraph of G on all its vertices. Vertices having no degree of freedom in G

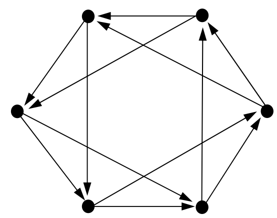

Figure 4.6: Example of persistent graph in $\Re^2$ in which no vertex has any degree of freedom.

have an out-degree D in S. A vertex i having $dof_{i,G} > 0$ degrees of in G has an out-degree $d^+_{i,S} = D - dof_{i,G}$ in S. The number of edges in S is thus

$$E_S = \sum_{i \in V} d^+_{i,S} = \sum_{i \in V} (D - dof_{i,G}) = nD - dof(S),$$

where $dof(S)$ is the sum over all vertices of S of their number of degrees of freedom. It follows from the persistence of G and Theorem 4.5 that S is rigid and contains therefore at least $nD - f_D$ edges. There holds thus $dof(S) \leq f_D$. □

There exist however graphs for which the total number of degrees of freedom is smaller than f_D, even if they contain more than $D-1$ vertices. See for example the graph in Figure 4.6. This may seem paradoxical as there is generically f_D degrees of freedom when representing a rigid graph in $\Re^D$, or when assigning a Euclidean displacement to this graph. To understand this apparent paradox, remember that the number of degrees of freedom that we have defined characterizes for each agent in a formation the dimension of the space in which it can choose its equilibrium position/displacement, *considering the other agents as fixed*. It concerns therefore the displacements resulting from the decision of a single agent, and not from a collective decision[8]. In Figure 4.6 for example, a collective decision could allow a 3-dimensional space of displacement, although no agent alone has the power to make the formation move. Collective decisions are however not considered in the persistence framework and are beyond the scope of this thesis.

We now show that, except if it belongs to some closed subgraph on less than D vertices, each vertex is connected by a directed path to all vertices having degrees of freedom. Intuitively, remember that in the autonomous agent formation, the information travels in the directions opposite to the edges. This

[8]For this reason, "degree of decision" would probably be a more appropriate terminology than "degree of freedom". However, we prefer here to use the latter phrasing so as to be consistent with previously published work [64, 69, 134].

means that the information about the decisions taken by the agent with degrees of freedom can reach every agent in the formation, except some agent having a sufficient freedom.

To prove this result, we introduce the notion of redundance number of a rigid graph for a particular dimension. The *redundance number* $r_D(G)$ of a graph G for $\Re^D$ is defined by $r_D(G) = |E| - nD + f_D$ if $n \geq D$ and by $r_D(G) = 0$ if $n < D$. For a rigid graph, it follows from Theorem 3.1 that the redundance number corresponds to the maximal number of edges that can be removed without breaking rigidity. More generally it is always possible to remove $r_D(G)$ edges to a graph without affecting the generic rank of its rigidity matrix, as this rank is never larger than $nD - f_D$ when $n \geq D$.

Lemma 4.3. *Let $G^*(V^*, E^*)$ be a subgraph of a rigid graph $G(V, E)$. There holds $r_D(G^*) \leq r_D(G)$.*

Proof. If G^* or G contains less than D vertices, the result is trivial. Otherwise, the rigidity of G together with Theorem 3.1 implies the existence of a subgraph $G'(V, E')$ of G on the same vertices such that $|E'| = D|V| - f_D$, and for any subgraph $G''(V'', E'')$ of which there holds $|E''| \leq D|V''| - f_D$. In particular, considering the subgraph $G' \cap G^*$, there holds $|E' \cap E^*| \leq D|V \cap V^*| - f_D = D|V^*| - f_D$. Moreover, it follows from the definition of $r_D(G)$ that G' is obtained from G by removing $r_D(G)$ edges. Therefore the intersection of G' with $G^* \subseteq G$ is obtained from G^* by removing at most $r_D(G)$ edges, and we have

$$|E^* \cap E'| \geq |E^*| - r_D(G).$$

Since $|E^* \cap E'|$ is no greater than $D|V^*| - f_D$, this leads to

$$|E^*| \leq D|V^*| - f_D + r_D(G),$$

and thus by definition of $r_D(G^*)$ to $r_D(G^*) \leq r_D(G)$. □

Theorem 4.6. *Let $G(V, E)$ be a persistent graph for $\Re^D$. If $i \in V$ does not belong to any closed subgraph on less than D vertices, there are directed paths starting at i and reaching all vertices with positive number of degree of freedom. As a consequence, every closed subgraph on at least D vertices contains all vertices with positive number of degrees of freedom.*

Proof. Let V_i be the set of vertices in G that can be reached from i via a directed path. Let then $F = \sum_{j \in V} dof_G(j)$ and $F_i = \sum_{j \in V_i} dof_G(j)$ be the sum over all vertices of respectively V and V_i of their number of degrees of freedom in G or equivalently in S. Obviously, $F_i \leq F$. We now show that $F_i \geq F$ also holds, so that the number of degrees of freedom in the set of vertices reachable from i is the same as the number of degrees of freedom in the whole graph, which implies the desired result. The relation between the different graphs and

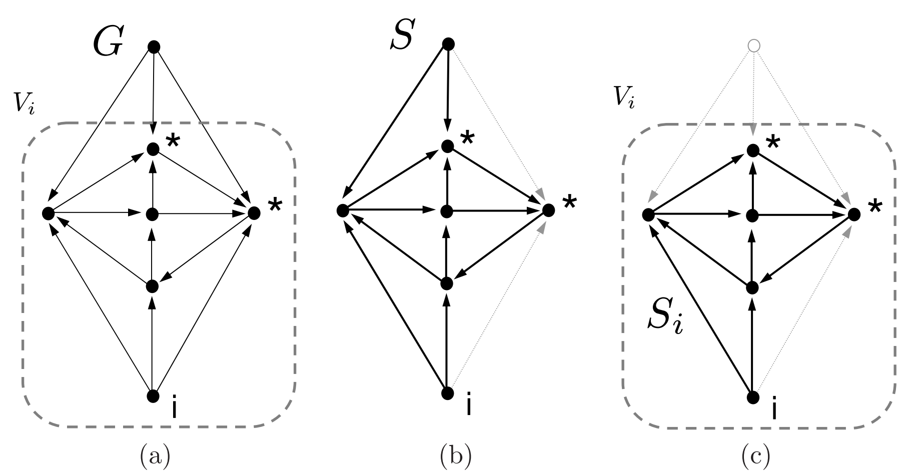

Figure 4.7: Illustration of different graphs and sets appearing in the proof of Theorem 4.6. In (a), V_i is the set of vertices that can be reached from a certain vertex i in a graph G. A strict min 2-subgraph S of G is shown in (b), and its restriction S_i to the vertices of V_i in (c). The symbol "*" represents one degree of freedom.

sets introduced in the sequel is illustrated by an example in Figure 4.7.

Consider a strict min D-subgraph $S(V, E_S) \subseteq G$ on all the vertices of G. It follows from Theorem 4.5 that S is rigid. Moreover, the definition of number of degrees of freedom implies that the vertices have the same number of degrees of freedom in S as in G. The number of edges in S is then $\sum_{j\in V} d_S^+(j) = \sum_{j\in V}(D - dof_S(j)) = |V|\,D - F$, and the redundance number of S is $r_D(S_i) = D\,|V| - F - |V|\,D + f_D = f_D - F$

Let now S_i be the restriction of S to the vertices of V_i. By construction every edge of $E_S \subseteq E$ leaving a vertex of V_i arrives at a vertex of V_i. The vertices of V_i have thus the same out-degree and number of degrees of freedom in S_i as in S, and their number of degrees of freedom is also the same as in G. The total number of edges in S_i is thus $|E_{Si}| = \sum_{j\in V_i} d_{S_i}^+(j) = \sum_{j\in V}(D - dof_{S_i}(j)) = |V|\,D - F_i$. By assumption, $|V_i| \geq D$. Therefore the redundance number of S_i is $r_D(S_i) = D\,|V_i| - F_i - |V_i|\,D + f_D = f_D - F_i$. It follows then from Lemma 4.3 that $F_i \geq F$, which as explained above proves our result. $\square$

Since a vertex with no degree of freedom has at least D neighbors and is therefore never in a closed subgraph of less than D vertices, this implies that every vertex without degree of freedom is connected by directed paths to all vertices with positive number of degrees of freedom. Moreover, it follows from

Theorem 4.5 that every min D-subgraph on all the vertices of a persistent graph is persistent. In each such subgraph, a vertex with at least D neighbors is thus connected to every vertex with positive number of degrees of freedom in the initial graph, as the degrees of freedom in the initial graph are also present in the min D-subgraph. From a formation point of view, remember that the agents may "ignore" the information travelling on some edges as long as the edges that are not ignored constitute a min D-subgraph. This means thus that the agents receive the information about the decisions made by those with degrees of freedom for any allowed choice of ignored constraints.

4.6 Minimal persistence

We say that a graph is *minimally persistent* in $\Re^D$ if it is persistent in $\Re^D$ and if no graph obtained from it by removing one or several of its edges is persistent in $\Re^D$. Clearly, every persistent graph contains a minimally persistent subgraph on all its vertices. The converse is however not true, as adding an edge to a (minimally) persistent graph does not necessarily lead to a persistent graph. The graph in Figure 4.5(a) contains for example the (minimally) persistent graph of Figure 4.5(c) as subgraph.

Theorem 4.7. *Let $G(V,E)$ be a directed graph. The following conditions are equivalent:*
(a) G is minimally persistent;
(b) G is minimally rigid and its largest out-degree is at most D;
(c) G is persistent and minimally rigid;
(d) G is persistent and contains $D\,|V| - f_D$ edges if $|V| \geq D$ and $\frac{1}{2}\,|V|\,(|V|-1)$ edges if $|V| < D$;
(e) G is persistent and contains a minimal number of edges, i.e., every persistent graph on $|V|$ vertices has at least $|E|$ edges.

Proof. (a) $\Rightarrow$ (b): It follows from Corollary 4.3 that a minimally persistent graph has no vertex with an out-degree larger than D, for otherwise one could obtain a smaller persistent graph by removing an edge leaving such a vertex. Moreover, if such a graph is not minimally rigid, there is at least one edge whose removal gives a rigid graph, which by Corollary 4.2 is also persistent.

(b) $\Rightarrow$ (c): Corollary 4.2 implies that every rigid graph graph with no out-degree larger than D is persistent.

(c) $\Rightarrow$ (a): If G is minimally rigid, the removal of any of its edges leads to a non-rigid graph, which by Corollary 4.1 is therefore not persistent. So if G is also persistent, is minimally persistent.

(c) $\Leftrightarrow$ (d): Suppose that G is persistent and therefore rigid. If $|V| \geq D$, it follows from Theorem 3.1 that G is minimally rigid if and only if $|E| = D|V| - f_D$ holds. If $|V| < D$, the result follows from the fact G is then rigid if and only if every pair of vertices is connected by an edge.

(d) $\Rightarrow$ (e): Every persistent graph is rigid, and it follows from Theorem 3.1 that a rigid graph contains at least $D|V| - f_D$ edges if $|V| \geq D$ and $\frac{1}{2}|V|(|V|-1)$ edges else. Therefore, if G satisfies the condition (d), no graph having the same number of vertices but less edges than G is persistent.

(e) $\Rightarrow$ (a): If G is persistent and no graph having the same number of vertices and less edges is persistent, then the removal of any edge of G leads clearly to a non-persistent graph, so that G is minimally persistent. □

This theorem implies that persistence can be checked in linear time when the graph is minimally rigid (b $\Leftrightarrow$ c), and that minimal persistence can be checked in linear time when the graph is persistent (a $\Leftrightarrow$ d). The following corollary follows immediately from Theorem 4.7 and Proposition 4.1

Corollary 4.4. *The total number of degrees of freedom in a minimally persistent graph on D or more vertices is always f_D.*

4.7 Double edges

To close this chapter, we would like to remove the possible ambiguity that could arise in the presence of so-called "double-edges". Since the graphs are directed, we may have a case in which i is linked to j by a directed edge (i,j), and j to i by another directed edge (j,i). As explained in Section 4.1, these two directed edges representing two unilateral distance constraints on the same distance are not equivalent to one undirected edge representing a bilateral distance constraint. In all the results about persistence and rigidity, they are considered as two different edges, and the fact that they connect the same pair of vertices is not taken into account. For rigidity however, such edges are exactly equivalent to one single edge, as the algebraic expression of their corresponding constraints are indeed identical. One of them can thus always be removed without losing rigidity. The following results shows that one of them can always be removed from a persistent graph without losing persistence. This means that such edges are never needed to ensure persistence. Moreover it has been argued in [11] that "double-edges" may cause some instabilities when some class of control laws is used.

Proposition 4.3. *Let i,j be two vertices of a persistent graph $G(V,E)$ such that $(i,j) \in E$ and $(j,i) \in E$. Then at least one of the two graphs obtained from G by removing (i,j) or (j,i) is persistent.*

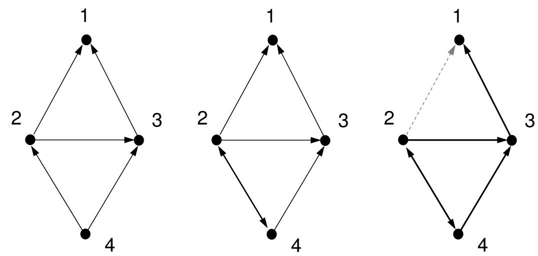

(a) (b) (c)

Figure 4.8: The graph represented in (a) is (minimally) persistent in $\Re^2$, as can be seen using Theorem 4.7. The graph represented in (b), obtained by adding one additional edge $(2, 4)$ between agents which were already connected by the edge $(4, 2)$ in the opposite direction, is not persistent in $\Re^2$. It contains indeed the non-rigid graph (c) as strict min D-subgraph.

Proof. If one vertex among i and j has an out-degree larger than D, the result follows directly from Lemma 4.3. If both have an out-degree no greater than D, then both (i, j) and (j, i) belong to all strict min D-subgraph of G on all its vertices. The removal of anyone of those two edge preserves therefore the rigidity of all these subgraphs. By Theorem 4.5, it also preserves the persistence of the graph. $\square$

The converse implication is however not true as shown by the example in Figure 4.8. Another consequence of Proposition 4.3 is that a minimally persistent graph never contains double-edges.

4.8 History and literature

The main results in this Chapter appeared first in [63, 64] for two-dimensional spaces and were then extended to higher dimensional spaces in [67,133,134] (except for Theorem 4.6 which was first published in [67, 133, 134]). The difference between the way they are obtained here and in those publications comes from the use of infinitesimal displacements to define persistence and constraint consistence. This allows indeed simpler proofs using linear algebraic tools. Besides, the proof of Theorem 4.6 and Lemma 4.3 on which it is based are new.

Chapter 5

Structural Persistence

5.1 Convergence to equilibrium

Persistence and constraint consistence defined in Chapter 4 characterize formations at equilibrium. They do not however take the convergence to this equilibrium into account. Indeed, although Assumption 4.1 states that every agent's control law leads it to an equilibrium *if all other agents are fixed*, it does not say anything about convergence of the entire formation. Neglecting the convergence of the formation could have dramatic consequences. Consider for example the graph represented in Figure 5.1(a). It follows from Corollary 4.2 that this graph is persistent in $\Re^3$. Observe however that agents 1 and 2 are not responsible for any distance constraint, and can thus freely choose their positions or displacements. Even if the distance between the positions they choose is different from the distance separating them in the reference representation as in Figure 5.1(b), they are at equilibrium and there is no reason for them to change their position. But in such a situation, the shape of the formation is obviously not preserved.

This apparent paradox is not in contradiction with the persistence of the graph. Agents 3, 4 and 5 are indeed responsible for three constraints each. Considering the positions of the four other agents as fixed, each of them is able to satisfy all its constraints. As a result, each of these agents is at equilibrium if and only if it satisfies its three constraints. But despite the ability of each agent to satisfy individually all its constraints, it may not be possible for 3, 4 and 5 to *simultaneously* satisfy all their constraints, so that they are *never simultaneously at equilibrium*. More formally, a representation in which the distance between 1 and 2 is not the same as in the reference representation is never an equilibrium representation, because there always is at least one vertex that has not an equilibrium position. The non-congruence of such a representation to

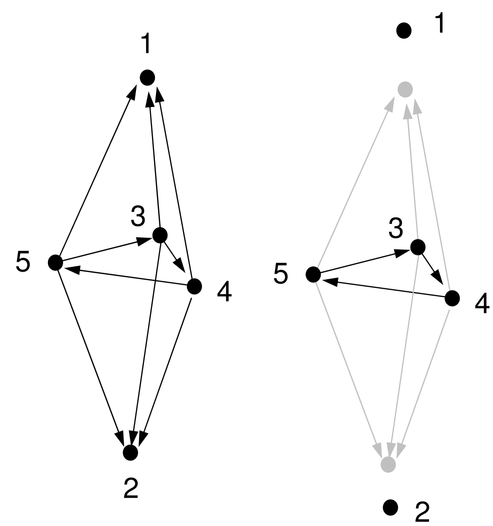

(a) (b)

Figure 5.1: (a) Persistent representation (in $\Re^3$) of a graph for which the convergence to an equilibrium is not guaranteed. 1 and 2 have indeed no constraint and can therefore freely choose their positions. They can then keep their positions independently of the positions of 3, 4 and 5. If, as in (b), the distance between the position they choose is not the same as in the reference representation, it becomes impossible for 3, 4 and 5 to simultaneously satisfy all their constraints, although each of them could satisfy its three constraints if the other were fixed. As a consequence, they never converge to equilibrium. No such example is possible in the plane.

the initial representation does thus not contradict the persistence or constraint consistence of the graph.

The example of Figure 5.1 shows the importance of taking the convergence issue into account when analyzing the ability of a structure of unilateral distance constraints to preserve a formation shape. A complete characterization should provide a guarantee that all agents converge to equilibrium. And a necessary condition for this is the absence of invariant set (in the position domain) whose closure contains no equilibrium representation. More formally, let $p(t) \in \Re^{nD}$ be the vector containing all agent positions $p_i(t) \in \Re^D$ at time t, One should make sure that there exists no $S \subseteq \Re^{nD}$ such that

- If $p(t) \in S$, then $p(t') \in S$ for all $t' \geq t$.
- No $p \in \bar{S}$ is an equilibrium representation.

In the example of Figure 5.1, the set $S := \{p_1 = p_1^*; p_2 = p_2^*; p_3, p_4, p_5 \in \Re^D\}$ is such a forbidden set for almost all $p_1^*, p_2^* \in \Re^D$, as agents 1 and 2 keep their position independently of the other agent's positions.

There is no full characterization available yet of the set of directed graphs for which each invariant set's closure contains at least an equilibrium representation. It is however possible to characterize the directed graphs ruling out a subset of those problematic invariant sets: those in which some of the agents positions or displacements are fixed while other are totally free. This leads to the notion of structural persistence presented below and introduced in [134]. No problematic invariant set has been found yet that cannot be treated using structural persistence and the graph of Figure 5.1 is for example not structurally persistent. Moreover, we will show below that every graph persistent in $\Re^2$ is structurally persistent, and a graph persistent in $\Re^3$ is structurally persistent if and only if it does not contain two vertices with three degrees of freedom each.

5.2 Structural persistence

Let p be a representation of a graph $G(V, E)$. We call *partial (infinitesimal) displacement* a subset of vertices $V_c \subseteq V$ together with an (infinitesimal) displacement δp_i for each $i \in V_c$. We denote it by δp_{V_c}. A partial displacement is thus the restriction of a displacement to a subset of the vertices. For a partial displacement δp_{V_c}, any displacement δp whose restriction to V_c is δp_{V_c} is called a *completion* of δp_{V_c}. Finally, we say that a partial displacement δp_{V_c} is a partial equilibrium (for p and G) if, for any completion δp of δp_{V_c} (i.e. for any assignment of displacements to the vertices of $V \setminus V_c$), every $i \in V_c$ is at equilibrium with its δp_i. In the a representation of the graph of Figure

5.1 for example, any assignment $\delta p_1, \delta p_2$ is a partial equilibrium. 1 and 2 have indeed an out-degree 0, and are thus always at equilibrium independently of their displacement and of those of the other vertices.

Definition 5.1. *A representation p of a graph G is* structurally persistent (respectively constraint consistent) *if it is persistent (respectively constraint consistent) and if every partial equilibrium can be completed to obtain an equilibrium displacement, that is an displacement for which every vertex is at equilibrium.*

Intuitively, a representation that is constraint consistent but not *structurally* constraint consistent is one in which

a) Some agents in a subset V_c could take positions such that they are at equilibrium independently of the positions of the agents in $V \setminus V_c$, and have thus no reason to move;

b) Due to the positions of the agents in V_c, it becomes impossible for the other agents to be all simultaneously at equilibrium.

This situation corresponds thus to a invariant set defined by fixing some agent positions while keeping other free, and that contains no equilibrium position for the whole formation.

5.3 Characterizing structural persistence

Although we have also defined structural constraint consistence, our analysis in the sequel focusses solely on structural persistence. We refer the reader to Section 7.1.3 for more information on constraint consistence in the absence of rigidity.

We have intuitively argued that the agents in a $\min D$-subgraph can ignore all information coming from outside the subgraph. Provided that they satisfy sufficiently many constraints inside the subgraph, they can thus be at equilibrium independently of the agents out of the subgraph. The following lemma establishes this formally, and also proves that being in a $\min D$-subgraph is necessary for the agents to behave independently.

Lemma 5.1. *Let p be a non-degenerate representation of a directed graph $G(V,E)$, and $V_c \subseteq V$ a subset of the vertices of G. A partial displacement δp_{V_c} on $V_c \subseteq V$ is a partial equilibrium for G and p if and only if it is admissible by at least one $\min D$-subgraph $G_c(V_c, E_c)$ of G on all the vertices of V_c and the restriction of p to it.*

Proof. Let p be a non-degenerate representation of G and $G_c(V_c, E_c)$ be a $\min D$-subgraph of G. Consider an displacement δp_{V_c} admissible by G_c, i.e., satisfying all constraints corresponding to edges of G_c. Since G_c is a $\min D$-subgraph of G, the out-degree in G_c of every vertex $i \in V_c$ is at least $\min(D, d^+_{i,G})$.

Therefore, for any completion δp of δp_{V_c}, at least $\min(D, d^+_{i,G})$ constraints corresponding to edges leaving i in G are satisfied for each $i \in V_c$. It follows then from Lemma 4.1 that every $i \in V_c$ is at equilibrium with its δp_i, so that δp_{V_c} is a partial equilibrium.

To prove the converse implication, consider now a partial equilibrium δp_{V_c} on the vertices of V_c. It follows from Lemma 4.1 that for any $i \in V_p$, at least $\min(D, d^+_{i,G_c})$ constraints corresponding to edges leaving i are satisfied, independently of the δp_j assigned to vertices j of $V \setminus V_c$. Remember that the constraint of (4.3) corresponding to an edge (i, j) is $(p_i - p_j)^T \delta p_i = (p_i - p_j)^T \delta p_j$. In a non-degenerate representation, there is for any δp_i a δp_j such that the constraint is not satisfied. Similarly, for a group of constraints corresponding to edges $(i, j_1), (i, j_2), \ldots (i, j_{d_i^+})$ leaving i and for any δp_i, it is possible to choose $\delta p_{j_1}, \delta p_{j_2}, \ldots \delta p_{j_{d_i^+}}$ in such a way that none (or any selection) of the constraints is satisfied, provided again that the representation is non-degenerate. No constraint corresponding to an edge arriving in $V \setminus V_c$ can thus be satisfied *for all completions δp of δp_{V_c}*. This implies that at least $\min(D, d^+_{i,G})$ of the edges leaving each $i \in V_c$ arrive at another vertex of V_c, and are satisfied by the partial displacement δp_{V_c}. Consider then the graph $G_c(V_c, E_c)$ obtained by taking V_c and all edges of E connecting two vertices of V_c and whose corresponding constraint is satisfied by δp_{V_c}. δp_{V_c} is clearly admissible by G_c and the restriction of p to V_c. Moreover, each vertex of V_c has in G_c at least an out-degree $\min(D, d^+_{i,G})$, so that G_c is a $\min D$-subgraph of G. □

Lemma 5.2. *A partial displacement is Euclidean if and only if can be completed to obtain a Euclidean displacement of all vertices. As a consequence, the restriction of a Euclidean displacement to a subset of vertices is Euclidean.*

Proof. Let p be a representation of a graph $G(V, E)$. A displacement is Euclidean if and only if there is a continuous Euclidean transformation $E(t, x)$ of $\Re^D$ (that is, a combination of rotations and translations of $\Re^D$), such that $\delta p_i = K \frac{dE(x,t)}{dt}|_{p_i,0}$ holds for all $i \in V$ for a same K. As a consequence, the restriction of a Euclidean displacement to a subset of vertices is Euclidean. Moreover, a Euclidean partial displacement can always be completed to obtain a (full) Euclidean displacement by assigning $\delta p_i = K \frac{dE(x,t)}{dt}|_{p_i,0}$ to all remaining vertices. □

Theorem 5.1. *Let p be a non-degenerate representation of a directed graph $G(V, E)$. The representation p is structurally persistent if and only if its restriction to every strict* $\min D$*-subgraph of G is rigid as a representation of this subgraph.*

Proof. We prove the result supposing that p is persistent as a representation of G. The complete results follows then directly from Theorem 4.3.

Suppose first that p is not structurally persistent but is persistent. There exists then a partial equilibrium displacement δp_{V_c} for which no completion is an equilibrium displacement for p and G. The admissibility of δp_{V_c} by a $\min D$-subgraph G_c of G and the restriction of p to G_c follows from Lemma 5.1. Moreover, δp_{V_c} is not a Euclidean displacement, for otherwise it could by Lemma 5.2 be completed to obtain a Euclidean displacement of all vertices, which would trivially also be an equilibrium displacement for all vertices. The restriction of p to G_c is therefore not rigid as a representation of G_c because it admits a non-Euclidean displacement δp_V. By removing some edges to G_c one can then obtain a strict $\min D$-subgraph of G, and the restriction of p to V_c is not rigid as a representation of this subgraph either, for otherwise it would be rigid as a representation of G_c.

Conversely, suppose now that there exists a strict $\min D$-subgraph G_c of G such that the restriction of p to G_b is not rigid as a representation of G_c. G_b and the restriction of p to it admits then a non-Euclidean displacement. It follows from Lemma 5.2 that this δp_{V_c} cannot be completed to obtain a Euclidean displacement. Therefore, none of its completions is an equilibrium displacement for p and G. The persistence of p as a representation of G implies indeed that every equilibrium displacement is Euclidean. And since it follows from Lemma 5.1 that δp_{V_c} is a partial equilibrium for p and G, this implies that p is not structurally persistent as a representation of G. $\square$

This theorem shows that structural persistence is a generic notion. The following characterization of *structurally persistent graphs* follows then directly.

Theorem 5.2. *A graph is structurally persistent if and only if each of its strict* $\min D$*-subgraphs is rigid.*

Observe the interesting parallelism between Theorem 5.2 and Theorem 4.5 characterizing persistent graphs. Checking the structural persistence of a graph is equivalent to checking the rigidity of all its strict $\min D$-subgraph, while checking the persistence of a graph is equivalent to checking the rigidity of a subclass only of these subgraphs, those containing all the vertices of the initial graph. From an intuitive point of view, the condition of Theorem 4.5 means that for any selection of constraints to be ignored by the agents, the remaining constraints need to be sufficient to maintain the formation shape. Theorem 5.2 adds the condition that if agents in a set V_c can ignore the other agents, then the constraints on distance between agents in V_c should be sufficient to maintain the shape of the sub-formation containing the agents of V_c.

5.4 Toward a simpler characterization

The necessary and sufficient condition of Theorem 5.2 involves detecting and checking a number of subgraphs that can grow exponentially with the size of the graph. It does therefore not lead to a polynomial time algorithm checking structural persistence. We now establish a simpler characterization of structural persistence for persistent graphs, leading to an algorithm requiring $O(D^2n)$ operations. This algorithm can not be used to test persistence, but it does detect some particular sort of non-persistent graphs.

Proposition 5.1. *Let G be a graph that is persistent in $\Re^D$. Then every strict* $\min D$*-subgraph on at least D vertices is rigid.*

Proof. Let $G(V,E)$ be a persistent graph and $G^*(V^*,E^*)$ one of its strict $\min D$-subgraph on at least D vertices. Let us build $S(V,E_S)$ by adding to G^* all vertices of $V\setminus V^*$, and for each of them, D of its outgoing edges or all of them if its out-degree is smaller than D. Observe that every edge of E_s leaving a vertex of V^* belongs to E^* and arrives at a vertex of V^*.

S is by construction a strict $\min D$-subgraph of G on all its vertices. The persistence of G and Theorem 4.5 imply then that S is rigid, and thus that there exists a minimally rigid subgraph $S'(V,E_{S'})$ on all its vertices. Let $S^*(V^*,E_{S^*})$ be the restriction of S' to the vertices V^* of G^* (obtained by removing edges that are incident to one or two vertices of $V\setminus V^*$). Since S' is minimally rigid and $|V^*|\geq D$, it follows from Proposition 3.1 that proving the relation $|E_{S^*}|\geq D\,|V^*|-f_D$ is sufficient to establish the rigidity of S^* and therefore also the rigidity of G.

The minimal rigidity of S' implies that $|E_{S'}|=D\,|V|-f_D$. Besides, every edge of $E_{S'}\setminus E_{S^*}$ is incident to at least one vertex of $V\setminus V^*$ because S^* is the restriction of S' to V^*. Since we know that every edge of $S'\subseteq S$ leaving a vertex of V^* arrives at a vertex of V^*, this implies that all edges of $E_{S'}\setminus E_{S^*}$ leave vertices of $V\setminus V^*$. $|E_{S'}\setminus E_{S^*}|$ is then upper-bounded by $\sum_{i\in V\setminus V^*} d^+_{i,S'}\leq D\,|V\setminus V^*|$, where we have used the fact the all out-degrees in S and a fortiori in S' are bounded by D. This implies then

$$|E_{S^*}|=|E_S|-|E_{S'}\setminus E_{S^*}|\geq D\,|V|-f_D-D\,|V\setminus V^*|=|V^*|-f_D.$$

□

This result leads to a stronger characterization of structural persistence, which we use to show that problems related to partial equilibrium never appear in $\Re^2$, and are very simply detected in $\Re^3$.

Corollary 5.1. *A graph that is persistent in $\Re^D$ is structurally persistent in $\Re^D$ if and only if all its closed subgraphs on less than D vertices are rigid.*

Proof. It follows from Proposition 5.1 and Theorem 4.5 that a persistent graph is structurally persistent if and only if all its strict min D-subgraph on less then D vertices are rigid.
The result follows then from the fact that any min D-subgraph on less than D vertices is always a closed subgraph. A min D-subgraph subgraph that is not a closed subgraph contains indeed at least a vertex with an out-degree at least D, and contains thus at least $D+1$ vertices. □

Corollary 5.2. *All graphs persistent in $\Re^2$ are structurally persistent. A graph persistent in $\Re^3$ is structurally persistent if and only if at most one of its vertices has three degrees of freedom, that is, an out-degree 0.*

Proof. A graph containing only one vertex is trivially always rigid. The result for $\Re^2$ follows then directly from Corollary 5.1.
The same corollary implies that a graph persistent in $\Re^3$ is structurally persistent if and only if it contains no non-rigid closed subgraph on less than 3 vertices. The only non-rigid graph on less than 3 vertices is the empty graph on 2 vertices. And, one can verify that it appears as a closed subgraph of a directed graph if an only if the latter graph contains 2 vertices with out-degree 0. Vertices have indeed the same out-degree in a closed subgraph as they have in the graph. □

Note that persistent but non-structurally persistent graphs may have more complicated structure in higher dimensions. One can for example easily build graphs in which the problematic closed subgraph is connected. We now achieve the characterization by considering the closed subgraph on less than D vertices.

Proposition 5.2. *Let G be a directed graph persistent in $\Re^D$, and χ the set of its closed subgraphs C_i on less than D vertices. The union of these subgraphs $C^* = \bigcup_{C_i \in \chi} C_i$ is a closed subgraph on less than D vertices.*

Proof. We prove that the union of two graphs of χ is also in χ, which due to the finite size of χ implies that the union of all graphs in χ is itself in χ. Consider $C_1(V_1, E_1), C_2(V_2, E_2) \in \chi$. Clearly, $C_1 \cup C_2$ is also a closed subgraph of G. Suppose now to obtain a contradiction that it contains D or more vertices. Then by Proposition 5.1 it should be rigid as $C_1 \cup C_2$ is a strict min D-subgraph of G. All vertices in C_1 and C_2 have indeed the same out-degree in G, C_1, C_2 and $C_1 \cup C_2$, and this out degree is smaller than D.

Let $x = |V_1 \cap V_2|$, $y = |V_1 \setminus (V_1 \cap V_2)| = |V_1| - x$ and $z = |V_2 \setminus (V_1 \cap V_2)| = |V_2| - x$. We have $|V_1 \cup V_2| = x + y + z \geq D$. Since it is rigid and contains at least D vertices, $C_1 \cup C_2$ has at least $D(x+y+z) - \frac{1}{2}D(D+1)$ edges (counting pairs of "double edges" as one single edge). On the other hand, the maximal number of edges (counting again pairs of "double edges" as one single edge) that $C_1 \cup C_2$ can contain is $\frac{1}{2}(x+y+z)(x+y+z-1) - yz$, obtained by connecting every pair of vertices except those composed by one vertex of $V_1 \setminus V_2$ and one of $V_2 \setminus V_1$. We now prove that there holds

$$\frac{1}{2}(x+y+z)(x+y+z-1) - yz < D(x+y+z) - \frac{1}{2}D(D+1),$$

which implies that $C_1 \cup C_2$ does not have enough edges to be rigid, contradicting our hypothesis. The inequality above is equivalent to

$$(x+y+z)^2 - (x+y+z) - 2yz < 2D(x+y+z) - D^2 - D,$$

which can be rewritten as

$$(x+y+z-D)^2 - (x+y+z-D) < 2yz. \tag{5.1}$$

Remember now that $x+y+z$ is no smaller than D. Moreover, since $x+y$ and $x+z$ are smaller than D, there holds $(x+y+z-D)^2 < yz$, so that the inequality (5.1) holds. As explained above, this proves our result. □

Theorem 5.3. *Let $G(V,E)$ be a persistent (in $\Re^D$) directed graph, and $V_c \subseteq V$ the set of vertices that belong to closed subgraphs of G on less than D vertices. G is structurally persistent if and only if its restriction to V_c is a complete graph.*

Proof. Let χ be the set of closed subgraphs C_i of G on less than D vertices. Note that the restriction of G to V_c is the union of all graphs in χ. By Corollary 5.1, G is structurally persistent if and only if every $C_i \in \chi$ is rigid. And by Proposition 3.2, graphs on less than D vertices are rigid if and only if there are complete graphs.
It follows from Proposition 5.2 that the union of the graphs in χ is itself in χ. The structural persistence of G implies thus that it is a complete graph. On the other hand, since each C_i is a closed subgraph of G, if the restriction of G to the union of the C_i is a complete graph then its restriction to any of the C_i is also a complete graph, which is trivially rigid. As a result, G is structurally persistent. □

The condition of this latter Theorem can be checked in linear time. Observe first that a vertex belongs to a closed subgraph on less than D vertices if and only if at most $D-2$ vertices can be reached from it by a directed path. Checking if a vertex belongs to V_c is thus done by a graph exploration limited to D vertices, which requires at most $O(D^2)$ operations. As a result, identifying V_c requires at most $O(nD^2)$ operations. Verifying whether the restriction of the graph to V_c is a complete graph requires then at most $O(D^2)$ operations since we know that V_c contains less then D vertices. The whole algorithm requires thus at most $O(D^2n)$ operations. Note that if the set V_c found contains more than $D-1$ vertices, it follows from Proposition 5.2 that the graph is not persistent. This algorithm could thus be used to filter some non-persistent graphs.

5.5 History and literature

The main results presented in this Chapter 5 were first published in [67, 133, 134]. Our presentation here is however different. First, we use infinitesimal displacements. Second, structural persistence is introduced in the more general framework of the formation convergence to an equilibrium. Besides, the graph-theoretical proofs in Section 5.4 are further simplified.

Chapter 6

Particular Classes of Graphs

Our characterization of persistence in Theorem 4.5 does in general not lead to a polynomial-time algorithm to check if a given graph is persistent, as it requires checking the rigidity of a potentially exponential number of subgraphs. As a related issue, it is non-trivial to see if a graphs remains persistent after some minor transformations such as the addition or removal of vertices and/or edges. We provide in this chapter stronger results on these issues for some particular classes of graphs. In Section 6.1, we show that the persistence of acyclic graphs can be efficiently checked, and that simple transformations preserve their persistence and their acyclicity. We have seen in Section 4.6 that (minimal) persistence can be checked in polynomial time for minimally rigid graph. We focus in Section 6.2 on two-dimensional minimally persistent graphs, and show how such graphs can be transformed one into another using simple operations, in such a way that all intermediate graphs are minimally persistent. Finally, we present in Section 6.3 a polynomial-time algorithm to check persistence of graphs having three degrees of freedom in a two-dimensional space.

6.1 Acyclic graphs

6.1.1 Introduction and relevance

A (directed) graph is said to be *acyclic* if it contains no path starting and arriving at the same vertex. Acyclic persistent graphs are an important subclass of persistent graphs, as the corresponding formations are easy to control and have good stability properties. Consider for example a two-dimensional formation corresponding to the graph of Figure 6.1, and observe that a move of an agent

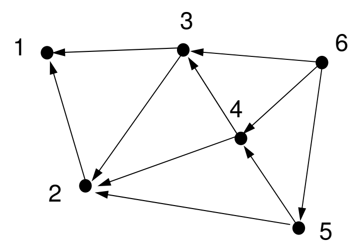

Figure 6.1: Example of acyclic graph with its topological sort. Besides, the graph is persistent in $\Re^2$.

can only influence the agents with higher labels. Intuitively, the agent 1 freely chooses its position. It is never influenced by other agents and is thus always at equilibrium. Agent 2 chooses its position on a circle centered on agent 1, and is then also at equilibrium since it can only be influenced by 1 which is already at equilibrium. The agent 3 can then uniquely determine its position, followed by 4, 5 and 6. Each of them remains at equilibrium once it has reached it. Note that the out-degree 3 of the sixth agent is not a problem, as once the five other agents are at equilibrium, their relative positions guarantee that the three constraints of 6 are compatible.

Some initial works on formation control by Baillieul and Suri only considered acyclic formation [11]. Recent studies have however shown that formations involving cycles could be stabilized and lead to equilibrium [3, 132]. Although we recognize the important stability advantage of acyclic formations, we believe that non-acyclic graphs are also worth studying, not only for theoretical completeness, but also because they could be useful for some applications. The absence of cycle removes indeed all possibility of feedback, guaranteeing the system stability. Feedback could however be desirable for some formations with large number of agents, in order to control the accumulation of small errors into large deformation.

Suppose for example that UAV's are required to fly around a circular surface without entering it, with the additional constraint that the distance between the UAV's should remain constant or at least not vary too much. This can theoretically be done using an acyclic topology of constraints such as the one presented in Figure 6.2(a). In the real world however, there would much likely be small errors on all measurements and relative positions of the agents. Even if these errors are insignificant for one single agent, their accumulation could cause the last agent to be at a position significantly different from its desired position, as represented in Figure 6.2(b). In particular, the distance between 1

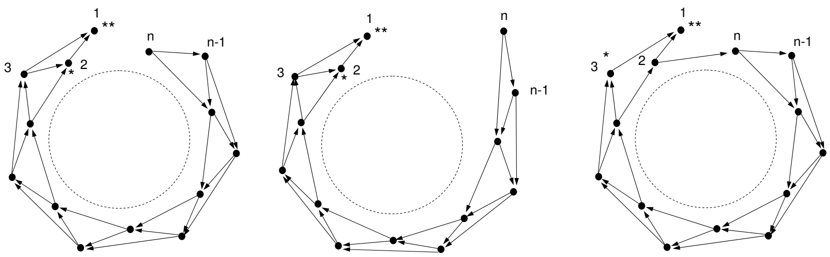

(a) (b) (c)

Figure 6.2: (a) Example of acyclic formation in which small errors can accumulate to produce large deformations, such as represented in (b). This phenomenon could be avoided by introducing cycles to created feedback, such as in (c). The symbol "*" represents one degree of freedom.

and n could be significantly larger than what is acceptable for the application. This could especially be the case if the angles made by the vectors joining each agent i to the agents $i-1$ and $i-2$ are small, as a small variation in the perception by i of the positions of $i-1$ and $i-2$ can then cause a proportionally large variation in the position of i. One solution to avoid this phenomenon could be to introduce a feedback by asking 1 and 2 to remain at a constant distance from n, as represented in Figure 6.2(c), and so introducing a cycle. In that case, two other constraints could be removed without losing persistence.

When analyzing acyclic graphs, it is convenient to use the notion of topological sort.

Definition 6.1. *Let G be a directed graph. A labelling $l : V \to \{1, \ldots, |V|\}$ of the vertices is a* topological sort *if for every edge (i, j) there holds $l(i) > l(j)$.*

An example of topological sort is shown in Figure 6.1. The following Proposition is a well known result on acyclic graphs. For a proof and more details on acyclic graphs, we refer the reader for example to [31, 80].

Proposition 6.1. *A directed graph is acyclic if and only if it admits a topological sort. As a consequence, $d^+_{i,G} < i$ and $d^-_{i,G} < n - i$ hold for each i in an acyclic graph where the vertices have been relabelled according to a topological sort, with in particular $d^+_{1,G} = 0$ and $d^-_{n,G} = 0$.*

Note that topological sort is in general not unique. In the sequel, we suppose that the vertices are always labelled by integers according to a topological sort. For every edge (i, j), there holds thus $i > j$.

6.1.2 Rigidity and persistence of acyclic graphs

To analyze the persistence of acyclic graphs, we need some results on addition and deletion of vertices

Proposition 6.2. *A minimally rigid graph $G(V,E)$ remains minimally rigid by addition of a vertex with degree $\min(D,|V|)$ or by deletion of a vertex with degree $\min(D,|V|-1)$.*

Proof. If $|V| \leq D$, the result follows from the equivalence, for graphs on no more than D vertices, between rigidity and being a complete graph (see Section 3.4). We now consider graphs with $|V| > D$, and first treat the case of the deletion.

The minimal rigidity of G implies that $|E| = D\,|V| - f_D$ holds. Let $G'(V',E')$ be the graph obtained from G by removing the vertex with degree D. There holds $|E'| = |E| - D = D(|V|-1) - f_D = D\,|V'| - f_D$. It follows from Proposition 3.1 that this relation together with the fact that G' is a subgraph of G on more than D vertices imply the minimal rigidity of G'.

Suppose now that $G'(V',E')$ is a graph obtained from G by adding a vertex $n+1$ connected by D edges to vertices of V. There holds again $|E'| = D\,|V'| - f_D$, so that the minimal rigidity of G' is equivalent to the (generic) independence of its edges. Let p' be a generic representation of G' and p its restriction to G. The minimal rigidity of G implies that the $|E| = D\,|V| - f_D$ lines of the rigidity matrix $R_{G,p}$ of G are linearly independent. The rigidity matrix is then

$$R_{G',p'} = \left(\begin{array}{c|c} R_{G,p} & 0 \\ \hline & (p'_{n+1} - p_{j_1})^T \\ X & \vdots \\ & (p'_{n+1} - p_{j_D})^T \end{array}\right),$$

which clearly contains only independent lines as the $D \times D$ sub-matrix on the lower right-hand side is (generically) nonsingular. □

Corollary 6.1. *Let $G(V,E)$ be a rigid graph. The graphs obtained by addition of a vertex connected to at least $\min(D,|V|)$ of the vertices of G or by deletion of a vertex with a degree at most $\min(D,|V|-1)$ are rigid.*

Proof. Let S be a minimally rigid subgraph of G, and i be a vertex of G having a degree no greater than $\min(D,|V|-1)$ in G and a fortiori also in S. Let then G' and S' be the graphs obtained from respectively G and S by the removal of i. S' is a subgraph of G. Moreover, it follows from Proposition 6.2 that it is minimally rigid, so that G' is rigid.

Let now G'' be a graph obtained from G by the addition of a vertex j connected to at least $\min(D,|V|)$ vertices, and S'' the graph obtained from S by adding j to S and connecting it to $\min(D,|V|)$ of the vertices to which it is connected in G''. S'' is by construction a subgraph of G''. And, it follows from Proposition 6.2 that it is minimally rigid, so that G' is rigid. □

Proposition 6.3. *Let $G(V,E)$ be an acyclic graph, and G' be a graph obtained from G by adding a vertex with in-degree 0 and connecting it to at least $\min(D,|V|)$ vertices of V. G' is persistent if and only if G is persistent.*

Proof. Let $\Sigma(G)$ be the set of all strict $\min D$-subgraph of G on all its vertices (that is, the set of all subgraphs that can be obtained from G by removing edges leaving vertices with out-degree larger than D until all vertices have an out-degree no greater than D), and $\Sigma(G')$ the corresponding set for G'. Call i the vertex added to G to obtain G'. In any graph $S' \in \Sigma(G')$, i has an in-degree 0 and an out-degree $d^+_{i,S'} = \max(D, d^+_{i,G'}) = \min(D,|V|)$. Observe that any graph in $\Sigma(G)$ can be obtained by removing i from one graph in $\Sigma(G')$. Similarly, any graph of $\Sigma(G')$ can be obtained by adding i and $\min(D,|V|)$ outgoing edges to a graph in $\Sigma(G)$.

Suppose first that G is persistent and therefore (by Theorem 4.5) that all graphs in $\Sigma(G)$ are rigid. Take a graph $S' \in \Sigma(G')$ and call S a graph in $\Sigma(G)$ from which S' can be obtained by adding i and connecting it to $\min(D,|V|-1)$ vertices of V. It follows from the rigidity of S and from Corollary 6.1 that S' is rigid, and therefore that G' is persistent since this is true for any $S' \in \Sigma(G')$. Suppose now that G' is persistent and therefore that all graphs of $\Sigma(G')$ are rigid. Take a graph $S \in \Sigma(G)$, and a graph $S' \in \Sigma(G')$ from which S can be obtained by removing i. It follows again from the rigidity of S and from Corollary 6.1 that S' is rigid, and therefore that G is persistent since this is true for any $S \in \Sigma(G)$. □

Using these results, we can now give a full characterization of persistence for acyclic graphs, constructing them by adding vertices one by one.

Theorem 6.1. *Let $G(V,E)$ be an acyclic graph whose vertices are labelled according to a topological sort. G is persistent if and only if any vertex $i \leq D$ has an out-degree $i-1$ and any other vertex has an out-degree no smaller than D.*

Proof. The result clearly holds if $|V| = 1$. To prove it inductively for other values of $|V|$, we suppose that it holds for $|V| = n$, and prove that it then holds for $|V| = n+1$. Note that $d^+_{i,G} < i$ and $d^-_{i,G} < |V| - i$ follows from Proposition 6.1 and the acyclicity of G.

Let G' be a graph of $n+1$ vertices, and G the graph obtained by the removal of the vertex $n+1$, which has an in-degree 0. Observe that if $d^+_{n+1} < \min(D,n)$ holds, then G' is not rigid nor persistent by Corollary 3.1, and does also clearly not satisfy the condition of this theorem. We now suppose that $d^+_{n+1} \geq \min(D,n)$. The graph G' can in that case be obtained from G by adding a vertex with an out-degree at least $\min(D,n) = \min(D,|V|)$ and an in-degree 0. It follows then from Proposition 6.3 that G' is persistent if and only if G is persistent, and our induction hypothesis implies that G is persistent if and only if it satisfies the condition of the theorem. As a consequence, and since $d^+_{n+1} \geq \min(D, n)$, G' is persistent if and only if it satisfies the condition of this theorem. □

Note that the necessary part of the theorem above could also easily be obtained by verifying that the total number of degrees of freedom in the graph is not larger than f_D.

It follows from Theorem 6.1 that all acyclic persistent graphs can be built by starting from a single vertex and sequentially adding vertices with $\min(D, |V| - 1)$ outgoing edges. In this process, all intermediate graphs are persistent and acyclic. Besides, the condition of this theorem can be checked in $O(E)$ and even in $O(V)$ if the out-degrees are known, providing a linear time algorithm to check persistence for acyclic graph. Note that acyclicity can be checked in $O(E)$. We can also particularize Theorem 6.1 to minimally persistent graphs.

Corollary 6.2. *An acyclic graph $G(V, E)$ is minimally persistent if and only if there exists an ordering $1, \dots, |V|$ of the vertices such that any vertex $i \leq D$ has an out-degree $i - 1$ and every other vertex (if any) has an out-degree D.*

Proof. An acyclic graph satisfying the condition of this corollary is persistent by Theorem 6.1. Moreover, the removal of any one or several of its edges diminishes the out-degree of one of its vertices, so that the graph obtained does not satisfy anymore the condition of Theorem 6.1 and is therefore not persistent. The initial graph is thus minimally persistent. Conversely, if an acyclic persistent graph does not satisfy the condition of this Corollary, it contains by Theorem 6.1 a vertex with an out-degree larger than D. The same theorem implies then that the graph obtained by removing an edge leaving this vertex is also persistent, preventing the initial graph from being minimally persistent. □

In recent works, Baillieul and McCoy [10] analyze the possibility of counting the number of two-dimensional acyclic minimally persistent graphs. They obtain an algorithmic procedure for carrying the enumeration based on a result presented in Corollary 6.2, but a general closed-form solution remains to be found. We now prove another consequence of Theorem 6.1, that all persistent acyclic graphs are structurally persistent. This is in agreement with the intuition that equilibrium can always be reached sequentially, starting with the first vertex and finishing with the last one according to a topological sort.

Corollary 6.3. *All persistent acyclic graphs are structurally persistent.*

Proof. Let G be an acyclic persistent graph labelled according to a topological sort. All vertices (if any) with a label $i \geq D$ have an out-degree at least $D - 1$ and do therefore not belong to any closed subgraph on less than D vertices.

Every other vertex i has by Theorem 6.1 an out-degree $i-1$, so that the restriction of G to these vertices or to any subset of them is a complete graph. The structural persistence of G follows then from Theorem 5.3. □

6.2 2-dimensional minimally persistent Graphs

6.2.1 Introduction and relevance

Minimally persistent graphs have a minimal number of edges and minimize thus the number of distances to be measured in the corresponding formation. Such formations are however not robust to the loss of communication, since the loss of one edge in a minimally persistent graph always results in a loss of persistence. Minimal persistence also implies the absence of redundant information, a property which has been argued to be necessary for the formation stability when some types of control laws are used [9, 11]. More complex control laws can however deal with redundant constraints. In view of Assumption 4.2 stating that an agent should satisfy a maximal subset of constraints, one can for example suppose that an agent facing incompatible constraints just temporarily discard some of them. Another solution is provided by the use of "death-zones" [51].

In this section, we study the construction and transformation of two-dimensional persistent graphs. Analogously to the powerful results about Henneberg sequences for minimally rigid graphs, we propose different types of directed graph operations allowing one to sequentially build any minimally persistent graph or to obtain it from any other minimally persistent graph, each intermediate graph being also minimally persistent.

Before starting, remember that a *minimally rigid* graph is a rigid graph such that no edge can be removed without losing rigidity. Particularizing Laman's Theorem (3.2) to minimally rigid graphs, we obtain the following criterion:

Proposition 6.4. *A graph $G = (V, E)$ with $|V| > 1$ is minimally rigid in $\Re^2$ if and only if the two following conditions hold*
a)$|E| = 2\,|V| - 3$,
b) for all $E'' \subseteq E, E'' \neq \varnothing$, there holds $|E''| \leq 2\,|V(E'')| - 3$.

A *minimally persistent graph* is a persistent graph such that no edge can be removed without losing persistence. The following Proposition is obtained by applying Theorem 4.7 to $\Re^2$.

Proposition 6.5. *A graph is minimally persistent in $\Re^2$ if and only if it is minimally rigid in $\Re^2$ and no vertex has an out-degree larger than 2.*

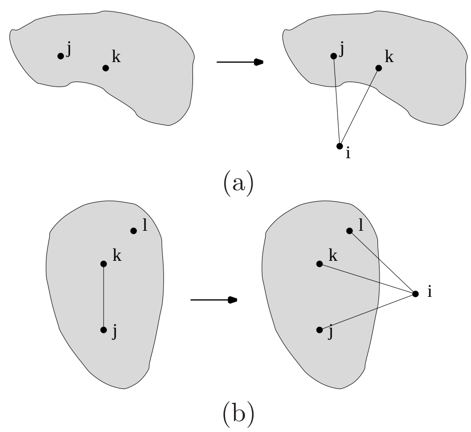

Figure 6.3: Representation of (a) the undirected vertex addition operation and (b) the edge splitting operation.

It follows from Proposition 4.1 that a vertex i has in $\Re^2$ $\max(0, 2 - d^+(i))$ degrees of freedom. In a minimally persistent graph, Proposition 6.5 implies then that this number is equal to $2 - d^+(i)$. Moreover, the total number of degrees of freedom in the graph is exactly 3, provided that the graph has at least 2 vertices. Note that all graphs considered in this section are supposed to be represented in $\Re^2$, although we do not always explicitly repeat it.

6.2.2 Henneberg sequences for minimally rigid graphs

Let j, k be two distinct vertices of a minimally rigid graph $G = (V, E)$. A *vertex addition* operation consists in adding a vertex i, and connecting it to j and k, as shown in Figure 6.3(a). It follows from Proposition 6.2 or 6.4 that this operation preserves minimal rigidity. The same proposition implies that if a vertex has a degree 2 in a minimally rigid graph, one can always perform the inverse vertex addition operation by removing it (and its incident edges) and obtain a smaller minimally rigid graph.

Let j, k, l be three vertices of a minimally rigid graph such that there is an edge between j and k. An *edge splitting* operation consists in removing this edge, adding a vertex i and connecting it to j, k and l, as shown in Figure 6.3(b). This operation provably preserves minimal rigidity [124]. Consider now a vertex i connected to three vertices j, k and l. A reverse edge splitting consists in removing i and adding one edge among (j, k), (k, l) and (l, j), *in such a way that the graph obtained is minimally rigid*. This operation can be performed on every vertex with degree 3 in a minimally rigid graph [84, 124],

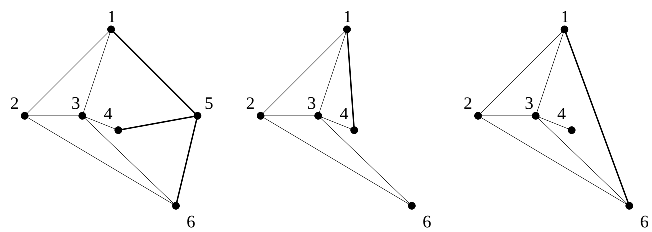

Figure 6.4: Example of unfortunate added edge selection in reverse edge splitting. After the removal of the vertex 5 from the minimally rigid graph (a), minimal rigidity can be preserved by the addition of the edge $(1,4)$ but not of $(1,6)$, as shown respectively on (b) and (c). The pair $(1,6)$ defines an implicit edge in the minimally rigid subgraph induced by 1, 2, 3 and 6.

but one cannot freely choose the edge to be added as shown on the example in Figure 6.4.

A *Henneberg sequence* is a sequence of graphs $G_2, G_3, \ldots, G_{|V|}$ with G_2 being the complete graph on two vertices K_2 and each graph G_i $(i \geq 3)$ can be obtained from G_{i-1} by either a vertex addition operation or an edge splitting operation. Since these operations preserve minimal rigidity and since K_2 is minimally rigid in $\Re^2$, every graph in such a sequence is minimally rigid in $\Re^2$.

A simple degree counting argument based on Proposition 6.4 shows that every minimally rigid graph $G_{|V|} = (V,E)$ with more than 2 vertices contains at least one vertex with degree 2 or 3. One can thus always perform either a reverse vertex addition or a reverse edge splitting operation and obtain a smaller minimally rigid graph $G_{|V|-1}$. Doing this recursively, one eventually obtains a minimally rigid graph on two vertices, which can only be K_2. It is straightforward to see that the sequence $K_2 = G_2, G_3, \ldots G_{|V|}$ is then a Henneberg sequence. We have thus proved the following result [124]:

Theorem 6.2. *A graph on more than one vertex is minimally rigid in $\Re^2$ if and only if it can be obtained as the result of a Henneberg sequence.*

Henneberg operations have been partly extended to minimally rigid graphs in three dimensions. Three minimal rigidity-preserving operations exist that add vertices with degree 3, 4 and 5 respectively. The reverse versions of the first two operations allow the removal of any vertex with a degree 3 or 4 while preserving minimal rigidity. But, no such operation is known yet to remove a vertex with a degree 5 (or more) and to always preserve minimal rigidity. Since there are three-dimensional minimally rigid graphs where all vertices have a degree at least 5, the argument proving Theorem 6.2 can thus not be generalized. The existence of a three-dimensional equivalent to Theorem 6.2 actually remains an open question. As a consequence there is no obvious way to generalize the result that we present in this section to higher dimensions. For more information on three-dimensional operations, we refer the reader to [124].

6.2.3 Natural extension of the Henneberg operations to directed graphs

Let j, k be two distinct vertices of a minimally persistent graph $G = (V, E)$. A *directed vertex addition* [44, 63] consists in adding a vertex i and two directed edges (i, j) and (i, k) as shown in Figure 6.5(a). A reverse (directed) vertex addition consists in removing a vertex with an out-degree 2 and an in-degree 0 from a minimally persistent graph.

Let now (j, k) be a directed edge in a minimally persistent graph and l a distinct vertex. A *directed edge splitting* [44, 63] consists in adding a vertex i, an edge (i, l), and replacing the edge (j, k) by (j, i) and (i, k), as shown in Figure 6.5(b). Let now i be a vertex with out-degree 2 and in-degree 1, call j the vertex left by an edge arriving at i, and k, l the other neighbors of i. The reverse directed edge splitting operation consists in removing i and its incident edges, and adding either (j, k) or (j, l) (k and l being interchangeable) *in such a way that the graph obtained is minimally rigid.*

Lemma 6.1. *The directed vertex addition and edge splitting operations preserve minimal persistence, and so do the reverse directed vertex addition and reverse directed edge splitting operations.*

Proof. All these operations preserve minimal rigidity as their undirected counterpart do. Moreover, they respectively add or remove a vertex with out-degree 2 without affecting the out-degree of the other vertices. It follows thus from Proposition 6.5 that they preserve minimal persistence. □

We denote by $\mathcal{S}$ the set of operations containing the directed vertex addition operation and the directed edge splitting operation, and by $\mathcal{S}^{-1}$ the set of operations containing their reverse versions (the same convention is used in the sequel for all operation sets). The smallest minimally persistent graph on more than one vertex consists in two vertices connected by one directed edge. We refer to this graph as a *leader-follower pair*, the leader being the vertex with an out-degree 0. Since the operations in $\mathcal{S}$ preserve minimal persistence, any graph obtained by performing a sequence of directed vertex addition or edge splitting operations on an initial leader-follower pair is minimally persistent. The following result establishes that to any minimally rigid graph corresponds a minimally persistent graph that can be obtained in that way.

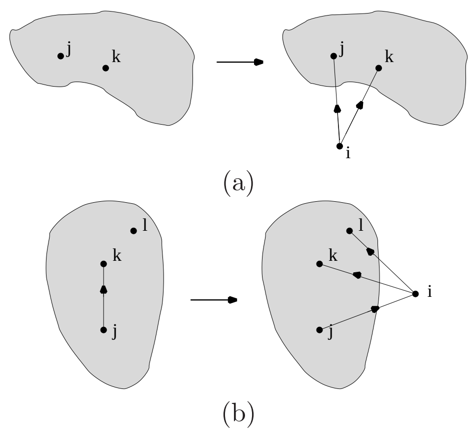

Figure 6.5: Representation of the directed vertex addition (a) and edge splitting (b) operations.

Proposition 6.6. *It is possible to assign directions to the edges of any minimally rigid graph such that the obtained directed graph is minimally persistent and can be obtained by performing a sequence of operation in $\mathcal{S}$ on an initial leader-follower pair. Moreover, all intermediate graphs are minimally persistent.*

Proof. Let G be a minimally rigid (undirected) graph. By Theorem 6.2, it can be obtained by performing a sequence of undirected vertex additions and edge splittings on K_2. By performing the same sequence of the directed version of these operations on an initial leader-follower pair, one obtains a directed graph having G as underlying undirected graph. Moreover, since this initial seed is minimally persistent and since the directed versions of both vertex addition and edge splitting preserve minimal persistence, the obtained graph and all the intermediate graphs are minimally persistent. □

The operations of $\mathcal{S}$ are however not sufficient to build all minimally persistent graphs. We have indeed shown the existence of infinitely many minimally persistent graphs that cannot be obtained from a smaller graph by performing an operation of $\mathcal{S}$ [66, 68, 69]. In the same works, we have also considered generalized versions of vertex addition and edge splitting operations, defined by allowing any direction for the added edges. These operations can generally not be used on any vertex as they may increase some vertices' out-degrees, but their use allows one to build more minimally persistent graphs. However a minimally persistent graph on 24 vertices was found that cannot be built even with those generalized operations.

6.2.4 A third purely directed operation

Edge reversal

Let (i, j) be an edge such that j has at least one degree of freedom, i.e., $d^+(j) = 0$ or $d^+(j) = 1$. The *edge reversal* operation consists in replacing the edge (i, j) by (j, i). As a consequence, one degree of freedom is transferred from j to i. This operation is its auto-inverse and preserves minimal persistence since it does not affect the underlying undirected graph and the only increased out-degree $d^+(j)$ remains no greater than 2. From an autonomous agent point of view j transfers its decision power or a part of it to i. We now define two macro-operations based on repeated applications of edge reversal.

Path reversal

Given a directed path P between a vertex i and a vertex j such that j has a positive number of degrees of freedom, a *path reversal* consists in reversing the directions of all the edges of P. As a result, j loses a degree of freedom, i acquires one, and there is a directed path from j to i. Moreover, the number of degrees of freedom of all the other vertices remain unchanged. Note that i and j can be the same vertex, in which case the path either has a trivial length 0 or is a cycle. In both of these situations, the number of degrees of freedom is preserved for every vertex.

The path reversal can easily be implemented with a sequence of edge reversals: Since j has a degree of freedom, one can reverse the last edge of the path, say (k, j), such that j loses one degree of freedom while k acquires one. One can then iterate this operation along the path until i, as shown in Figure 6.6. At the end, i has an additional degree of freedom, j has lost one, and all the edges of the paths have been reversed. Note that the sequence of edge reversals can usually not be performed in another order, for the condition requiring the availability of a degree of freedom would not be satisfied. The final result would be the same, but all the intermediate graphs would not necessarily be minimally persistent.

The following lemma, which is a particular case of Theorem 4.6, implies that a degree of freedom can be transferred from any vertex having at least one to any other having less than two of them using a path reversal.

Lemma 6.2. *Let G be a minimally persistent graph, i and j two vertices of G with $d^+(i) \geq 1$ and $d^+(j) \leq 1$. Then, there is a directed path from i to j.*

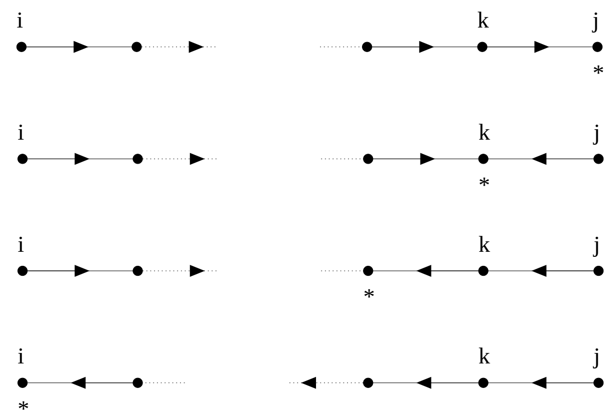

Figure 6.6: Implementation of the path reversal by a sequence of edge reversals. The symbol "*" represents one degree of freedom.

Cycle reversal

A *cycle reversal* consists in reversing all the edges of a directed cycle. This operation does not affect the number of degrees of freedom of any vertex nor the underlying undirected graph, and preserves therefore minimal persistence. Besides, as opposed to a path reversal operation applied to the particular case of a cycle, it does not require the presence of a degree of freedom.

A cycle reversal on a minimally persistent graph can be implemented by a sequence of edge reversals. Let us indeed first suppose that there is a vertex i in the cycle that has at least one degree of freedom. In that case, the cycle reversal is just a particular case of the path reversal, with $i = j$. We now assume that no vertex in the cycle has a degree of freedom. Let l be a vertex in the cycle, and m a vertex that does not belong to the cycle but has a degree of freedom. The existence of a directed path from l to m follows from Lemma 6.2. Let i be the last vertex in this path belonging to the cycle. There is trivially a path P from i to m such that every other vertex of this path does not belong to the cycle. The implementation of a cycle reversal by three path reversals is then represented in Figure 6.7. One begins by reversing the path P into P' such that i acquires a degree of freedom. As explained above, the cycle can then be reversed since it is a particular case of path reversal, and finally, one reverses the path P' back to P such that the degree of freedom acquired by i is re-transmitted to m.

Both cycle reversal and path reversal are their auto-inverse, as is the case for edge reversal. Moreover, the fact that they can be implemented using only

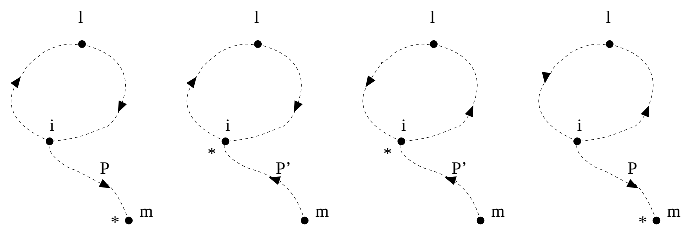

Figure 6.7: Implementation of the cycle reversal operation using path reversals. The "*" represents one degree of freedom.

edge reversals is another way to show that they preserve minimal persistence.

It follows from Lemma 6.2 that one can arbitrarily reposition degrees of freedom using path reversals. The following result implies that two minimally persistent graphs having the same underlying undirected graph and the same positions for their degrees of freedom (all vertices having therefore the same out-degree in the two graphs) can differ only by cycles of opposite edges, and its proof provides a greedy algorithm to find such a cycle.

Lemma 6.3. *Let $G_A = (V, E_A)$ and $G_B = (V, E_B)$ be two graphs having the same underlying undirected graph and such that every vertex has the same out-degree in both graphs. If an edge of G_A has the opposite direction to that in G_B, then it belongs to a cycle of such edges in G_A.*

Proof. Suppose that $(i_0, i_1) \in E_A$ and $(i_1, i_0) \in E_B$ (i.e., this edge has opposite directions in G_A and G_B); then there exists at least one vertex $i_2 \neq i_0$ such that $(i_1, i_2) \in E_A$ and $(i_2, i_1) \in E_B$. For if the contrary holds, we would have $d^+(i_1, G_A) = d^+(i_1, G_B) - 1$, which contradicts our hypothesis. Repeating this argument recursively, we obtain an (infinite) sequence of vertices $i_0, i_1, i_2, \dots$ such that for each $j \geq 0$, $(i_j, i_{j+1}) \in E_A$ and $(i_{j+1}, i_j) \in E_B$. Since there are only a finite number of vertices in V, at least one of them will appear twice in this sequence. By taking the subsequence of vertices (and induced edges) appearing in the infinite sequence between any two of its occurrences we obtain then a cycle of edges of G_A having opposite directions to those in G_B. This cycle does not necessarily contain (i_0, i_1). But if it does not, we can re-apply the same argument to G'_A, G'_B obtained from G_A and G_B by removing the edges of the cycle found. (i_0, i_1) has indeed an opposite direction in G'_A to that in G'_B, and these graphs satisfy the other hypotheses of the lemma. Moreover, they contain less edges than G_A, G_B. Therefore by doing this recursively, we eventually obtain a cycle containing (i_0, i_1) since the number of edges in the graphs is finite. $\square$

6.2.5 Obtaining all minimally persistent graphs

Using the results on the two macro-operations defined above, we can now show the following proposition.

Proposition 6.7. *By applying a sequence of edge reversals to a given minimally persistent graph, it is possible to obtain any other minimally persistent graph having the same underlying undirected graph. Moreover, all the intermediate graphs are then minimally persistent.*

Proof. Let G_A and G_B be two minimally persistent graphs having the same underlying undirected graph. Suppose that there is a vertex i which has less degrees of freedom in G_A than in G_B. Since at most three vertices have positive degree of freedom, there are at most three such vertices i. And since the total number of degrees of freedom is 3 in all minimally persistent graphs, there exists a vertex j which has more degree(s) of freedom in G_A than in G_B. In G_A i has thus necessarily less than two degrees of freedom and j has at least one degree of freedom. It follows then from Lemma 6.2 that there exists a directed path from i to j in G_A. The reversal of this path transfers a degree of freedom from j to i without affecting the number of degrees of freedom of the other vertices. Doing this at most two more times, the two graphs will have the same positions for their degrees of freedom.

We now show that the following algorithm, which uses only cycle reversals, transforms then G_A into G_B:

while $\exists$ e having opposite direction in G_A to that in G_B **do**

> Select a cycle C of such edges
> Reverse C in G_A

end do

existence of C when $G_A \neq G_B$: This is a direct consequence of Lemma 6.3 since both graphs have the same underlying undirected graph and since all the vertices have the same out-degrees in both of them.

end of the algorithm: At each step of the loop, the number of edges having opposite directions in G_A and G_B is strictly reduced because all the edges for which directions are changed in G_A initially had an opposite direction in G_B (and because Proposition 6.4 forbids the presence of cycles of length 2 in a minimally persistent graph). Since there are only a finite number of edges, the algorithm finishes, and all the edges have then the same directions in both graphs.

The result of this proposition then follows from the fact that both path reversal and cycle reversal can be implemented by a sequence of edge reversals, which preserves minimal persistence. □

From an autonomous agent formation perspective, suppose that a reorganization of the distance constraints distribution needs to be performed, and that this reorganization preserves the structure of constraints from an undirected point of view, i.e., the reorganization only involves changes of some directions. Proposition 6.7 implies that this can be done by a sequence of local degree of freedom transfers, in such a way that during all the intermediate stages, the formation shape is guaranteed to be maintained.

Let $\mathcal{T}$ be the set of operations containing vertex addition, edge splitting, and edge reversal. We can now state the main result of this section.

Theorem 6.3. *Every minimally persistent graph can be obtained by applying a sequence of operations in $\mathcal{T}$ to an initial leader-follower seed. Moreover, all the intermediate graphs are minimally persistent.*

Proof. Consider a minimally persistent graph G. This graph is also minimally rigid. By Proposition 6.6, there exists thus a (possibly different) minimally persistent graph having the same underlying undirected graph that can be obtained by performing a sequence of operations in $\mathcal{S} \subset \mathcal{T}$ on an initial leader follower seed. By Proposition 6.7, G can then be obtained by applying a sequence of edge reversals on this last graph. Moreover, since all the operations in $\mathcal{T}$ preserve minimal persistence, all the intermediate graphs are minimally persistent. □

To illustrate Theorem 6.3, consider the graph G represented in the right hand side of Figure 6.8(c). This graph cannot be built using operations of $\mathcal{S}$, as none of its vertices can be removed using an operation of $\mathcal{S}^{-1}$. Only 4 satisfies indeed the necessary condition on the out-degree. And if 4 is removed by a reverse edge splitting, either $(3,5)$ or $(3,2)$ should be added, leading in both cases to a non-rigid graph. The graph can however be obtained by applying a sequence of operations in $\mathcal{T}$ on an initial leader-follower seed. Let us take 1 and 2 as respectively leader and follower of this initial seed. One can begin by adding 3, 4 and 5 using three vertex additions as shown in Figure 6.8(a). The graph obtained has the same underlying undirected graph as G, but the degrees of freedom are not allocated to the same vertices. By reversing the path $(5,4,2,1)$ using a sequence of edge reversals, one can then transfer one degree of freedom from 1 to 5 as shown in Figure 6.8(b) such that in the obtained graph, all vertices have the same number of degrees of freedom (and therefore same out-degree) as in G. As stated in Lemma 6.3, any edge of this graph that does not have the same direction as in G belongs to a cycle of such edges. The only such cycle here is C. By reversing it using a sequence of edge reversals, one

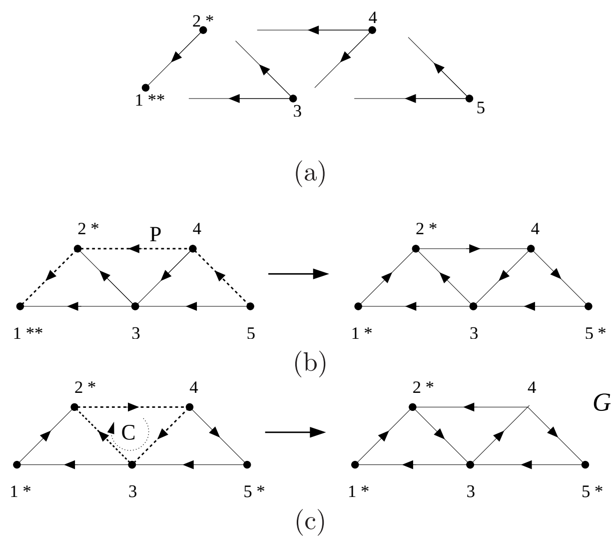

Figure 6.8: Example of obtaining of a minimally persistent graph by applying a sequence of operations in $\mathcal{T}$ on a leader-follower seed. The graph G is obtained from the leader-follower seed by (a) three vertex additions, (b) the reversal of the path P and (c) of the cycle C.

finally obtains the graph G, as shown in Figure 6.8(c). Note that, consistently with Theorem 6.3, all the intermediate graphs are minimally persistent.

Corollary 6.4. *Every minimally persistent graph can be transformed into any other minimally persistent graph using only operations in* $\mathcal{T} \cup \mathcal{T}^{-1}$.

Proof. Let G_A and G_B be two minimally persistent graphs. Since G_A can be obtained by applying a sequence of operations in $\mathcal{T}$ on a leader-follower pair, the leader-follower pair can be re-obtained from G_A by applying the reverse versions of these operations (which are all in $\mathcal{T}^{-1}$) in the reverse order. By Theorem 6.3 one can then obtain G_B from this leader-follower pair by a sequence of operations in $\mathcal{T}$. $\square$

The method proposed in the proof of Corollary 6.4 is generally not optimal in terms of the number of operations. Note also that unlike in the case of undirected Henneberg sequences, the number of operations to build a minimally persistent graph is not uniquely fixed by its number of vertices, although it is bounded in $O(|V|^2)$. The first part of our construction requires indeed $|V|-2$ operations of $\mathcal{S}$. The second part requires up to 3 path reversals (one for each degree of freedom), and up to $\frac{1}{3}|V|$ cycle reversals (the smallest possible cycles have a length 3 in a minimally persistent graph). The bound follows then from

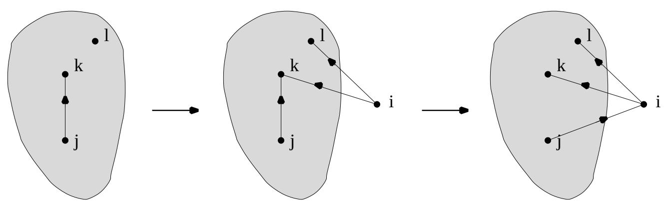

Figure 6.9: Implementation of the edge splitting by a vertex addition and an edge reorientation. The vertex i is first added with two outgoing edges by vertex addition, and the edge (j, k) is then reoriented and becomes (j, i).

the fact that a path reversal requires at most $|V|$ edge reversals, and a cycle reversal can be implemented by two path reversals. The existence of a smaller bound or of minimally persistent graphs that requires $O(|V|^2)$ operations to be built remains an open question.

Finally, observe that the three operations in $\mathcal{T}$ are basic operations that can be performed locally. They can thus easily be implemented in a local way on an autonomous agent formation. It might however be possible to improve this basic character using for example an operation such as an *edge reorientation*, i.e., an operation consisting in changing the arrival vertex of an edge. As shown in Figure 6.9, a vertex addition operation and an edge reorientation operation can indeed implement an edge splitting operation which could thus be discarded. However, this would require an efficient and simple criterion to determine when such an edge reorientation operation can be performed, and no such criterion is presently available.

6.3 Two-dimensional graphs with three degrees of freedom

6.3.1 Introduction and relevance

In this section, we consider graphs representing two-dimensional formations that have three degrees of freedom. We show how their persistence can be checked in polynomial time, a result recently obtained by Bang-Jensen and Jordán [12]. Graphs with three degrees of freedom are those in which all vertices have an out-degree 2 or more, except either three vertices who have an out-degree 1 or one vertex with an out-degree 1 and one vertex with an out-degree 0. It follows from Proposition 4.2 that 3 is the maximal number of degrees of freedom that a persistent graph can have in $\Re^2$. Although some persistent graphs have less than 3 degrees of freedom, the class of those having three of them is in practice the most important one. They are indeed the only ones whose corresponding formation can be fully controlled by individual actions of the agents. As explained in Section 4.5, formations with less degrees of freedom can indeed not perform certain moves unless a common decision is made by several agents in a collaborative way, an hypothesis that we do not consider in this thesis.

Particular cases of two-dimensional persistent graphs with three degrees of freedom include acyclic graphs (see Theorem 6.1), minimally persistent graphs (see Corollary 4.4 or Section 6.2.1) and leader-follower graphs. A leader-follower graph is one in which one "leader" vertex has an out-degree 0, and one "first follower" vertex has an out-degree one, with its single outgoing edge arriving at the leader vertex. These graphs are advisable for many practical applications as the displacements of the corresponding formation are easy to control. The leader is indeed responsible for the formation's position and the first follower for its orientation. Recent results have shown that any minimally persistent two-dimensional leader-follower formation can efficiently be controlled [3]. However, the study and use of graphs with less degrees of freedom should not be a priori ruled out, especially in applications where the orientation for example of the formation is less relevant than its robustness with respect to loss of edges.

All results in this section were obtained by Bang-Jensen and Jordán [12], but the presentation and notations are modified to be consistent with the other chapters. Besides, all graphs in this section are supposed to be represented in $\Re^2$.

6.3.2 Polynomial-time algorithm

We first give a simpler criterion for persistence that can be applied to graphs with three degrees of freedom.

Proposition 6.8 (Bang-Jensen and Jordán [12]). *Let G be a directed graph for which the total number of degrees of freedom in $\Re^2$ is 3. G is persistent if and only if there is no subgraph $G'(V',E')$ of G for which $\sum_{i\in V'} \max(0, 2-d^+_{i,G'}) < 3$, i.e. a subgraph in which the total number of degrees of freedom is smaller than 3.*

Proof. Let $\Sigma(G)$ be the set of all strict $\min D$-subgraph of G on all its vertices, i.e., the set of the subgraphs S of G on all its vertices such that $d^+_{i,S} = \min(2, d^+_{i,G})$. Theorem 4.5 characterizing persistent graphs states that G is persistent if and only if every graph in $\Sigma(G)$ is rigid. Since $\sum_{i\in V} \min(2, d^+_{i,G}) = 2|V| - 3$ as G has three degrees of freedom, every graph $S \in \Sigma(G)$ contains $2|V| - 3$ edges. So by Theorem 3.2, it is rigid if and only if it contains no subgraph $S'(V', E_{S'})$ for which there holds $|E_{S'}| > 2|V'| - 3$.

Suppose first that there exists a subgraph G' of G for which $\sum_{i\in V'} \max(0, 2- d^+_{i,G'}) < 3$ holds. By removing edges leaving vertices with out-degree larger than 2, one can obtain a subgraph of $S'(V', E_{S'})$ of G' for each vertex i of which there holds $d^+_{i,S'} = \min(2, d^+_{i,G'})$ and thus $\max(0, 2 - d^+_{i,G'}) = 2 - d^+_{i,S'}$. For this graph, there holds then

$$|E_{S'}| = \sum_{i\in V'} d^+_{i,S'} = 2\,|V'| - \sum_{i\in V'} (2 - d^+_{i,S'}) > 2\,|V'| - 3.$$

Since for each $i \in V'$ there holds $d^+_{i,S'} = \min(2, d^+_{i,G'}) \leq \min(2, d^+_{i,G})$, it is possible to obtain a subgraph $S \in \Sigma(G)$ from S' by adding all the vertices of $V \setminus V'$ and some outgoing edges. Since S' is a subgraph of S for which $|E_{S'}| > 2\,|V'| - 3$, S is not rigid and G is not persistent.

Conversely, suppose that G is not persistent, and thus that there is a graph $S \in \Sigma(G)$ admitting a subgraph $S'(V', E_{S'})$ for which there holds $|E_{S'}| > 2\,|V'| - 3$. There holds thus $\sum_{i\in V'} d^+_{i,S'} \geq 2\,|V'| - 3$. Since no vertex has in $S' \subseteq S$ has an out-degree larger than 2, this implies that $\sum_{i\in V'} \max(0, 2 - d^+_{i,S'}) < 3$ holds for this subgraph S'. □

The following lemma describes a simple algorithm, a variation of which appears in [12].

Lemma 6.4. *Let $G(V, E)$ be a directed graph, and assign an integer b_i to each vertex i. There exists a polynomial time algorithm checking the existence of a non-empty subgraph $G'(V', E')$ such that for each $i \in V'$ there holds $d^+_{i,G'} \geq b_i$. Moreover, the algorithm provides the largest of these subgraphs when the answer is positive.*

Proof. The idea of the algorithm is to remove vertices having a too small out-degree as long as such vertices can be found. More formally:

Let $t = 0$, $G_0(V_0, E_0) = G(V, E)$
while $\exists\ i \in V_t$ such $d^+_{i,G'} < b_i$ **do**

> Take $G_{t+1} = G_t$ without i and its incident edges
> $t = t + 1$

end do $G^* = G_t$.

Executing this algorithm clearly takes a polynomial time, which can be as large as $O(|E|)$ if the initial graph is connected. Moreover, if G^* is not empty at the termination of the algorithm, each of its vertices i satisfies $d^+_{i,G'} \geq b_i$. Conversely, suppose that there is a subgraph G' of G for each vertex of which there holds $d^+_{i,G'} \geq b_i$, and let us show inductively that G' is a subgraph of G^*

obtained at the end of the algorithm. Obviously, G' is a subgraph of the initial graph $G_0 = G$. Therefore, $d^+_{i,G_0} \geq d^+_{i,G'} \geq b_i$ holds for each $i \in V'$. So if a vertex is removed by the algorithm, it does not belong to V'. G' is thus also a subgraph of G_1. Repeating this argument for each iteration, we see that no vertex of G' is ever removed by the algorithm. □

We can now use this algorithm to check the persistence of the graph by checking the existence of a subgraph as described in Proposition 6.8.

Theorem 6.4 (Bang-Jensen and Jordán [12]). *The persistence in $\Re^2$ of a directed graph having 3 degrees of freedom can be checked in $O(|V|^2 |E|)$.*

Proof. It follows from Proposition 6.8 that the non-persistence of G is equivalent to the existence of a subgraph $G'(V', E')$ for which $\sum_{i \in V'} \max(0, 2 - d^+_{i,G'}) \leq 2$ holds, that is, a subgraph in which all vertices have an out-degree no smaller than 2, with the possible exception of either one vertex with an out-degree 0 or two vertices with an out-degree 1. To check persistence, one needs thus just to check first for each vertex i if there is a subgraph of G in which each vertex has an out-degree at least 2 except possibly i, and then to check for each pair of vertices i, j if there is a subgraph of G in which each vertex has an out-degree at least 2 except possibly i and j who only need to have an out-degree at least 1. Each of these tests can be done in polynomial time using Lemma 6.4, with either $b_i = 0$ and $b_k = 2$ for every other k, or with $b_i = b_j = 1$ and again $b_k = 2$ for every other k. Since at most $|V| + \frac{1}{2} |V| (|V| - 1)$ tests need to be done, the total execution time is bounded by $O(|V|^2 |E|)$. □

Remember that persistence and structural persistence are equivalent notions when graphs are represented in $\Re^2$ (See Corollary 5.2). The algorithm provided in Theorem 6.4 therefore also checks structural persistence. Note that the complexity of the algorithm presented is probably not optimal. Moreover, its efficiency can in practice be improved by first sequentially removing all vertices with in-degree 0 using Proposition 6.3. Besides, the existence of a polynomial-time algorithm to check persistence without assuming that the graph has three degrees of freedom is an open question.

6.4 History and literature

The results in Section 6.1 on acyclic graphs were presented in [63, 64] for two-dimensional spaces and partly extended in [133, 133] for higher dimensions. The construction of all two-dimensional minimally persistent graphs results in Section 6.2 are appeared in [66, 68, 69] which contains also other set of operations, and detailed proof of the impossibility to build all minimally persistent graph with certain type of operation sets. Some preliminary results had also been obtained in [64, 133]. Finally, the persistence testing results of Section 6.3 for two-dimensional graphs with three degrees of freedom were obtained by Bang-Jensen and Jordán [12].

Chapter 7

Further Research Directions

This research has lead to many open questions and ideas for possible further investigation, some of which have already been mentioned in the previous chapters or in published papers. We present them in the rest of this chapter, together with some tracks that could be followed to approach them and some partial results when available. Section 7.1 is dedicated to further research directions on shape maintenance in formation control. Section 7.2 contains more precise open questions appearing in the context of persistence, and whose formulation is mostly graph-theoretical.

7.1 Further research directions in formation control

7.1.1 Convergence to equilibrium

This issue is related to the content of Sections 4.3, 5.1 and 5.2

Constraint consistence characterizes the fact that once a formation reaches equilibrium, all constraints are satisfied. If the formation is also rigid, this implies that the formation shape is the desired one. This description does not take the issue of convergence to equilibrium into account. Progress in this direction is made by defining in Chapter 5 structural constraint consistence and persistence. These notions ensure that if a set of displacements in which some agents have fixed displacements and others are free is invariant[1], then it contains an

[1]Remember that a set S of system states is an invariant set if when the system state is in S it remains in S for all further time.

equilibrium. Even if that would not necessarily be sufficient to guarantee the convergence to equilibrium, it would be desirable to have a stronger notion, ensuring the fact that every invariant set's closure[2] contains an equilibrium. If possible, this notion should be formulated in a unified way: One could for example require every invariant set's closure to contain an admissible displacement, that is, a displacement for which all constraints are satisfied. This would first imply that a formation can always get arbitrarily close to an admissible displacement and thus to an equilibrium. Since an equilibrium displacement constitutes a invariant set, it would also imply that all constraints are satisfied at every equilibrium displacement. We do however not know if such notion could be formalized without making strong assumptions on the agents' internal control laws. As a related issue, one should also make sure that when a formation converges to an equilibrium, this equilibrium is "close" to the initial conditions. Persistence, constraint consistence and rigidity are indeed local notions. Persistence guarantees for example that the formation shape is the desired one provided that all agents are at equilibrium, and that the formation is in *a certain neighborhood of a reference configuration*. One should thus make sure that the control laws do not drive the formation to an equilibrium lying out of this neighborhood, for otherwise the formation shape might not be the desired one, and some constraints might even not be satisfied. Note that this issue could be related to global constraint consistence and persistence, that we mention in Section 7.1.4.

For convenience, we use the infinitesimal displacement formalism as for example in Sections 3.2 and 4.3. In the definition of constraint consistence, we require the admissibility of all equilibrium (infinitesimal) displacements. Stronger notions could be obtained by requiring the admissibility of all (infinitesimal) displacements in larger sets. We present here such a notion, and show that it encapsulates both constraint consistence and structural constraint consistence. How this notion relates to the invariant sets issue, and what we abusively call its "physical meaning", need however still to be determined.

Remember that the displacements considered for constraint consistence are those where each agent is at equilibrium, that is, those for which no agent could satisfy an additional constraint without breaking one that it already satisfies, considering the other agents' displacement as fixed. This idea can be generalized to groups of agents.

[2] Considering the set's closure is needed as one could have a formation converging to equilibrium without ever reaching it. The formation's trajectory would then be an invariant set containing no equilibrium, but its closure would contain one.

Definition 7.1. *Let p be a representation of a directed graph $G(V,E)$, and consider a subset of vertices $V_1 \subseteq V$. We say that the set V_1 is at equilibrium with a certain displacement δp if there is no $\delta p'$ such that*

a) $\delta p'_j = \delta p_j$ for every $j \in V_1$,
b) The set of constraints corresponding to edges leaving vertices of V_1 that are satisfied by $\delta p'$ strictly contains the set of those satisfied by δp.

Intuitively, a set of agents is at equilibrium if no concerted displacement of the agents leads to the satisfaction of an additional constraint without breaking one that they already satisfy, considering the displacement of the agents outside the set as fixed. In particular, a set of one single agent is at equilibrium if and only if the agent is at the equilibrium according to our usual definition (see Section 4.3).

However, a set of agents at equilibrium is not necessarily at equilibrium as a set. A collective decision could indeed possibly improve the situation for everybody while no agent could improve its own situation by an isolated decision. Conversely, if a set of agents is at equilibrium as a set, each agent is not necessarily at equilibrium individually. It can indeed be that some agents could improve their situation, but that by doing so they break other agent's constraint, so that the situation of the group is not improved. The notion of set at equilibrium does thus not exactly correspond to the intuitive notion of equilibrium. We say that a displacement is a *set-wise equilibrium* if there exists a partition of the vertices $V = V_1 \cup \cdots \cup V_c$ ($V_i \cap V_j = \varnothing$ if $i \neq j$) such that every set V_i is at equilibrium for this displacement.

Definition 7.2. *A graph representation is* set-wise constraint consistent *if every set-wise equilibrium displacement is admissible by the representation, i.e., is such that all linear constraints implied by the rigidity matrix are satisfied.*

As an example of application of this definition, consider the representation in Figure 7.1. Let us first assign δp_1 and δp_2 such that $(p_1 - p_2)^T(\delta p_1 - \delta p_2) \neq 0$, that is, such that the distance between 1 and 2 varies. We have seen in Section 5.1 that it is then impossible for the three other agents to simultaneously satisfy all their constraints. One can however choose δp_3, δp_4 and δp_5 such that all constraints are satisfied but one, say the one associated to $(3,4)$. Consider then the partition $V = \{1\} \cup \{2\} \cup \{3,4,5\}$. The two single-vertex sets are obviously at equilibrium as they have no constraints. The third set is also at equilibrium as a set (although 3 is not at equilibrium as an agent) as it is impossible for $\{3,4,5\}$ to simultaneously satisfy all their constraints. So this displacement is a set-wise equilibrium displacement, but is not admissible by the graph representation, which is thus not set-wise constraint consistent.

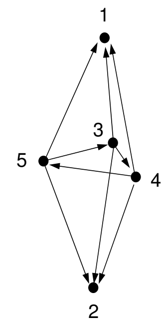

Figure 7.1: Example of three-dimensional representation that is constraint consistent but not structurally constraint consistent nor set-wise constraint consistent.

To further analyze set-wise constraint consistence, we look at some relevant classes of set-wise equilibrium displacements. Obviously, every admissible displacement is a set-wise equilibrium displacement, as if all constraints are satisfied, all sets of any partition of the vertices are at equilibrium. Moreover, any (normal) equilibrium displacement is also a set-wise equilibrium displacement. Indeed, if every vertex is at equilibrium, then the set of vertices V can be decomposed in $|V|$ single-vertex sets which are each at equilibrium as sets.

Proposition 7.1. *Every set-wise constraint consistent representation is constraint consistent.*

Proof. A graph representation is constraint consistent if every equilibrium displacement is admissible. The result follows thus directly from the fact that every equilibrium displacement is also a set-wise equilibrium displacement. $\square$

We now show that set-wise constraint consistence also implies structural constraint consistence.

Lemma 7.1. *Let p be a representation of a graph $G(V,E)$ in $\Re^D$, and $V = V_1 \cup V_2$ a partition of the vertex set. For every $\delta p_{V_2} \in \Re^{D|V_2|}$, there is a $\delta p_{V_1} \in \Re^{D|V_1|}$ such that V_1 is at equilibrium as a set for the displacement obtained by aggregating δp_{V_1} and δp_{V_2}.*

Proof. Let us fix a δp_{V_2} and take an initial arbitrary δp_{V_1}. We call E_1 be the set of edges leaving vertices of V_1, whether they arrive in V_1 or in V_2. We suppose that the vertices are labelled in such a way that the displacement obtained by aggregating δp_{V_1} and δp_{V_2} is $[\delta p_{V_1}^T, \delta p_{V_2}^T]^T$. If there is a $\delta p'_{V_1}$ such that $[\delta p_{V_1}^T, \delta p'^T_{V_2}]^T$ satisfies all constraints corresponding to edges of E_1 and at least an additional one, take this $\delta p'_{V_1}$ as new δp_{V_1}. Repeating this procedure, we obtain after at most $|E_1|$ iterations a δp_{V_1} satisfying a maximal subset of constraints of E_1, so that V_1 is at equilibrium as a set with $[\delta p_{V_1}^T, \delta p_{V_2}^T]^T$. $\square$

Proposition 7.2. *Every set-wise constraint consistent representation is structurally constraint consistent*

Proof. By Proposition 7.1 and the definition of structural constraint consistence, it suffices to prove that a constraint consistent but not structurally constraint consistent representation is not set-wise constraint consistent either. Let p be such a representation. Since it is not structurally constraint consistent, there exists a set V_2 and a partial equilibrium δp_{V_2} (that is, displacements for every completion of which the vertices of V_2 are at equilibrium) that cannot be completed into an equilibrium displacement. In particular, it cannot be completed into an admissible displacement. Observe now that the completion of δp_{V_2} by the δp_{V_1} provided by Lemma 7.1 is a set-wise equilibrium displacement. V_1 is indeed at equilibrium as a set by Lemma 7.1, and every other vertex is at equilibrium as a single-vertex set because δp_{V_2} is a partial equilibrium. There is thus a set-wise equilibrium displacement that is not admissible by p, preventing it from being set-wise constraint consistent. $\square$

Studying this notion of set-wise constraint consistence could thus be promising. Among the many open questions, one can wonder if this notion is a generic one, i.e., one that almost only depends on the graph. Also, since persistence and structural persistence are equivalent notions in $\Re^2$, maybe set-wise persistence (i.e. set-wise constraint consistence and rigidity) is equivalent to persistence in $\Re^2$. However, one would also need to understand the meaning and the relevance of the set-wise equilibrium displacements in terms of autonomous agent systems. Besides, it uses the idea of satisfying a maximal subset of constraints, which is partly arbitrary (see Section 7.1.2). Finally, it excludes representations that are structurally persistent, such as the one obtained from the representation in Figure 7.1 by connecting 1 to 2. Proving that this representation is not set-wise constraint consistent can be done using the same displacement as above for the representation in Figure 7.1.

7.1.2 Different equilibrium criterion for the agents

How much relies on Assumption 4.2 as compared to Assumption 4.1? This issue is related to the content of Sections 4.1 and 4.3.

Our characterization of persistence relies on Assumption 4.2, stating that an agent is at equilibrium if and only if it satisfies a maximal set of constraints. As explained in Section 4.1, this assumption is reasonable but partly arbitrary.

We therefore briefly analyze now how our results would be extended or modified if a different assumption was used. We first consider the results that are valid for any reasonable assumption, that is, any assumption particularizing Assumption 4.1. The latter states that an agent is at equilibrium when it satisfies all its constraints, and not at equilibrium if it does not but could satisfy all its constraints. We then propose such an assumption based on the least-square satisfaction of the linear system of constraints. Before starting, remember that modifying or replacing Assumption 4.1 implies a different definition of equilibrium, and therefore of constraint consistence. It has thus no effect on rigidity. Moreover, we see below that it has no effect either on the constraint consistence and persistence of acyclic graphs and of graphs where the maximal out-degree is D. As a consequence, the issue of convergence to equilibrium presented in Section 7.1.1 is relevant independently of the assumption used, as it concerns among others such graphs with maximal out-degree D (see Figure 5.1).

Generic assumptions

Let us take a representation p of a directed graph $G(V,E)$, and an (infinitesimal) displacement δp. Remember that if a vertex i is connected by directed edges to vertices $j_1, \dots, j_{d_i^+}$, the corresponding agent faces the following linear system:

$$\begin{pmatrix} (p_i - p_{j_1})^T \\ (p_i - p_{j_2})^T \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \end{pmatrix} \delta p_i = \begin{pmatrix} (p_i - p_{j_1})^T \delta p_{j_1} \\ (p_i - p_{j_2})^T \delta p_{j_2} \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \delta p_{j_{d_i^+}} \end{pmatrix}. \tag{7.1}$$

Assumption 4.1 states that an agent satisfying all its constraints is at equilibrium, and that an agent that does not satisfy all its constraints but can move to a position where it would satisfy all of them (considering all other agents as fixed) is not at equilibrium. We thus say that if the system (7.1) for i is satisfied by δp, then i is at equilibrium. On the other hand, if δp does not satisfy the system while it admits a solution $\delta p'_i$ (considering every other δp_j as fixed), then i is not at equilibrium. This does of course not fully define the set of equilibrium displacements, as it does not treat the case of systems admitting no solution. We however show that some results hold for any definition of equilibrium satisfying these properties. Note that a third part of Assumption 4.1 states that, considering its δp_j as fixed for every neighbor j of i, there always exists a δp_i such that i is at equilibrium.

Suppose now that we fix a definition of equilibrium, without explicitly specifying it. We use the same formal definition of (infinitesimal) constraint consistence and persistence as in Section 4.3. For a representation p of a graph

G, an *equilibrium (infinitesimal) displacement* is an (infinitesimal) displacement for which all vertices are at equilibrium, and the set of all equilibrium displacements is denoted by $\text{Equil}^*_{G,p}$. A representation is *(infinitesimally) constraint consistent* if every equilibrium displacement is admissible by the representation and the graph, that is if such displacement satisfies all constraints ($\text{Equil}^*_{G,p} \subseteq \text{Ker}R_{G,p}$). A representation is *(infinitesimally) persistent* if every equilibrium displacement is Euclidean ($\text{Equil}^*_{G,p} \subseteq Eu_p$). Note that these notions are not necessarily generic notions for graphs, which leads us to a first open question.

Open question 1. *Does any definition of equilibrium consistent with Assumption 4.1 lead to notions of persistence and constraint consistence that are generic notions for graphs?*

Obviously, every admissible displacement is an equilibrium displacement, so that the inclusion $\text{Ker}R_{G,p} \subseteq \text{Equil}^*_{G,p}$ holds. Indeed, if all constraints are satisfied, the system (7.1) associated to each vertex is also satisfied. The following decomposition of persistence is thus still valid as its proof solely relies on this inclusion.

Theorem 7.1. *For any definition of equilibrium consistent with Assumption 4.1, a representation is persistent if and only if it is rigid and constraint consistent.*

Observe now that if p is not constraint consistent, there exists an equilibrium displacement δp for which at least one constraint is not satisfied, and thus for which the system (7.1) of at least one vertex is not satisfied. Since this vertex is at equilibrium this implies that this system admits no solution (considering other agents as fixed). Graph representations for which the system (7.1) always admits a solution are thus always constraint consistent. Consider for example a non-degenerate representation of a graph containing no vertex with out-degree larger than D. The system (7.1) of each vertex i contains $d_i^+ \leq D$ linearly independent equations involving the same $D-$dimensional variable δp_i, and admits thus always a solution. This together with Theorem 7.1 proves the next proposition.

Proposition 7.3. *For any definition of equilibrium consistent with Assumption 4.1, any non-degenerate representation of a graph whose largest out-degree is bounded by D is constraint consistent. Such a representation is thus persistent if and only if it is rigid.*

To avoid any possible ambiguity, we do not omit the term "generic" in this section. We thus say that a graph is *generically constraint consistent (respectively generically persistent)* if almost all its representations are constraint consistent (respectively persistent). Since constraint consistence and persistence have not been proved to be generic, we also say that a graph is *generically not constraint consistent (respectively not persistent)* if almost none of its representations are constraint consistent (respectively persistent). This double phrasing is necessary as there could exist some graphs that are neither generically constraint consistent nor generically not constraint consistent. The result above can now be restated for graphs.

Proposition 7.4. *For any definition of equilibrium consistent with Assumption 4.1, any graph whose largest out-degree is bounded by D is generically constraint consistent. Such a graph is thus generically persistent if it is generically rigid and generically not persistent otherwise.*

Similarly as in Section 4.6, one can use this proposition to prove that a graph whose largest out-degree is not larger than D is minimally persistent if and only if it is minimally rigid. This condition is however only proved valid for graphs with bounded out-degree, as there could exist persistent graphs with larger out-degrees, and to which the removal of any edge would result in a loss of persistence. Proving that a directed graph is minimally persistent if and only if it is minimally rigid and contains no vertex with an out-degree larger than D is equivalent to providing a positive answer to the following open question.

Open question 2. *For any definition of equilibrium consistent with Assumption 4.1, does every generically persistent graph contain a bounded*[3] *minimally rigid subgraph on all its vertices?*

A sufficient condition for a positive answer to this question would be that any generically persistent graph contains at least one generically persistent strict min D-subgraph on all its vertices.

The representations of acyclic graphs are another class of representations for which persistence does not depend on the particular definition of equilibrium. The proof of this relies on the following lemma.

Lemma 7.2. *Let $G'(V',E')$ be a directed graph obtained from a graph $G(V,E)$ by adding a vertex with no incoming edge and left by at least $\min(D,|V|)$ edges. For any definition of equilibrium consistent with Assumption 4.1, G' is generically persistent if and only if G is generically persistent, and generically not persistent if and only if G is generically not persistent.*

Proof. (sketch). Let p' be a non-degenerate generic representation of G' and p its restriction to G. We prove that p' is persistent if and only if p is persistent, and our result for graphs can then easily be deduced.

Suppose first that p is not persistent, i.e. there exists an equilibrium displacement δp that is not Euclidean. Call $n+1$ the vertex that was added to G

[3] Remember that a graph is bounded if its largest out-degree is not greater than D.

to obtain G'. Since no edge arrives at $n+1$, its displacement does not influence the fact that the other vertices are at equilibrium. Moreover, it follows from Assumption 4.1 that if all other agent's displacements are fixed, there exists a δp_{n+1} for which $n+1$ would be at equilibrium. So $\delta p' = [\delta p^T, \delta p_{n+1}^T]^T$ is an equilibrium displacement for p' and G'. This displacement is clearly not Euclidean as δp is not Euclidean, which implies that p' is not persistent.

Conversely, suppose that p is persistent. It follows from Corollary 6.1 that p' is rigid, so that we just need to prove its constraint consistence. Let $\delta p'$ be a displacement of p' and δp its restriction to p. Clearly, δp is also an equilibrium displacement. The persistence of p implies that it is Euclidean. One can prove that all constraints corresponding to edges leaving $n+1$ are then compatible, i.e., that the system (7.1) for $n+1$ admits a solution. Since $n+1$ is at equilibrium, all constraints in this system are thus satisfied, so that $\delta p'$ is admissible. The representation p' is thus constraint consistent and persistent. $\square$

Using this lemma and the fact that for any graph $G(V,E)$, no graph obtained from G by adding a vertex and connecting it to the vertices of G by less than $\min(D, |V|)$ is rigid, we obtain by induction the following result.

Proposition 7.5. *For any definition of equilibrium consistent with Assumption 4.1, an acyclic graph is generically persistent if every vertex i has an out-degree at least* $\min(D, i-1)$ *(where it is assumed that the vertices are labelled such that every edge leaves a vertex with a label larger than the vertex at which it arrives), and generically not persistent else.*

We believe that for any notion of equilibrium, a persistent graph should contain at most $f_D = \frac{1}{2}D(D+1)$ degrees of freedom. Intuitively, if it has more degrees of freedom, the set of equilibrium displacement has a larger dimension that the set of Euclidean displacement. For this argument to hold as a formal proof, one would need to prove that there exist non-trivial equilibrium displacements even when one linear constraint is added on each vertex displacement for each degree of freedom that it has in the graph. Besides, Theorem 4.6 should also hold for any reasonable definition of equilibrium. It states that every vertex with a positive number of degrees of freedom can be reached by a directed path from any other vertex in the graph, unless the latter vertex belongs to a closed subgraph on less than D vertices. By an argument similar to the one in Theorem 4.6, this would be proved if one could provide a positive answer to the following open question.

Open question 3. *For any definition of equilibrium consistent with Assumption 4.1, does every persistent graph contain a strict* $\min D$*-subgraph on all its vertices?*

Note that a positive answer to this question would also provide a positive answer to the Open Question 2. Finally, one could try to prove that this property holds for any rigid directed graph having at most f_D degrees of freedom (which is the case in $\Re^2$), which would be equivalent to providing a positive answer to the following question.

Open question 4. *Does the following hold? Let G be a rigid directed graph with at most f_D degrees of freedom. If the maximal out-degree is larger than D, then there is an edge leaving a vertex with an out-degree larger than D and whose removal results in a rigid graph.*

We have thus seen that two important classes of persistent graphs remain persistent independently of the particular condition of equilibrium used, as long as it is consistent with Assumption 4.1. We believe that this issue should be further investigated and that more results could be obtained independently of the assumptions used. These could include the potential results that we mention in our open questions, but also other results, for example on structural persistence.

Least-square solution of constraints system

Instead of requiring an agent to satisfy a maximal subsystem of (7.1), we require here its displacement to be a least-square solution of that system, which seems to be another natural behavior. A vertex i is thus at equilibrium for a displacement δp if and only if the following system of D equations holds:

$$\begin{pmatrix} (p_i - p_{j_1})^T \\ (p_i - p_{j_2})^T \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \end{pmatrix}^T \left(\begin{pmatrix} (p_i - p_{j_1})^T \\ (p_i - p_{j_2})^T \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \end{pmatrix} \delta p_i - \begin{pmatrix} (p_i - p_{j_1})^T \delta p_{j_1} \\ (p_i - p_{j_2})^T \delta p_{j_2} \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \delta p_{j_{d_i^+}} \end{pmatrix} \right) = 0, \tag{7.2}$$

Note that one could have normalized the lines before. Besides, the system (7.2) is equivalent to the system (7.1) when the latter admits a solution. For the sake of simplicity, let us restrict our attention to non-degenerate representations, that are representations for which (7.1) always admits a solution when it contains no more than D equations. For each i, let then

$$N_i = \begin{pmatrix} (p_i - p_{j_1})^T \\ (p_i - p_{j_2})^T \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \end{pmatrix}^T \qquad \text{if } d_i^+ > D$$

and the identity matrix of order d_i^+ if $d_i^+ \leq D$. Let then $N_{G,p}$ be the diagonal matrix whose blocks are $N_1, \ldots, N_{|V|}$. Supposing that the lines of the rigidity matrix are properly ordered (i.e., that the first d_1^+ lines correspond to edges leaving 1, that the next d_2^+ lines correspond to edges leaving 2, etc.), one can verify that δp is an equilibrium displacement for p and G if and only if $N_{G,p}R_{G,p}\delta p = 0$. As a result, p is constraint consistent if and only if $\mathrm{Ker} N_{G,p}R_{G,p} = \mathrm{Ker} R_{G,p}$, and persistent if and only if $\mathrm{Ker} N_{G,p}R_{G,p} = Eu_p$. Observe now that $N_{G,p}R_{G,p}$ contains $D\,|V|$ lines minus the number of degrees of freedom of the graphs. So if a graph representation is persistent, the graph contains at most f_D degrees of freedom for otherwise the dimension of $\mathrm{Ker} N_{G,p}R_{G,p}$ would be larger than f_D, the maximal dimension of Eu_p. Besides, remember that since the definition of equilibrium is consistent with Assumption 4.1, all results obtained above with generic assumptions can be applied.

We believe that the study of persistence and constraint consistence based on this particular least-square condition for equilibrium should be pushed further, and that the differences between these notions and those based on Assumption 4.2 should also be investigated. In particular, we would like to mention the two following open questions.

Open question 5. *Is the notion of persistence obtained by using the least-square definition of equilibrium a generic notion for graphs?*

Open question 6. *Is every rigid graph with less than f_D degrees of freedom persistent if we use the least-square definition of equilibrium?*

7.1.3 Non-rigid constraint consistent formations

This issues is partly related to the content of Sections 4.2 and 4.3.

The major part of this thesis has until here been devoted to persistence, and not to constraint consistence. In particular, we derived results on constraint consistence only when they were of immediate relevance to persistence or were providing a useful intuition. For example, Theorem 4.3 was proved for constraint consistence and then particularized to persistence as its proof only uses constraint consistence ideas, but Theorem 5.1 was not proved for structural constraint consistence as the ideas behind its proof use properties of Euclidean displacements and require thus the rigidity of the graph.

Our reason for doing so is that, in the absence of rigidity, constraint consistence as defined in Section 4.2 is of little use alone when analyzing the properties of autonomous agent system. It characterizes indeed the ability of a formation to satisfy all its distance constraints, *provided that the agents are initially sufficiently close from a reference position in which all constraints are satisfied.* In terms of definition, this is translated into the fact that only a neighborhood of representations is considered in Definition 4.2, and to the first-order only analysis in Section 4.3. When the graph is not rigid, the local character of these definitions can be a problem, because large deformations may occur even when all constraints are satisfied. Consider a two-dimensional formation whose corresponding graph is represented in Figure 7.2(a). This graph is clearly constraint consistent by Corollary 4.2, which means here that for any sufficiently small displacement of 2 and 3, agent 1 can find a position where its two constraints are satisfied. But, since 2 and 3 have no constraints, they can move to any position they want, and could on the long term get separated by a large distance, so that there would be no position at which 1 would satisfy its constraints. So, although the graph is constraint consistent, the satisfaction of all constraints is not guaranteed if large continuous displacements are admissible. This problem would however not happen if 2 was connected by a directed edge to 3, making the graph rigid. In that case, 2 would indeed always remain at a distance of 3 for which 1 can satisfy its two constraints.

It makes however sense to analyze the ability of formation to satisfy all its distance constraints, even if these constraints are not sufficient for maintaining the formation shape. First, one can imagine that an agent has to remain at constant distance from another in a formation, but that its relative position does not need to remain constant. An agent could for example need to be surrounded by protecting agents, whose positions need to remain constant with respect to the other protecting agents but not necessarily with respect to the rest of the formation. Second, suppose that two persistent formations need to be merged into a persistent meta-formation. This can be done by the addition of several directed edges linking the two formations (see [70, 71] for example). When some edges have been added but not all yet, the set of all agents constitutes a flexible formation containing two rigid sub-formations, corresponding to the two initial formations to be merged. At this intermediate stage, it is important to guarantee that all constraints can be satisfied, for otherwise the shape of the sub-formations could not be preserved. The same issue appears when performing different operations with one or several persistent formations. Note that in all those cases, one should also make sure that collisions between agents are avoided.

We therefore introduce a new notion, analogous to constraint consistence, but that takes the possibility of large continuous[4] deformations into account.

[4]One could also look at large possibly discontinuous deformations, and obtain a notion of global constraint consistence, analogous to global rigidity. We think here that the continuity of the agent displacements should be taken into account.

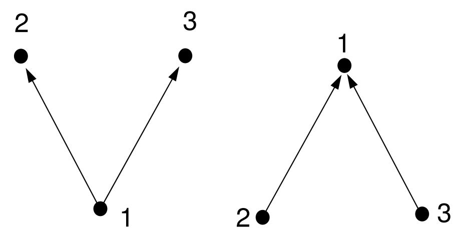

Figure 7.2: Example of a constraint consistent but not continuously constraint consistent representation (a). Agent 1 can indeed not satisfy its two constraints if the other agent move too far away one from each other. This problem cannot happen to the representation (b), which is continuously constraint consistent.

Definition 7.3. *Let $p : V \to \Re^D$ be a representation of a graph $G(V, E)$. The representation p is* continuously constraint consistent *if for every $\tilde{p} : V \times \Re^+ \to \Re^D : (i, t) \to \tilde{p}_i(t)$ Lipschitz continuous with respect to t and such that*

a) $\tilde{p}(0) = p$,
b) $\tilde{p}(t)$ is an equilibrium representation (with respect to the distance set induced by p and G) for every $t \geq 0$,

$\tilde{p}(t)$ is a realization of the distance set induced by p for every t, i.e. all constraints are satisfied for every t.

Intuitively, a representation is continuously constraint consistent if all constraints remain satisfied during any (Lipschitz) continuous deformation during which all agents are trying to satisfy all their constraints. The representation in Figure 7.2(a) is for example not continuously constraint consistent as 1 cannot satisfy its two constraints if the other agents get too far away from each other. The representation in Figure 7.2(b) on the other hand is continuously constraint consistent as, for any position of 1, the other agents can satisfy their constraints.

We believe that the study of continuous constraint consistence would be very interesting, not only for its practical relevance, but also because it has some surprising properties. In particular, it *is not a generic notion* as we show below, although there are some graphs that are generically continuously constraint consistent or generically not continuously constraint consistent. We now detail some first results obtained on continuous constraint consistence. Due again to the possible ambiguity that could arise, we do not omit the word "generically" in the rest of this section. Remember that a graph is generically P when the property P holds for almost all its representations, and generically not P when it holds for almost none of them.

Proposition 7.6. *Continuous constraint consistence is not a generic notion for graphs: there exist graphs for which the set of continuously constraint consistent representations and the set of not continuously constraint consistent representations both have positive measures.*

Proof. To prove this result, it suffices to exhibit a graph having a positive measure set of continuously constraint consistent representations, and a positive measure set of non continuously constraint consistent representations. Figure 7.3 shows two representations of such a graph. Observe that if the distances separating 4 from 2 and 3 are large as compared to those separating 1 from 2 and 3 as in Figure 7.3(a), the representation is continuously constraint consistent. Indeed, agent 4 is able to satisfy its two constraints for any displacement of 1, 2 and 3, as represented in Figure 7.3(b). On the other hand if the distances separating 4 from 2 and 3 are small as compared to those separating 1 from 2 and 3 as in Figure 7.3(c), the agents 2 and 3 can move to positions where it becomes impossible for 4 to satisfy its constraints, as represented in Figure 7.3(d). Such representation is thus not continuously constraint consistent. □

There are however graphs that are generically continuously constraint consistent or generically not continuously constraint consistent, i.e. that almost all their representations are or are not continuously constraint consistent. Obviously, continuous constraint consistence implies constraint consistence (We do not present a formal proof here as it would require proving the existence of a continuous path from p to any sufficiently close equilibrium representation). So, every generically non constraint consistent graph is generically not continuously constraint consistent, because almost none of its representations are continuously constraint consistent. We now prove that persistence implies continuous constraint consistence, and thus that every generically persistent graph is generically continuously constraint consistent.

Proposition 7.7. *Every persistent representation is continuously constraint consistent.*

Proof. (sketch). Consider a persistent representation p of a graph G, and a Lipschitz continuous displacement $\tilde{p}$ with a constant L. Since p is persistent, there is an $\epsilon > 0$ such that any equilibrium representation (with respect to p and G) in the bowl $B(p, \epsilon)$ is congruent to p. Besides, the Lipschitz continuity of $\tilde{p}$ implies that $\tilde{p} \in B(p, \epsilon)$ for every $t \in [0, \epsilon/L]$. So for every $t \in [0, \epsilon/L)$, $\tilde{p}(t)$ is congruent to p. Let now $p' = \tilde{p}(\epsilon/2L)$. Since p' can be obtained from p by a translation and rotation, for the same $\epsilon > 0$ as above, every equilibrium representation (with respect to p and G) in the bowl $B(p', \epsilon)$ is congruent to p. As a result, $\tilde{p}(t)$ is congruent to p for every $t \in [\epsilon/2L, 3\epsilon/2L)$. Repeating this argument shows that $\tilde{p}(t)$ is congruent to p for every $t \geq 0$, and thus also a realization of the distance set induced by p for every $t \geq 0$. As a consequence, p is continuously constraint consistent. □

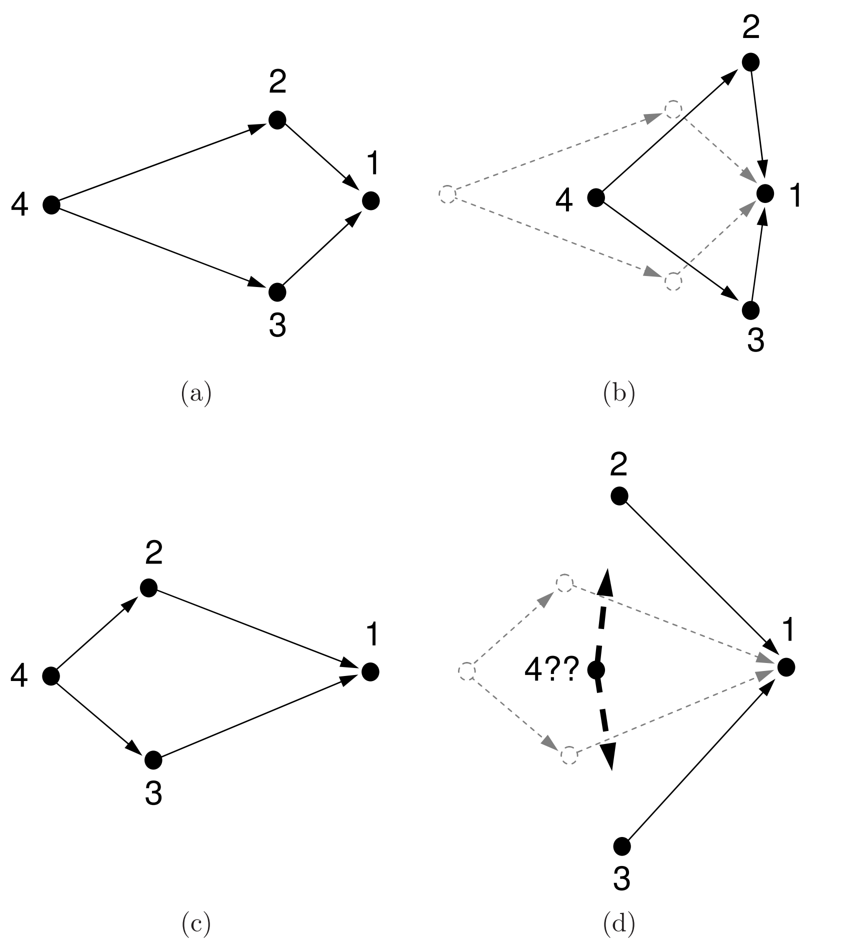

Figure 7.3: Example of graph admitting positive measure sets of continuously constraint consistent representations (a) and not continuously constraint consistent representations (c). Agents 2 and 3 can always satisfy their single constraints. In (a), the agent 4 can also always satisfy its two constraints, as shown in (b). This is not true for the representation (c), as shown in (d).

Note that some authors define rigidity with respect to continuous displacement, analogously to continuous constraint consistence. The notion obtained in that way is exactly equivalent to the rigidity defined in Section 3.1 [129]. The next corollary follows from the Proposition above and from Theorem 4.1.

Corollary 7.1. *All rigid graphs are either generically constraint consistent or generically not constraint consistent.*

Finally, let us mention that a graph in which no vertex has an out-degree larger than 1 is always continuously constraint consistent. Besides, any one-dimensional constraint consistent formation is continuously constraint consistent. This follows from the fact that rigidity is equivalent to connectivity in one-dimensional space, and thus that all connected components of a constraint consistent graph are persistent.

7.1.4 Global persistence

Persistence, constraint consistence and rigidity are local notions. They can be used to characterize the properties of a formation under the condition that the agent initial positions are in a neighborhood of some reference configuration in which all constraints are satisfied. As mentioned in Section 3.1, there exists a notion of global rigidity. A graph representation is *globally rigid* if all realizations of its induced distance set are congruent. We could make extend persistence and constraint consistence in the same way. A representation p of a graph G would be *globally constraint consistent* if every representation that is at equilibrium for the distance set induced by p and G is a realization of this distance set. A representation p of a graph G would be *globally persistent* if every representation at equilibrium for the distance set induced by p and G is congruent to p.

Intuitively these notions would be similar to constraint consistence and persistence, but allowing any (possibly remote) initial position for the agents. We believe that it would be interesting to investigate the properties of these notions and their exact meaning in terms of multi-agent formations. Note that global constraint consistence should not be mixed with continuous constraint consistence introduced in Section 7.1.3.

7.1.5 Stability and condition number of rigidity matrix

We have seen in Section 4.4 that persistence is a generic notion: all representations of a persistent graph are persistent, except those lying in some zero-measure set. But one can wonder if practical control problems arise for representations out of this set but close to it. Consider for example the representation

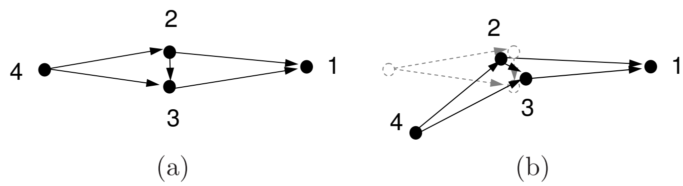

Figure 7.4: Example of persistent representation close to the set of non-rigid representations (a), which may cause stability problems. A variation in the positions of 2 and 3 can cause a proportionally large displacement of 4, as shown in (b).

in Figure 7.4(a). If the positions of 2 and 3 were the same, it would not be rigid, and therefore not persistent. Here, 2 and 3 have different positions so that the representation is persistent, but their positions are close to each other. Observe that a small variation of their positions, whether real or perceived by 4, may cause a very large variation of the position of 4. Mathematically, remember that the zero-measure set contains (but is not limited to) all the non-rigid representations of the graphs. The non-rigid representations are those for which the rank of the rigidity matrix is smaller than it is for other representations, due to the singularity of some sub-matrices. Since the determinant depends continuously on the matrix, this implies that it takes small values in the neighborhood of such representations, and that the corresponding matrices are thus ill-conditioned.

This topic has to the best of our knowledge not been formally studied yet for autonomous agent formations, but interesting information can be found in Laura Krick's thesis [79]. Results can also be obtained for undirected rigid graphs when the force transmission approach is used (see Section 3.2 and [124]). It can be shown that if a rigid representation is close to the set of non-rigid representations, the rigidity matrix becomes ill-conditioned, so that the application of small external forces can cause proportionally large internal forces.

7.1.6 Use of other geometric measures

We have analyzed the formation shape maintenance by means of unilateral distance constraints. One could however imagine mixing unilateral and bilateral distance constraints, as a formation can contain some agents with a limited visibility cone and some others with unlimited vision. This could be modelled by a graph containing both undirected[5] and directed edges. Moreover, inter-

[5]Remember that undirected edges are not equivalent to two edges having opposite directions, as the first models a bilateral distance constraints, for the satisfaction of which the agent can collaborate, while the second models two unilateral constraints on the same

agent distances are not the only geometrically relevant quantities that can be used. One can for example measure and constrain the angle between the perceived positions of two neighbors as in [45], the angle between the neighbor's perceived position and the agent velocity vector. It would be interesting to see if our persistence theory can directly be extended to those other types of constraints.

One could also take the physical extensions of the agents into account. Note that the position and orientation of an agent with some physical extension can be described by a single point in a higher-dimensional space. The position and orientation of a two-dimensional body can for example be represented by a point in a three-dimensional space.

7.1.7 Relaxed constraints and shape maintenance

The idea behind rigidity and persistence is to maintain exactly a formation shape by constraining sufficiently many inter-agent distances to prescribed values. Some applications may only require the formation shape to be approximately maintained or to be kept in some acceptable set. Suppose for example that two agents need to be within communication range and need to avoid collisions. The distance separating them does not need to be constant, but should remain in some interval.

A tool to keep a formation shape into a certain set is to use a partly flexible formation, as already mentioned in Section 7.1.3. Another one could be to constrain some inter-agent distance to intervals instead of unique values. One should then generalize the notion of rigidity, persistence and constraint consistence to take this new type of constraints into account. Such generalization could bee related to the notion of tensegrity presented in Section 3.6.

7.1.8 Unilateral and bilateral distance constraints

Persistence characterizes the efficacy of a structure of unilateral distance constraints to maintain a formation shape. If the constraints are bilateral instead of unilateral, that is, if the agents can collaborate in order to maintain the distance between them constant, then the graph of constraints is undirected, and rigidity theory can be applied as such (Assuming that control law make sure that all constraints are eventually satisfied, which is maybe not such a trivial assumption). One may however imagine hybrid formations in which some evolved agents are able to collaborate, while other can only handle unilateral constraints. Such formations would be modelled by graphs containing both directed and undirected edges. Rigidity would not be enough to characterize

distance.

their shape maintenance ability, but persistence could not be applied as it does not consider undirected edges. Remember indeed that an undirected edge is not equivalent to two directed edges with opposite directions. Some hybrid theory considering both unilateral and bilateral constraints would thus be needed.

7.2 Open questions in graph theory

7.2.1 Testing rigidity in three and higher dimensions

Open question 7. *How to characterize graphs that are rigid in a three or higher-dimensional space? (Section 3.2)*

Laman's theorem (Theorem 3.2) provides a necessary and sufficient condition to check if a graph is rigid in two-dimensional space. No equivalent characterization of rigid graphs is known for higher dimensions, and Theorem 3.1 only provides a necessary condition (Figure 3.9 shows a counterexample to its sufficiency). Another necessary condition for rigidity in $\Re^3$ is 3-connectivity. In $\Re^2$, it is automatically satisfied by any graph satisfying the condition of Laman. It has however recently been shown that 3-connectivity of a subgraph satisfying the necessary condition extending Laman's theorem (and even 4- and 5-connectivity) is not sufficient to determine rigidity [96]. Besides, several sufficient conditions are known, see [124] for example, or [130] for a survey.

In practice one can however easily determine whether a given structure is rigid or not, as once positions are given to vertices, rigidity can be determined by computing the rank of the rigidity matrix. Moreover, since the (infinitesimal) rigidity of one representation implies the rigidity of the graph [6, 7], a simple method to determine rigidity is then to take a random representation of the graph, and to compute the rigidity matrix rank. If the rank shows that the representation is rigid, then the graph is rigid. If it does not, then either the graph is not rigid, or the graph is rigid and the randomly chosen representation lies in the zero-measure set its non-rigid representation. The method can thus produce a certificate of generic rigidity for a graph, but not of the absence of generic rigidity, as it could produce a false negative with a probability 0. Another solution is to compute the maximal rank of the matrix parameterized by the vertices positions. This can indeed be done by checking the emptiness of several semi-algebraic sets, which can be done using Tarski's method [122] or more recent ones. The computational cost of such methods is however very high. In any case, these methods provide ways to check the rigidity of a graph, but they do not provide a characterization on which a further theory can be built, as Laman's theorem does.

7.2.2 Adding an edge to a persistent graph

Open question 8. *Given a pair of vertices in a persistent graph, is there a simple criterion to determine whether the connection of this pair of vertices leads to a persistent graph? Does this criterion lead to a polynomial time algorithm? (Sections 4.6 and 4.4)*

We have seen in Figures 4.5 and 4.8 that the addition of an edge to a persistent graph may result in a non-persistent graph, although this is obviously not always the case. In general, no non-trivial criterion to determine if the connection of a given ordered pair of vertices preserve the persistence is known. Such a criterion would however be useful to characterize maximally persistent graphs (see Open Question 10 in Section 7.2.4), and also to provide a polynomial-time criterion to check persistence (see Open Question 9 in Section 7.2.3). A weaker question would be to analyze this problem for two-dimensional space, for graphs with three degrees of freedom. In that case indeed, persistence is equivalent for a graph to the absence of a particular type of subgraph, as detailed in Section 6.3. A polynomial-time algorithm follows then from Theorem 6.4, but there exists maybe a simpler criterion that the one behind this algorithm.

Let us finally mention the following partial result:

Proposition 7.8. *Let $G(V,E)$ be a persistent graph and $i \in V$ be a vertex whose out-degree is smaller than D. The graph G' obtained by adding to G a directed edge leaving i is persistent.*

Proof. Call j the vertex at which arrives the edge added to G to obtain G'. Let $\Sigma(G)$ and $\Sigma(G')$ be the set of all strict $\min D$-subgraph of G and G' on all their vertices. Since the out-degree of i is no greater than D in both G and G', one can verify that each graph of $\Sigma(G')$ can be obtained by adding the edge (i,j) to some graph of $\Sigma(G')$. Since, by Theorem 4.5, all graphs in $\Sigma(G)$ are rigid, all graphs in $\Sigma(G')$ are also rigid, which by Theorem 4.5 implies that G' is persistent. □

7.2.3 Polynomial-time algorithm

Open question 9. *Is there a polynomial-time algorithm to check if a graph is persistent? (Sections 4.4 and 6.3)*

The criterion provided in Theorem 4.5 to check the persistence of a graph leads to an algorithm requiring checking the rigidity of a potentially exponentially growing number of subgraphs. Besides, checking rigidity may already be a problem in three dimensions. Polynomial-time algorithms exist however for acyclic graphs (Theorem 6.1), for minimally rigid graphs (Theorem 4.7), and for graphs with three degrees of freedom in $\Re^2$ (Theorem 6.4). A related question is to know whether the complexity of this algorithm ($O(|V|^2 |E|)$) is optimal.

This problem is equivalent to the existence of a polynomial-time algorithm deciding if the addition of an edge to a persistent graph lead to a persistent graph (see Section 7.2.2). Solving the latter problem with a polynomial-time algorithm to check persistence is indeed trivial. Suppose now that we can efficiently check if the addition of an edge preserves persistence, and let $G(V, E)$ be the graph of which we want to check the persistence. It suffices to check first the rigidity (and thus persistence) of one of its strict $\min D$-subgraph on all its vertices S. We can then add, one by one, all edges of $G \setminus S$, and check at each step if their addition preserve persistence. If all these additions preserve persistence, G is persistent, otherwise it is not. It follows indeed from Corollaries 4.2 and 4.3 that G is persistent if and only if S is rigid (and thus persistent) and every graph S' such that $S \subseteq S' \subseteq G$ is persistent. This method requires only $O(|E|)$ tests of persistence keeping by addition of an edge. It is important to note that, even if the order in which the additions are made has no importance, they should be added one after each other. The persistence of all graphs obtained by adding one single edge of $G \setminus S$ to S does indeed not imply the persistence of G obtained by adding all edges of $G \setminus S$ to S.

7.2.4 Maximal persistence

Open question 10. *Are there non-trivial maximally persistent graphs, and how to characterize them? (Sections 4.6 and 6.2.3)*

Unlike rigidity and many other graph notions, persistence is not preserved by edge addition. Adding an edge to a persistent graph can indeed lead to a non-persistent graph (see Figures 4.5 and 4.8). We say that a graph is *maximally persistent* if it is persistent and if the addition of any one or several edges would render it non-persistent. This definition can be restricted or not to graphs without double-edges. Note that by "non-trivial maximally persistent graphs", we exclude complete graph to which it is impossible to add one edge, independently of the persistence issue.

If double edges are allowed, the existence of non-trivial maximally persistent graphs can be established. If the number n of vertices is sufficiently large (larger than 3 in $\Re^2$), the graph K_n^* in which every pair of vertices is connected by two edges with opposite directions is not persistent, due to a too important redundance of the different constraints. This graph contains as subgraphs all graphs on n vertices, and thus all persistent graphs on n vertices. At least one of them is thus maximally persistent, for otherwise one could always increase the number of edges of a persistent graph while keeping it persistent, and eventually re-obtain K_n^*.

The issue of maximal persistence is related to the problem of keeping persistence by addition of an edge, presented in Section 7.2.2. Using an argument similar to the one used in Section 7.2.3 we now prove that if no single edge can be added to a persistent graph without losing persistence, then it is maximally persistent. In other words, there are no graphs from which one cannot obtain a persistent graph by adding *one* but one could obtain a persistent graph by adding *several* edges.

Proposition 7.9. *Let G be a persistent graph. If the addition of any single edge to G lead to a graph that is not persistent, then G is maximally persistent.*

Proof. Consider such a G and suppose, to obtain a contradiction, that there exists a persistent graph G' obtained from G by adding $k > 1$ edges. It follows from Proposition 7.8 that all vertices of G and thus of G' have an out-degree at least D. Therefore, G can be obtained from G' by removing edges leaving vertices with an out-degree larger than D, which by a repeated application of Corollary 4.3 implies that every graph G'' such that $G \subseteq G'' \subseteq G'$ is persistent. In particular, there is one persistent graph that can be obtained by adding one edge to G, contradicting our hypothesis. $\square$

Note that a similar argument proves that, if G_m and G_M are two persistent graphs on the same vertices and with the same degrees of freedom, and if $G_m \subseteq G_M$, then every graph G with $G_m \subseteq G \subseteq G_M$ is persistent. The result does however not hold if the graphs do not have the same degrees of freedom.

7.2.5 Redundant persistence and robustness

Open question 11. *Which persistent graphs remain persistent after deletion of any one edge? or any k edges? Can something be said about the probability of remaining persistent after deletion of k randomly selected edges?*

A redundantly rigid graph is a rigid graph that remains rigid by deletion of any one edge (see [46, 73] for example). A similar notion could be defined for persistent graphs, and extended to the removal of a larger number of edges. More generally, one could characterize the robustness of a persistent graph with respect to the loss of edges by computing its probability of remaining persistent after deletion of a certain number of randomly selected edges or vertices. Different probabilities could also be assigned to the edges, making for example sure that the edge between a first follower and a leader is never deleted. Observe that minimally persistent graphs would have the lowest degree of robustness, as the deletion of any one of their edges always lead to a loss of persistence.

7.2.6 Complexity of building 2D minimally persistent graphs

Open question 12. *How many vertex addition, edge splitting and/or edge reversal operations are needed to build a given minimally persistent graph? (Section 6.2)*

Theorem 6.3 proves that this construction can be done in at most $O(|V|^2)$ operations, but does not prove that this bound is optimal. The quadratic character comes from the cycle-reversal operation. Actually, if an edge belongs to a cycle that is reversed, it does not belong to any other cycle that has been or will be reversed. The sum over all cycles being reversed of their number of edges is thus at most $|V|$. But the cycle reversal operation also involves the reversal of some edges out of the cycles, and these edges may be used for several cycle reversals. Figure 7.5 shows a graph for which the method of Theorem 6.3 may require $O(n^2)$ edge reversals. This does however not prove that a quadratic number of operations is needed, as one could build this graph in $O(n)$ operations using a different sequence of vertex addition and edge splitting operations. Besides, our method for reversing cycles is not necessarily optimal.

A related issue is the one of determining the sequence of operations. Theorem 6.3 proves the existence of a sequence of operations building the graph, but does not explicitly provide them. It indeed relies on the undirected Henneberg sequence that build the corresponding minimally rigid graph. To the best of the our knowledge, no simple way of explicitly finding these operations is available.

7.2.7 Directions assignment

Open question 13. *Is it possible to assign directions to the edges of any rigid graph in such a way that the graph obtained is persistent? (Section 6.2.3)*

Proposition 6.6 provides an affirmative answer to this question for two-dimensional minimally rigid graphs. Bang-Jensen and Jordán obtain the same results for graphs with $2n-2$ edges [12]. Affirmative answers (in $\Re^2$) have also been provided in [52] for other particular classes of graphs, such as complete graphs and wheel-graphs.

It is known however that the directed graph obtained cannot always be made acyclic. Consider indeed a two-dimensional minimally rigid graph where all vertices have a degree at least 3, as for instance the complete bipartite graph $K_{3,3}$. It follows from Theorem 4.7 that assigning directions to the edges of this graph leads to a persistent graph if and only if no vertex obtains an out-degree larger than 2. As a result, all vertices must get an in-degree at least 1, and the graph contains thus necessarily a cycle.

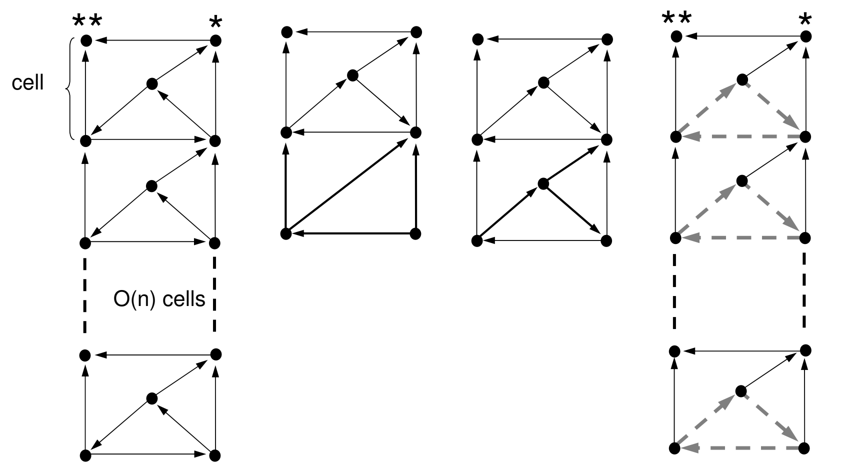

(a) (b) (c) (d)

Figure 7.5: Class of minimally persistent graphs (in $\Re^2$) whose construction using the method of Theorem 6.3 may require $O(n^2)$ edge reversals (a). Vertex addition (b) and edge splitting (c) operations are sufficient to build the graph represented in (d). That graph contains then $O(n)$ disconnected cycles, whose distance to the degree of freedom grows linearly. $O(n)$ cycle reversals are thus needed, and they each require on average $O(n)$ edge reversals. Note though that a different sequence of vertex addition and edge-splitting operations directly leads to the graph (d).

Chapter 8

Conclusions

In the first part of this thesis, we have proposed a notion of persistence, generalizing the notion of rigidity to directed graphs. This notion seeks to characterize the ability of a structure of unilateral distance constraints to maintain a formation shape. It is based on an analysis of the formations at equilibrium. We have further analyzed the issue of convergence to this equilibrium, and introduced the notion of structural persistence. We have then focussed on some particular classes of subgraphs, for which stronger results can be obtained.

Unlike rigidity which has a physical meaning, the notions that we have introduced are solely motivated by control and algorithmic issues. It would be interesting to know if there exist physical or mechanical systems to which they can be applied. The inherent asymmetry in the relation between vertices in directed graphs and the possible presence of cycles make us however think that such systems would be hard to conceive.

Writing this thesis was an opportunity to summarize, re-think, and reprocess the core of the research on persistence undertaken during the last three years. The proofs of many of the results appearing in this text have been simplified, and the writing has been an occasion to clearly redefine what persistence, constraint consistence and structural persistence are, and how they relate to each other.

We have clearly identified the main hypothesis on which persistence and constraint consistence relies, namely Assumption 4.2 that an agent is at equilibrium when it satisfies a maximal set of constraints. Looking back in time, we think that the importance of this assumption had not sufficiently been emphasized. It would therefore be relevant to analyze to which extent our result would be affected if we only use a weaker or a different hypothesis. We show some prospective results on this in Section 7.1.2.

Reformulating the notions also made clear that structural persistence is really about the possibility of converging to equilibrium, as it rules out the presence of a certain class of invariant sets containing no equilibrium points. This pins down the fact that a further analysis should be made to guarantee the convergence to an equilibrium. It is not totally clear if such analysis could be done as independently of the control laws as we want to. Ideally, we would like it to lead to a result such as "If the graph has that property, the shape can be maintained provided that the control laws used satisfy some basic assumptions. If the graph does not have that property, its shape cannot be guaranteed to be maintained, independently of the control law".

Finally, the reader will have noticed that persistence of a graph is not introduced at once, but trough different notion such as persistence of representation and infinitesimal persistence. This undirect definition was to be expected due to the close relation between persistence and rigidity. Although graph rigidity is a rather intuitive notion, its definition is indeed not aesthetical. It requires the definition of either infinitesimal rigidify or rigidity of a representation, the second being more intuitive, but the second allowing a simpler analysis. One could naturally think that a simpler definition should be found, bypassing the graph representations. Remember however that some properties similar to rigidity and persistence are not generic notion. We have seen for example in Section 7.1.3 that some graphs have positive measure sets of continuous constraint consistent representations and of non continuous constraint consistent representations. Tensegrity that we have briefly described in Section 3.6 is not a generic notion either. The fact that rigidity does only depend on the graph is thus a nontrivial result.

Part II

Consensus in Multi-Agent Systems with Distance Dependent Topologies

Chapter 9

Introduction

Consensus issues in multi-agent systems have attracted a lot of interest in the last years [5, 18, 49, 75, 76, 100, 103, 111, 112, 116, 117, 125]. In such multi-agent systems, every agent holds a value. Starting from some initial value, the agents communicate with each other, and tend to modify their values so that the difference with their neighbors decrease. These communications often take the form of agents averaging other agents values to update theirs. Such a system may reach or converge to *consensus*, that is, a situation where all agents hold the same value. The nature of the values on which consensus is sought, the way averages are performed and the communications that take place are part of each particular system definition.

Many properties of these systems depend on their communication topology, which is usually represented by a sequence of graphs. Vertices abstracting agents i and j are connected by a directed edge (j, i) in a graph G_t if the value of j is available to i at time t. Almost all sufficient conditions for convergence of consensus systems require for example some form of connectivity of the sequence of graphs $(G_t)_{t\geq 0}$.

We would like to distinguish between two of the possible approaches to analyze a system with a varying communication topology. A first approach is to prove results that are valid for any time-varying topology satisfying some basic assumptions. It corresponds thus to consider a random or exogenously given topology evolution. A second approach is to consider that the *topology does not depend on time but on the system state, which evolves with time.* Two robots or animals can for instance influence each other only when they are sufficiently close to each other, so that their communication abilities depend on their positions and not on time. But systems with state-dependent communications topologies are much more complex to analyze, so that the first approach is usually preferred in the literature. In this part of the thesis we will show that the classical results so obtained, although powerful, fail to explain central properties of a simple sorts of multi-agent systems. We will then analyze these systems taking explicitly the topology dependence into account, which will allow us to explain some phenomena that could otherwise not be understood.

In the rest of the introduction, we first present in Section 9.1 examples of systems in which consensus plays a central role. We then give in Section 9.2 convergence results for multi-agent systems, representative of those usually found in the literature. In Section 9.3, we expose phenomena appearing for some multi-agent systems that these results fail to explain because they do not take the dependence of the communication topology on the system state into account. Finally, we outline in Section 9.4 the rest of the second part of this thesis, in which we obtain stronger result for some particular multi-agent systems by explicitly using the topology dependence.

9.1 Examples of systems involving consensus

Before presenting examples of systems in which consensus plays a role, we would like to make a distinction between the design of consensus algorithms and the analysis of consensus seeking system models. In what we call a consensus algorithm, a communication protocol and a way of averaging neighbors' values are designed to ensure that all agents eventually reach consensus, possibly under some assumptions on the initial values. One may also require the value reached at consensus to be the exact average of all agents' initial values, a situation which is referred to as *average consensus*. These algorithms are usually relevant in the design of decentralized control laws. On the other hand, in a model of consensus seeking system the way agents interact with each other is supposed to model some real or imaginary behavior, and one asks whether the system eventually converges to an equilibrium and to a consensus. Such models naturally appear in the study of biological and social systems presenting self-organizing properties.

There is of course a close relationship between consensus algorithms and models, and they are often inspired from each other. The major difference is not a mathematical one. It lies in what someone studying those systems is allowed to do. In the design of a consensus algorithm, it is perfectly acceptable to modify the agent behavior in any way if it allows reaching a consensus in a better way, and if it is consistent with some possible physical or computational constraints of the system. In the analysis of a consensus model, the behavior of the agents is given and cannot be altered. Moreover it should be intuitively consistent with the behavior expected from the modelled agents, whether particles, human beings or animals. As a simple example, when modelling the way human opinions interact with each other, one should not expect a human being to compute the sum of two integers modulo 5 even though he/she would be capable of doing so. Such computations would however be totally acceptable in a consensus algorithm if they increase the capacity to reach a consensus.

In addition to the distinction between model and algorithm, we believe that three of the important characteristics of the consensus model are the nature of the value on which consensus is sought, the way the averages are computed, and the way the interaction topology is decided. We specify these characteristics for each of the systems presented below, and summarize them in Section 9.1.9.

9.1.1 Consensus in a wired sensor network

Type of system:	*algorithm for average consensus*
Nature of value:	*real number or vector*
Average:	*weighted linear average*
Interaction:	*fixed topology*

Consider a network of n sensors, each of them sensing a value x_i, and suppose that the average $\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$ needs to be computed. This could for example be the case of a set of decentralized sensors-controllers whose input variable should be the average temperature or pressure in a certain chemical process. But it can also be useful to minimize the measurement error risk, especially in the case of low-quality measurements. Suppose indeed that all sensors sense the same value y plus some (zero-mean) noise ϵ_i. The average sensed value

$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n}(y+\epsilon_i) = y + \sum_{i=1}^{n}\frac{\epsilon_i}{n}$$

is clearly a better measurement of y than any individual x_i, as its variance is much smaller than any individual measurement. Similarly, distributed sensors or robots may also need to compute an average reference value for calibration purpose or to synchronize their clock-time.

For robustness purpose or in order to avoid a possibly costly central controller, it might be desirable to compute these averages in a totally decentralized way. If the sensors are connected by wires allowing them to send their values to some other sensors, this can be done by repeated local averaging. Suppose for example that every sensor i is connected to the sensors $i-2$,$i-1$, $i+1$ and $i+2$ (with an appropriate convention to avoid border effects). If all sensors update their value synchronously by

$$x_i(t+1) = \frac{1}{5}\left(x_{i-2}(t) + x_{i-1}(t) + x_i(t) + x_{i+1}(t) + x_{i+2}(t)\right),$$

then they all converge (exponentially) to $\bar{x}$, the average of the initial values. The convergence follows indeed for example from Theorem 9.2 in Section 9.2, and one can verify that the average of all x_i is preserved at each iteration.

There exist of course other communication graphs for which the system converges to the average value. Optimal sets of graphs allowing convergence in finite time have for example been provided by Delvenne et al. [36]. One can also impose quantization of the exchanged information to represent the cost and limitation of digital communications, as in [49, 76]. Finally, the general issue of consensus has been extended to variables lying in space that are topologically different from $\Re^d$, see for example [116, 117] and the references therein.

9.1.2 Gossip algorithm for wireless sensor networks

Type of system:	*algorithm for average consensus*
Nature of value:	*real number or vector*
Average:	*pairwise average*
Interaction:	*randomized (possibly restricted)*

Some sensor networks rely on wireless connections, and communicate thus intermittently. This is especially the case when the sensors are embedded in moving agents who can communicate only when they are sufficiently close to each other. Provided that the communications are sufficiently frequent, it is however possible for such sensors to compute the average of the values they sense in a decentralized way, even if agents are never able to communicate all simultaneously. A simple way of achieving this is to use a so-called gossip algorithm based on pairwise interaction [22, 23]. When two agents i, j with values x_i and x_j interact, they both take as new value the average $\frac{1}{2}(x_i + x_j)$. One can verify that this operation preserves the global average. Moreover, all agent values converge to a common consensus equal to the average of the initial values, provided that the agents cannot be divided into small groups between which only a finite number of interactions occurs. This can be proved using for example Theorem 9.3 presented in Section 9.2.

9.1.3 Rendezvous problem for multi-agent formation with limited visibility

Type of system:	*algorithm for consensus*
Nature of value:	*position, real vector*
Average:	*usual average with nonlinearities to maintain connectivity*
Interaction:	*distance dependent (geometric graph)*

Consider a group of n autonomous agents with limited visibility. Typically, two agents can sense each other if they are distant by less than a certain radius R. The rendezvous problem consists in designing control laws for the agents allowing them to all gather in one single position. These control laws are expected to be the same for all agents and invariant under (at least) translations. It is usually further assumed that an agent cannot discriminate other agents, and must treat all of them in the same way. In variations of this problem, all agent eventually form a pre-specified shape by implicitly agreeing on a common reference position.

When the visibility radius R is finite, this problem can in general not be solved by deterministic control laws. Consider indeed a system of two agents initially separated by a distance larger than R, so that none of them can see any other agent. Deterministic control laws identical for all agents would dictate them exactly the same actions, so that the distance between them would remain constant, and they would never reach a common position.

The visibility relation is often modelled by a graph in which two agents are connected if and only if they can see each other. Different methods solve the rendezvous problem under the condition that this graph is initially connected [4, 27, 28, 32, 87–89, 120], depending on the assumed capacities and limitations of the agents. Lin et al. [87–89] proposes for example a simple law where each agent moves to a point in the interior of its neighbors positions' convex hull, in such a way that the distance to the neighbors remains smaller than R. This is a particular sort of consensus algorithm in which a common value is reached by averaging neighbors values. They prove that the visibility graph remains connected and that all agents eventually converge to a common position.

9.1.4 Deffuant's model of opinion dynamics

Type of system:	*model*
Nature of value:	*real number or vector*
Average:	*usual average*
Interaction:	*randomized with distance dependent constraint*

In this model introduced by Deffuant et al. [35], each agent i is a human being that has an opinion on some issue, represented by a real number $x_i \in \Re$. Pairs of agents meet randomly, and influence each other provided that their opinions are not too different. So, if the opinions x_i, x_j of two agents that meet do not differ by more than a pre-specified constant R they both take the average $\frac{1}{2}(x_i + x_j)$ as new opinion, otherwise they keep their previous opinions. This model is also often referred to as the Deffuant-Weisbuch model.

It can be proved, using for example Theorem 9.4 in Section 9.2, that for any sequence of meetings every agent opinion converges to a limiting value. These limiting values are generally not the same for all agents, and when they are different they are distant by at least R. It is experimentally observed that the distance between the limiting values are usually close to $2R$. [35, 128].

9.1.5 Krause's model of opinion dynamics

Type of system:	*model*
Nature of value:	*real number or vector*
Average:	*usual average*
Interaction:	*distance dependent*

This model of opinion dynamics was introduced by Krause [77, 78], and is also sometimes referred to as the Hegselmann-Krause model following an article by these two authors [60]. It presents many similarities with Deffuant's model, but its updates are deterministic and synchronized. Each agent i represents again a human being having an opinion represented by a real number. An agent finds another one sensible if their opinions differ by less than a certain constant R, and all agents update synchronously their opinions by computing the average opinions of the agents that they find sensible. More formally, for every i

$$x_i(t+1) = \frac{1}{|N_i(t)|} \sum_{j \in N_i(t)} x_j(t), \tag{9.1}$$

where $N_j(t) = \{j : |x_i(t) - x_j(t)| \leq R\}$. An example of this system evolution is shown in Figure 9.1. It has been proved [90] that all agent opinions converge to limiting values, and that when two different limiting values are not equal, they differ by at least R. This can also be seen as a consequence of Theorem 9.4 in Section 9.2. Experimental results show however that the distances between limiting opinions are closer to $2R$ than to R. The analysis of this model is the object of Chapter 11.

9.1.6 Vicsek's swarming model

Type of system:	*model*
Nature of value:	*direction, angle*
Average:	*direction average by vectorial sum*
Interaction:	*distance dependent, with distances depending on value history*

Vicsek's model was introduced in [127] to represent to behavior of flocks. It considers agents having a position x_i in the plane and a velocity v_i of fixed and common norm $||v_i|| = v \in \Re^+$. These agents all have a same visibility

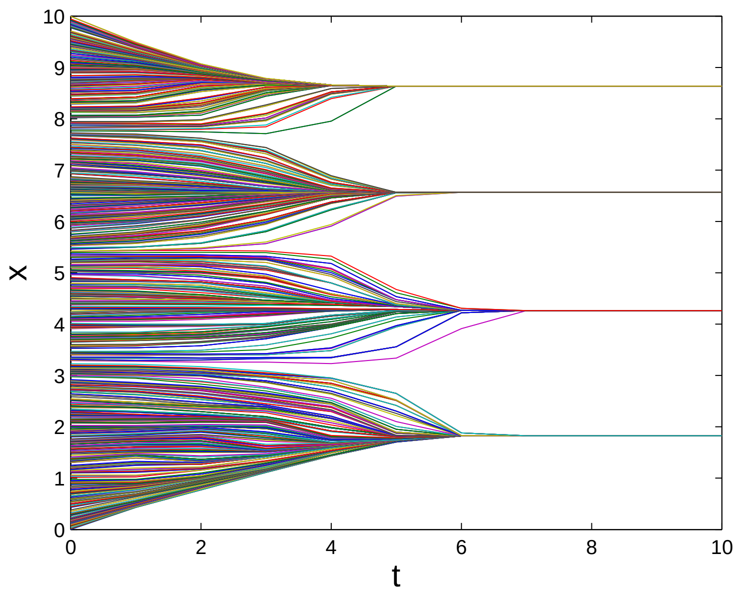

Figure 9.1: Evolution with time t of 1000 agent opinions initially randomly located on an interval of length 10, according to Krause model (9.1) with $R = 1$. The agent opinions converge to different clusters. Note that the distance between those clusters are significantly larger than the vision range 1.

radius R, and are said to be neighbors if they can see each other. Formally, i and j are neighbors at time t if $||x_i(t) - x_j(t)|| \leq R$. The agents update synchronously their positions and velocities. Each agent i takes as new position $x_i(t+1) = x_i(t) + v_i(t)$, and as new velocity the normalized sum of all its neighbors velocities:

$$w_i(t) = \sum_{j \in N_i(t)} v_i(t) \qquad v_i(t+1) = \frac{w_i(t)}{||w_i(t)||}, \tag{9.2}$$

where $N_i(t)$ is the set of neighbors of i at time t, which always includes i itself. Note that an additional rule has to be specified if $w_i = 0$. Vicsek et al. also add some random noise on the agent heading. They observe experimentally phase transitions in the limiting behavior of the system when the agent lie on a torus, depending on the level of noise and the density of agents. When the noise is important relatively to the density, agents tend to evolve randomly, in an uncorrelated way. For lower noise, strong correlations appear, and groups of agent moving in the same direction can be observed. For ever lower noise, all agent eventually move in a common direction.

This model has attracted a considerable attention after Jadbabaie et al. provided sufficient conditions for convergence to consensus for a linearized version [75]. These conditions are however generally not checkable in practice. In the linearized version, all velocities are $v_i = (v \cos \theta_i, v \sin \theta_i)$, and the agents update their angles θ_i by averaging their neighbors' angles

$$\theta_i(t+1) = \frac{1}{|N_i(t)|} \sum_{j \in N_i(t)} \theta_j(t).$$

This linearization introduces strong distortions. The linear average of 0 and 2π which represent the same angle is indeed π which is the exactly opposite angle. Jadbabaie et al. proved that the linearized system converges to consensus if there exists an infinite sequence of contiguous non-empty time-intervals of bounded length across each of which all agents are "linked together", that is, the union of all interaction graphs on the interval is a connected graph. This result was actually a particular case of earlier results by Tsitsiklis [125]. Using the symmetry of the neighborhood relation, it was proved later that for any initial condition, agents eventually gather in one or several groups (containing possibly one single agent), in which the distances between agents remain bounded and in which all agent directions tend to a common value (see for example [65,86,100]). This latter result does however not hold for the initial non-linear model. There exist indeed initial conditions for which the system does not converge [118], as the one in Figure 9.2. This exemplifies the importance of the difference between averaging real numbers and averaging elements over manifolds such as a torus.

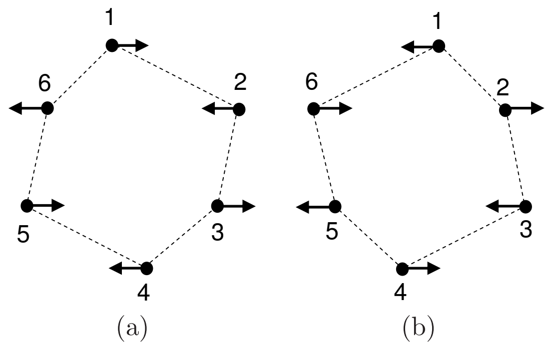

Figure 9.2: Example provided by Savkin [118] of cyclic behavior for Vicsek's nonlinear model. The six agents are initially located on the vertices of a regular hexagon whose edge length is between $1.5R$. Each agent i thus a neighbor of $i-1$ and $i+1$, as represented in (a). Agents 1, 3 and 5 have velocities exactly opposite to those of 2, 4 and 6. If the norm of these velocities is sufficiently small, the neighborhood graph remains unchanged at the next time step. Moreover, each agent's new velocity obtained by equation (9.2) is exactly opposite to its previous one (b), as $v_{i-1}+v_i+v_{i+1} = v_i - 2v_i = -v_i$. At the next time step, all agents return thus to their initial state (a).

9.1.7 Kuramoto-Sakaguchi oscillators

Type of system:	*model*
Nature of value:	*frequency, angular velocity*
Average:	*attraction on integral over time of value*
Interaction:	*all to all, strength depending on difference between the integrals over time of values*

The Kuramoto-Sakaguchi model was introduced in [115] to describe synchronization in systems of coupled oscillators, such as swarms of flashing fireflies or groups of pacemaker cells in the heart [81]. Each oscillator or agent has a phase $\theta_i(t)$, and a fixed natural frequency ω_i. The evolution of the phase is described by

$$\dot{\theta}_i(t) = \frac{K}{n}\sum_{j=1}^{n} \sin\left(\theta_j(t) - \theta_i(t) - \alpha\right),$$

with $|\alpha| \leq \frac{\pi}{2}$, and where $K > 0$ is the coupling strength. Observe that the system is invariant by addition of $2k\pi$ to any θ_i. The model extends the former Kuramoto model, where $\alpha = 0$. For sufficiently large values of K, a phenomenon called *partial entrainment* can be observed. It consists in the emergence of group of oscillators having bounded phase difference. The dependence of this phenomenon on the coupling strength K is sometimes complex.

For some configurations, increasing K may indeed result in the loss of partial entrainment. For more information ont he Kuramot-Sakaguchi model, we refer the reader to [34, 81, 113].

9.1.8 Aeyels-De Smet clustering model

Type of system:	*model*
Nature of value:	*real (velocity)*
Average:	*attraction on integral over time of value*
Interaction:	*all to all, strength depending on difference between the integrals over time of values*

The clustering model was introduced by Aeyels and De Smet [1, 2, 34]. It is similar to the Kuramoto-Sakaguchi model, but is defined on real values, with saturating interactions. As a result, it is more tractable, and allows for a more extended theoretical analysis. Each agent has a real value $x_i(t)$ and a natural velocity b_i. The evolution of the values is described by

$$\dot{x}_i(t) = b_i + \frac{K}{n} \sum_{j=1, j\neq i}^{N} f\left(x_j(t) - x_i(t)\right), \tag{9.3}$$

where f is odd, nondecreasing, and saturated, i.e., there exist $D \geq 0$ and $F \geq 0$ such that $x \geq D \Rightarrow f(x) = F$. Observe that the interactions between agents with significantly different velocities become constant after a certain time, unlike in the case of Kuramoto-Sakaguchi model where agent interactions keep varying due due the periodicity of the values.

This system produces clusters in both velocities $\dot{x}_i$ and positions x_i: The agents get partitioned in groups. In each group, all agent velocities converge to a same value, and the distance between any two agents remain bounded. The distance between agents in different groups grows unbounded, and their limiting velocities are different.

To be more precise, let $\bar{b}_S = \frac{1}{|S|}\sum_{i\in S} b_i$ be the average value of b on any set $S \subseteq \{1, \dots, n\}$ of agents, where $|S|$ denotes the cardinality of the set. For any initial condition, there exists a partition of the agents $\{1, \dots, n\} = N_1 \cup \dots \cup N_M$ ($N_a \cap N_b = varnothing$ if $a \neq b$) such that

- For every a, and every $i \in N_a$, there holds $\lim_{t\to\infty} \dot{x}_i = v_a$, with $v_a = \bar{b}_{N_a} + \frac{KF}{n}(\sum_{k>a}|N_k| - \sum_{k<a}|N_k|)$.
- For every a, and every $i, j \in N_a$, the difference between x_i and x_j is bounded: there exists a $B_{a,ij}$ such that $|x_i(t) - x_j(t)| < B_{a,ij}$ for all t.

- For every $a \neq c$ and every $i \in N_a, j \in N_c$, $\lim_{t\to\infty} |x_i(t) - x_j(t)| = \infty$ and there is a time after which $x_i > x_j$ if $a > c$.

Moreover, if $b_i < b_j$, then there is a time after which $x_i < x_j$. It suffices indeed to see that as long as $x_j \leq x_i$, there holds $\dot{x}_j - \dot{x}_i \geq b_j - b_i$.

Observe that the limiting velocities of the agents do not depend on the initial values, but only on the natural velocities b_i and on the composition of the groups. More surprisingly, it can be proved that the decomposition in groups is *uniquely determined* by the set of b_i and the value of the coupling parameter K. And, the asymptotic inter-agent distances in a group are independent of the initial conditions, provided that the considered agents have different natural velocities. The structure of clusters can moreover be determined using the following result.

Theorem 9.1 ([34]). *Let $\{b_i\}$ be a set of n natural velocities. The system 9.3 exhibits a clustering behavior with with sets $N_1, \ldots, N_m$ if and only if the two following conditions are satisfied:*

- *For all $k = 1, \ldots m-1$, there holds $\bar{b}_{N_{k+1}} - \bar{b}_{N_k} > \frac{KF}{n}(|N_{k+1}| + |N_k|)$.*
- *For all $k = 1, \ldots m-1$ and $N^* \subseteq N_k$, there holds $\bar{b}_{N^*} - \bar{b}_{N_k \setminus N^*} \leq \frac{KF}{n}\,|N_k|$.*

Note that the necessity of this condition can be obtained by computing the limiting average velocity in each group and ensuring that the velocity in N_{k+1} is larger than the velocity in N_k, and that for any partition of the group N_k in two subgroups, the average velocities in the two subgroups are identical. It has been proved that for any set of b_i and coupling parameter K, there is a unique partition $N_1 \cup \cdots \cup N_m$ satisfying the two conditions of Theorem 9.1. Moreover, the decomposition in groups follows a bifurcation scheme. For a large coupling parameter K, there is only one group. When K decreases, the number of groups increases, each new decomposition being (generically) obtained from the previous one by splitting one group into two subgroups. Using this, the composition of the different groups can be computed based on on K and on the natural velocities. Moreover, If all natural velocities are different, one can then also compute the relative positions of the agents in each group, the complexity of this last computation depending on the complexity of f.

Interestingly, in the experiments that we have conducted with random initial positions and natural velocities, the conditions of Theorem 9.1 were close to be tied when a non-trivial clustering behavior is observed. We think that this might be due to the small range of parameters leading to each particular decomposition in groups. In other words, if the conditions were not close to be tied, then some other conditions would not be satisfied, so that the considered decomposition would not be valid.

The asymptotic behavior of the clustering model is thus very well understood, partly thanks to the fact that the system tends to "forget" its initial condition. For more information on this model, on its extensions and on its applications, we refer the reader to [1, 2, 34].

9.1.9 Summary

The main characteristics of the systems presented above are summarized in the following table.

System	Type	Value	Average	Interaction
Wired sensors	algo.	real	weighted linear	fixed
Gossip	algo.	real	pairwise average	randomized, possibly restricted
Rendezvous	algo.	real	linear, nonlinearities to maintain connectivity	value dependent
Deffuant	model	real	linear	randomized, value constrained
Krause	model	real	linear	value dependent
Vicsek	model	angle	direction average by vectorial sum	dependent on value history
Jad. et al. linearized Vicsek	model	real	linear	dependent on value history
Kuramoto-Sakaguchi	model	angle	attraction on integral of value	all to all, strength dependent on difference between integrals of values
Aeyels - De Smet clustering	model	real	attraction on integral of value	all to all, strength dependent on difference between integrals of values

9.2 Representative convergence results

We now focus on systems where the agent values are real numbers, and where weighted averages are used. For such systems, we show some simple convergence results, first in discrete time and then in continuous time. Our goal is not to review the abundant literature on the domain (see [103] or [112] for surveys), but to show the type of convergence results that can be obtained for such systems.

9.2.1 Discrete-time consensus systems

We consider n agents, each of them having a real value $x_i(t)$, $i \in \{1, \ldots, n\}$. At each time step, every agent updates its value by taking a convex combination of other agents' values. The evolution of the agent values can thus be described by

$$x_i(t+1) = \sum_{j=1}^{n} a_{ij}(t) x_j(t).$$

where all $a_{ij}(t)$ are nonnegative, and $\sum_{j=1}^{n} a_{ij}(t) = 1$ holds for any i and t. Note the possible time-dependence of a_{ij}. This iteration can be rewritten in the more compact form

$$x(t+1) = A_t x(t), \tag{9.4}$$

where $[A_t]_{ij} = a_{ij}(t)$ and $[x(t)]_i = x_i(t)$. Every A_t is thus a stochastic matrix. Remember that a matrix A is stochastic if all its elements are nonnegative, and if $A\mathbf{1} = \mathbf{1}$ holds, where $\mathbf{1}$ is a vector of $\Re^n$ of which every entry is 1. Observe also that $\max_i x_i(t+1) \leq \max_i x_i(t)$ and $\min_i x_i(t+1) \geq \min_i x_i(t)$ holds for all t.

Almost all sufficient conditions for the convergence of (9.4) available in the literature do actually not depend on the particular values $a_{ij}(t)$ but only on the fact that they are positive or zero. It is therefore convenient to associate to every stochastic matrix A_t a directed graph $G_t(V, E_t)$ on n vertices, where j is connected to i by a directed edge (i, j) if $a_{ij}(t) > 0$. In view of (9.4), $(i, j) \in E_t$ means that the value $x_i(t+1)$ is influenced by the value $x_j(t)$. Self-loops are thus to be considered, and mean that the agent's new value is influenced by its previous value. Considering a sequence (G_t), we say that j *is connected to* i *over an interval* $[t_1, t_2]$ if there exists a sequence of T vertices $j = v_{t_1}, v_{t_1+1}, \ldots, v_{t_2} = i$ such that $(v_t, v_{t+1}) \in E_t$ holds for any $t \in [t_1, t_2]$, i.e., if $a_{v_{t+1}v_t}(t) > 0$ holds for every $t \in [t_1, t_2 - 1]$. In view of (9.4), the value of $x_i(t_2)$ is influenced by the value of $x_j(t_1)$ if and only if j is connected to i over $[t_1, t_2]$. Example of graphs associated to stochastic matrices and of paths over time intervals are provided in Figure 9.3.

$$A_0 = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad A_1 = \begin{pmatrix} 0.8 & 0 & 0.2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0.4 & 0.6 \\ 0 & 1 & 0 \end{pmatrix}$$

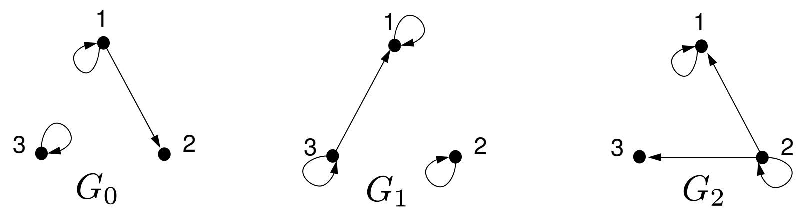

Figure 9.3: The graph G_0, G_1 and G_2 are associated to the matrices A_0, A_1 and A_2 respectively. Although 1 is connected to 3 by a directed path in neither of the three graphs, there is a path between 1 and 3 over $[0, 3]$, as $[A_0]_{12}$, $[A_1]_{22}$ and $[A_2]_{23}$ are positive. There is however no path over $[0, 3]$ between 2 and itself, nor between 3 and 2, although these paths would be present in $G_0 \cup G_1 \cup G_2$.

The convergence of $x(t)$ solution of (9.4) for time-invariant $A_t = A$ was studied by De Groot in 1974 [33]. A sufficient convergence condition for the time-varying case was then given in 1984 by Tsitsiklis [125]. A sufficent convergence condition was independently re-obtained by Jadbabaie et al. in 2003 [75], based on Wolfowitz's Theorem [131] on convergence of inhomogeneous products of stochastic matrices. Further proofs and results were then obtained, weakening or modifying these conditions [5, 25, 26, 86, 90, 100], or introducing additional variations such as time delays. We present here simple examples of such results, the proofs of which are largely inspired by those in [18].

Theorem 9.2. *Let $(A_t)_{t\geq 0}$ be a sequence of $n \times n$ stochastic matrices and $(G_t)_{t\geq 0}$ the associated sequence of graphs. For any initial condition $x(0) \in \Re^n$, there exists a $x^* \in \Re$ such that the sequence $x(t)$ solution of (9.4) converges exponentially to $x^* \mathbf{1}$ provided that the following two conditions are satisfied.*

a) All nonzero $a_{ij}(t)$ are larger than a certain $\alpha > 0$.
b) There exists a sequence of contiguous intervals of bounded length over each of which at least one vertex (possibly different for each interval) is connected to all others.

Proof. Let s be a vertex connected to all other vertices over the interval $[0, t^*]$, with $t^* \leq B$, and suppose first that $\min_i x_i(0) = 0$ and that $x_s(0) = 1$. Let then i be a vertex to which s is by hypothesis connected over $[0, t^*]$, i.e., there exists a sequence of vertices $s = v_0, v_1, \ldots, v_{t^*} = i$ such that $a_{v_{t+1}v_t}(t)$ is positive for every $t \in [0, t^* - 1]$, and thus no smaller than α by the hypothesis (a). Since all $x_i(t)$ are nonnegative, there holds

$$x_{v_1}(1) = \sum_{j=1}^{n} a_{v_1 j}(0) x_s(0) \geq a_{v_1 s}(0) x_s(0) \geq \alpha 1 = \alpha \tag{9.5}$$

By the same argument we obtain $x_{v_2}(2) \geq \alpha x_{v_1}(1) \geq \alpha^2$, $x_{v_3}(3) \geq \alpha^3$, and eventually $x_i(t^*) = x_{v_{t^*}}(t^*) \geq \alpha^{t^*} \geq \alpha^B$.

By two appropriate translations and scalings, this first result implies that for a general $x(0)$ there holds

$$\begin{array}{ll} x_s(0) - \min_i x_i(t^*) & \leq (1 - \alpha^B)\left(x_s(0) - \min_i x_i(0)\right), \\ \max_i x_i(t^*) - x_s(0) & \leq (1 - \alpha^B)\left(\max_i x_i(0) - x_s(0)\right), \end{array}$$

and thus

$$\max_i x_i(t^*) - \min_i x_i(t^*) \leq (1 - \alpha^B)\left(\max_i x_i(0) - \min_i x_i(0)\right).$$

Since there always holds $[\min_i x_i(t+1), \max_i x_i(t+1)] \subseteq [\min_i x_i(t), \max_i x_i(t)]$, a repeated application of this result based on the hypothesis (b) implies the convergence of all $x_i(t)$ to some $x^* \in \Re$. □

The following corollary provides a slightly weaker result, but its conditions are often simpler to check. It is proved by verifying that its three conditions imply those of Theorem 9.2.

Corollary 9.1. *Let $(A_t)_{t\geq 0}$ be a sequence of $n \times n$ stochastic matrices and $(G_t)_{t\geq 0}$ the associate sequence of graphs. For any initial condition $x(0) \in \Re^n$, there exists a $x^* \in \Re$ such that the sequence $x(t)$ solution of (9.4) converges exponentially to $x^* \mathbf{1}$ provided that the following three conditions are satisfied.*

a) All nonzero $a_{ij}(t)$ are larger than a certain $\alpha > 0$.
b) All diagonal elements $a_{ii}(t)$ are positive (and larger than α).
c) There exists a B such that for any t, the graph $\bigcup_{t'=t}^{t+B} G_{t'}$ contains a directed spanning tree, i.e., one of its vertices is connected to all others by a directed path.

Among the systems presented in Section 9.1, many have symmetric interaction or communication topologies. In such systems, an agent cannot influence another agent without being itself influenced by it, although not necessarily with the same strength. The following result shows that in the case of symmetric interaction topologies, no upper bound on the interval across which all agents communicate is required. It was proved by Li et al. [86] in a particular case, and in a more general case in [18, 65, 90, 100].

Theorem 9.3. *Let $(A_t)_{t\geq 0}$ be a sequence of $n \times n$ stochastic matrices and $(G_t)_{t\geq 0}$ the associated sequence of graphs. If every G_t is symmetric, then for any initial condition $x(0) \in \Re^n$, there exists a $x^* \in \Re$ such that the sequence $x(t)$ solution of (9.4) converges to $x^* \mathbf{1}$ provided that the following three conditions are satisfied.*

a) All nonzero $a_{ij}(t)$ are larger than a certain $\alpha > 0$.
b) All diagonal elements $a_{ii}(t)$ are positive (and larger than α).
c) The graph $\bigcap_{t\geq 0} \bigcup_{t'\geq t} G_{t'}$ is connected. Equivalently, the graph $\bigcup_{t'\geq t} G_{t'}$ is connected for each t.

Proof. We suppose first that $\max_i x_i(0) = 1$ and $\min_i x_i(0) = 0$, and show that there exists a time t^* at which $\min_i x_i(t^*) \geq \alpha^{n-1}$, while $\max_i x_i(t^*) \leq 1$ trivially holds. By appropriate translations and scaling, the repetition of this argument implies the existence of an infinite sequence $t_1 < t_2 < \dots$ such that

$$(\max_i x_i(t_k) - \min_i x_i(t_k)) \leq (1 - \alpha^{n-1})(\max_i x_i(t_{k+1}) - \min_i x_i(t_{k+1}))$$

holds for each t_k. Since in addition there always holds $[\min_i x_i(t+1), \max_i x_i(t+1)] \subseteq [\min_i x_i(t), \max_i x_i(t)]$, this is sufficient to prove the convergence of all x_i to some $x^* \in \Re$.

Let M be an index for which $x_M(0) = \max_i x_i(0) = 1$, and let $S_0 = \{M\}$ be a set of indices. For any $t \geq 1$, define S_t by adding to S_{t-1} all indices i for which there exists a $j \in S_{t-1}$ such that $(i,j) \in E_{t-1}$, i.e., all i for which the value of $x_i(t)$ is influenced by the value of $x_j(t-1)$ for some $j \in S_{t-1}$. The set S_t clearly increases with t. Let t^* be the time at which it reaches its largest size. S_{t^*} contains all indices $1, \dots, n$, for otherwise there would be a group of vertices that are connected to none of those in S_{t^*} for any time $t' \geq t^*$. As a result $\bigcup_{t'\geq t^*} G_{t'}$ would not be connected, contradicting our hypothesis (c).

We now prove by induction on t that for any $t \geq 0$ there holds $x_i(t) \geq \alpha^{|S_t|-1}$ for all $i \in S_t$, where $|S_t|$ is the number of elements in S_t. The property is clearly true for S_0 which contains only M, as $x_M(0) = 1$. Suppose now that it holds for some t. If $S_{t+1} = S_t$, then for all $i \in S_{t+1}$, the value $x_i(t+1)$ is a convex combination of values $x_j(t)$ where $j \in S_t$, and is thus at least as large as $\alpha^{|S_t|-1}$. If S_{t+1} and S_t are different, then for every $i \in S_{t+1}$, $x_i(t+1)$ is a convex combination of different $x_j(t) \in [0,1]$, with at least one j belonging to S_t. Indeed, if $i \in S_{t+1} \setminus S_t$ it is by construction connected to at least one $j \in S_t$ over $[t, t+1]$, and every $i \in S_t$ is by the hypothesis (b) always connected to itself over any interval. Since all (positive) coefficients $a_{ij}(t)$ are by hypothesis (a) lower-bounded by α, this together with the induction hypothesis implies that $x_i(t) \geq \alpha^{|S_t|-1+1} \geq \alpha^{|S_{t+1}|-1}$ holds for all $i \in S_{t+1}$. As a result, we have

$x_i(t^*) \geq \alpha^{n-1}$ for all i, which is sufficient to prove our result as explained above. □

As compared to the conditions of Corollary 9.1, the main advantage of the conditions of Theorem 9.3 is that the interval across which all agents need to be influenced by a same agent does not need to be bounded. Although this may seem a minor difference, it allows proving the convergence of the system even in the absence of consensus, under the sole conditions (a) and (b). This proof appears in [65, 90].

Theorem 9.4. *Let $(A_t)_{t\geq 0}$ be a sequence of $n \times n$ stochastic matrices and $(G_t)_{t\geq 0}$ the associated sequence of graphs. If every G_t is symmetric, then for any initial condition $x(0) \in \Re^n$, the sequence $x(t)$ solution of (9.4) converges to a limit vector provided that the following two conditions are satisfied.*

a) All nonzero $a_{ij}(t)$ are larger than a certain $\alpha > 0$.
b) All diagonal elements $a_{ii}(t) > 0$ are positive (and larger than α).

Moreover, if i and j belong to the same connected component of $\bigcap_{t\geq 0} \bigcup_{t'\geq t} G_{t'}$, their value converge to the same limit: $\lim_{t\to\infty} x_i(t) = \lim_{t\to\infty} x_j(t)$.

Proof. Suppose that a system satisfies conditions (a) and (b). If $G^* = \cap_t \left(\cup_{t'\geq t} G_t\right)$ is connected, Theorem 9.3 applies. If not, consider one of its connected component $G_A^*(V_A^*, E_A^*)$. There is a time t^* after which no vertex in V_A^* is connected to any vertex of $V \setminus V_A^*$ in any G_t, $t \geq t^*$. Otherwise, the finite number of vertices implies indeed that there would be at least one edge connecting V_A^* to $V \setminus V_A^*$ in $G^* \cap_t (\cup_{t'\geq t} G_t)$, contradicting the fact that G_A^* is a connected component of G^*. This means that after that time t^*, the agents corresponding to V_A^* do not interact in any way with the other agents. They constitute thus an isolated subsystem, to which Theorem 9.3 can be applied. It indeed trivially satisfies conditions (a) and (b), and also satisfies (c) as one can verify that the graph needing to be connected is actually G_A^*, which is connected by definition. There exists thus a $x_A^* \in \Re$ such that $\lim_{t\to\infty} x_i(t) = x_A^*$ holds for any $i \in V_A$. Since this is valid for every connected component of G^*, and since every vertex belongs to one connected component, it follows that every x_i converges to a limiting value, and that this value is the same for all i within a same connected component of G^*. □

This corollary cannot be extended in general to systems in which the communication topology is not symmetric. There exist indeed simple examples of systems that do not converge although they satisfy conditions (a) and (b) of Corollary 9.1. Consider for example three agents 1, 2 and 3, with initially $x_1(0), x_2(0) < -1$ and $x_3(0) > 1$. By applying a sufficient number iterations

$$x_1(t+1) = \frac{1}{2}\left(x_1(t) + x_3(t)\right), \qquad x_2(t+1) = x_2(t), \qquad x_3(t+1) = x_3(t),$$

one drives the system into a state where $x_1(t), x_3(t) > 1$ and $x_2(t) < -1$. By a similar process, one can then obtain after a sufficiently large number of iterations $x_1(t) > 1$, $x_2(t), x_3(t) < -1$, and then $x_1(t), x_2(t) > 1$ and $x_3(t) < 1$, a condition opposite to the initial one. By repeating this process indefinitely, one builds a sequence of interactions for which the system never converges.

Finally, note that all results presented in this section give conditions for the convergence of the solution of (9.4) for any initial condition. The same conditions imply thus each time the convergence of the inhomogeneous matrix product $A_t A_{t-1} \dots A_1$. Moreover, when the convergence to a consensus is guaranteed, then the product $A_t A_{t-1} \dots A_1$ converges to a rank one matrix of the form $y^* \mathbf{1}^T$, with $y^* \in \Re^n$.

9.2.2 Continuous-time consensus systems

The consensus issues can also be considered in continuous time. For n agents having real values $x_i(t)$, typical systems have the form

$$\dot{x}_i(t) = \sum_{j=1}^{n} a_{ij}(t) \left(x_j(t) - x_i(t)\right), \tag{9.6}$$

where all $a_{ij}(t)$ are nonnegative. This can be rewritten under the more compact form

$$\dot{x}(t) = A(t)x(t), \tag{9.7}$$

where $A(t)$ has zero rows sums, and nonnegative off-diagonal elements. In case $A(t)$ has zero column sums, observe that the average value $\bar{x}$ of x is preserved over time, as $\frac{d}{dt}\bar{x} = \frac{d}{dt}\mathbf{1}^T x = \mathbf{1}^T A(t)x(t) = 0$. The convergence of the system 9.7 has been the object of several studies. We present here without proof a result from [99].

Theorem 9.5 (Moreau [99])**.** *Consider the linear system (9.7), and assume that the system matrix is bounded and piecewise continuous (with t), and has zero row sums. The system converges exponentially, and there is a $x^* \in \Re$ such that $\lim_{t\to\infty} x_i(t) = x^*$ holds for every i, provided that there is a $k \in \{1, \dots, n\}$, a threshold value $\delta > 0$ and an interval length $T > 0$ such that for all $t \in \Re$, the digraph on n vertices obtained from*

$$A_t^* = \int_t^{t+T} A(s)ds,$$

by connecting j to i by the directed edge (j, i) if $[A_t^*]_{ij} \geq \delta$, has the property that all nodes may be reached from the node k.

Unlike in the discrete time case, no result is known guaranteeing convergence in the absence of consensus if the connectivity condition is not satisfied, even if the interactions are symmetric. A reason for this is that the system may not necessarily be separated into independent subsystems when this condition is not satisfied.

9.3 State-dependent communication topology

The results in Section 9.3 guarantee convergence of systems and matrix products under very mild assumptions. One of their main limitations is that they give conditions on the matrix sequence or evolution, but that they do not take the way those matrices are built into account. By ignoring this information one may indeed fail to see important characteristics in the behavior of the system. Moreover, it may be difficult for such systems to check whether the sufficient conditions on the matrices are satisfied since the sequence is not given a priori but depends on the evolution of the system state.

To show the importance of this limitation, we consider in the second part of this works two particularly simple systems involving state-depending communication topologies. We show that the generic results of Section 9.2 only allow proving the convergence of one of them, and fail to explain some peculiar behavior. We then analyze these two systems using explicitly the dependence of their communication topology on the system state, and show how stronger results can be obtained.

The first system is Krause's opinion dynamics model presented in Section 9.1.5, with $R = 1$. There are n agents and every agent $i \in \{1, \dots, n\}$ has a real value x_i representing its opinion. Each agent updates its value according to

$$x_i(t+1) = \frac{\sum_{j:|x_i(t)-x_j(t)|<1} x_j(t)}{|\{j : |x_i(t) - x_j(t)| < 1\}|}, \tag{9.8}$$

where $|\{j : |x_i(t) - x_j(t)| < 1\}|$ is the number of elements in the set $\{j : |x_i(t) - x_j(t)| < 1\}$. Two agents i, j for which $|x_i(t) - x_j(t)| < 1$ are said to be *neighbors* or *connected*. It follows from this definition that an agent is always its own neighbor. So in this system, each agent updates its value by computing the average values of its neighbors, the agents that are distant from it by less than 1.

As second system, we consider a continuous time variation of Krause's model, in which the evolution of the agent values is described by

$$\dot{x}_i(t) = \sum_{j:|x_i(t)-x_j(t)|<1} \left(x_j(t) - x_i(t)\right). \tag{9.9}$$

Here, each agent is continuously attracted by the values of those agents that are distant from it by less than 1. Note that the model (9.8) does not correspond to a direct discretization of (9.9), as the influences between agents are normalized by the number of neighbors in the former and not in the latter. In the sequel, we refer to the agent value x_i indistinctly as their positions, values or opinions.

Observe that the discrete-time system (9.8) can be rewritten $x(t+1) = A(x(t))x(t)$, where $[A(x(t))]_{ij} = \frac{1}{|\{j:|x_i(t)-x_j(t)|<1\}|}$ if i and j are neighbors, and 0 otherwise. These matrices $A(x(t))$ are stochastic, and their positive elements are bounded from below by $\frac{1}{n}$. Since the neighborhood relation is symmetric, $[A(x(t))]_{ij} > 0$ implies that $[A(x(t))]_{ji} > 0$. The conditions of Theorem 9.4 are thus satisfied and it follows that the system (9.8) always converges to a limiting vector $x^* \in \Re^n$. Moreover, $x_i^* \neq x_j^*$ can only hold if there is a time after which $[A(x)]_{ij}$ is always zero, that is, a time after which i and j are never neighbors. After this time, x_i and x_j must thus be separated by at least 1, and so are then their limiting value x_i^* and x_j^*. The results in Section 9.2 implies therefore that all opinions converge to clusters of opinions, that are separated from each other by at least 1.

It has however been experimentally observed that opinions initially uniformly distributed on an interval tend to converge to clusters of opinions separated by a distance larger than (but close to) 2 as shown in Figure 9.4(a). A similar phenomenon appears for the continuous-time model (9.9) as shown in Fig.9.4(b). But no explanation of this phenomenon nor any nontrivial lower bound on the inter-cluster distance have been provided so far. Moreover, for the continuous-time model, the results in Section 9.2 do not even imply convergence, as there is no continuous-time equivalent to Theorem 9.4.

Interestingly, inter-cluster distances significantly larger than 1 have also been observed for the stochastic version of the discrete-time model by Deffuant et al., presented in Section 9.1.4 [35]. The behavior of this system - and more particularly the final positions of the clusters - can also be studied by approximating the evolution of the opinion density using a partial differential equation [13, 14]. Besides, different systems involving discrete or continuous time, opinions and agent density and opinions have also been proposed modelling the same phenomenon [53, 126]. For a survey, see for example [92].

In the next chapters we study the convergence properties of the models (9.9) and (9.8). We introduce a particular notion of equilibrium stability, and prove that an equilibrium is stable if and only if a certain nontrivial lower bound on the distance between clusters holds. Although not proved yet, it is experimentally observed and conform to the intuition that the probability of converging to a stable equilibrium increases with the number of agents. To better understand the system behavior for large numbers of agents, we introduce and

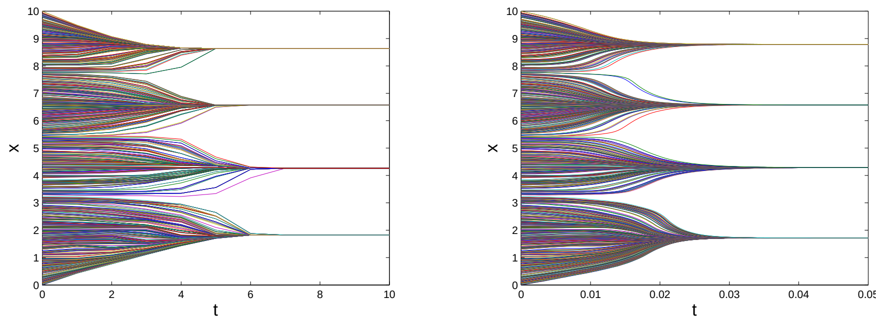

(a) (b)

Figure 9.4: Evolution with time t of 1000 agent opinions initially randomly located on an interval of length 10, according to the discrete-time model of Krause (9.8) in (a) and to its continuous-time counterpart (9.9) in (b). In both cases, all opinions converge to different clusters that are separated by distances significantly larger than the vision range 1. Note that the time-scale of the two systems are not comparable, as the typical velocity in (9.9) is proportional to the number of agents.

study versions of the models allowing continuous opinion distributions. We give partial convergence results for these, and provide lower bound on the distance between clusters at the equilibrium under some continuity assumptions.

Our models for agent continuum are obtained by indexing the agents by a real number instead of an integer. This allows for a more detailed formal analysis than with the earlier model, which are usually defined using agent density functions. Nevertheless, our model for discrete time and agent continuum, which we introduced in [19] is equivalent to the so-called "discrete-time density based HK model" proposed independently in [92], using a density-based formalism. The latter system is similar to a continuous-time system [53] also using the density formalism, and can finally be viewed as the limit when the number of discrete opinions tends to infinity of the "interactive Markov chain model of Lorenz" introduced in [91]. In that model, there is a continuous distribution of agents, but the opinions are discrete.

9.4 Outline of part II

The discrete time model (9.9) is more standard and has been studied more than the continuous-time model (9.8). The analysis of the latter is however simpler in several aspects, and several results can be re-used for the discrete-time case. We therefore begin by studying the continuous-time model in Chapter 10 before studying the discrete-time model in Chapter 11. In both chapters, we first consider the model for discrete agents, for which we prove convergence properties and give necessary and sufficient conditions for stability. We then introduce the systems on a continuum of agents, for which we give partial convergence results, and provide lower bounds on the distance between clusters at equilibrium. In Chapter 11 we also explore more formally the relation between the discrete-time systems (9.8) defined for discrete agents and for an agent continuum. The main results obtained in these two chapters are summarized in Table 9.1. In Chapter 12 we consider possible extensions of this analysis to systems in higher dimensional spaces, and to distance-dependent influences. The second part of this thesis is closed by the concluding remarks and open questions in Chapter 13.

The work presented in Chapters 10 and 11 is also available in [19, 20]. The order preservation results of Chapter 12 also appear in [62].

agents	$\{1, \ldots, n\}$	$[0, 1]$
$\dot{x} = \ldots$	*Section 10.1* - Convergence to an equilibrium - Equilibrium stability $\Leftrightarrow$ lower bound on inter-cluster distances - Conjecture that convergence to stable equilibrium for large n	*Section 10.2* - Convergence to a *set* of equilibrium - Equilibrium stability $\Rightarrow$ lower bound on inter-cluster distances - Convergence $\Rightarrow$ lower bound on inter-cluster distances
$x(t{+}1) = \ldots$	*Section 11.1* - Convergence in *finite* time to an equilibrium - Equilibrium stability $\Leftrightarrow$ lower bound on inter-cluster distances - Conjecture that convergence to stable equilibrium for large n	*Section 11.2* - Convergence to a *set* of equilibrium - Equilibrium stability $\Rightarrow$ lower bound on inter-cluster distances - Convergence $\Rightarrow$ lower bound on inter-cluster distances

Table 9.1: Summary of the main results obtained for each variation of the model. The mentioned lower bound on each inter-cluster distance is $1 + \frac{\min(W_A, W_B)}{\max(W_A, W_B)}$, where W_A and W_B are the weights of the two clusters considered, i.e. their number of agents.

Chapter 10

Continuous-time Opinion Dynamics

10.1 Finite number of discrete agents

Remember that the generic results of Section 9.2 do not prove the convergence of the continuous-time system. We propose here a simple convergence proof relying on the particular way in which the topology changes. Before doing so, let us observe four properties of the system.

First, the order between the agent opinions is preserved: If $x_i(t) > x_j(t)$ holds for some time t, it holds for any further time. Since the opinions evolve continuously with time, if an order inversion could happen we would necessarily first have $x_i - x_j < 1$. Call then $N_i(t)$ the set of agents connected to i and not to j, $N_j(t)$ the set of those connected to j and not to i, and $N_{ij}(t)$ the set of those connected to both i and j. As long as $x_i(t) - x_j(t) < 1$, there holds $x_{k_1}(t) > x_{k_2}(t) > x_{k_3}(t)$ for any $k_1 \in N_i(t), k_2 \in N_{ij}(t), k_3 \in N_j(t)$. It follows from (9.9) that

$$\dot{x}_i(t) \geq \sum_{k \in N_{ij}(t)} x_k(t) - x_i(t), \qquad \text{and} \qquad \dot{x}_j(t) \leq \sum_{k \in N_{ij}(t)} x_k(t) - x_j(t).$$

Therefore we have

$$\dot{x}_i(t) - \dot{x}_j(t) \geq |N_{ij}(t)| \left(x_j(t) - x_i(t)\right) \geq - |N_{ij}(t)| \left(x_i(t) - x_j(t)\right),$$

which implies that $x_i(t) - x_j(t)$ cannot decrease faster than $e^{-|N_{ij}(t)|t}$ when it is smaller than 1, and is therefore always positive. Besides, if initially

$x_i(0) = x_j(0)$, then it follows from (9.9) that x_i remains always equal to x_j[1]. We assume therefore in the sequel that the agent opinions are initially sorted, and that they remain so: $i \leq j \Rightarrow x_i(t) \leq x_j(t)$.

Second, the opinion of the first agent is always non-decreasing, and the opinion of the last one is always non-increasing. This follows directly from (9.9).

Third, if at some time the distance between two consecutive agent opinions x_i and x_{i+1} is larger than or equal to 1, it remains so forever. In such case, the system can be decomposed into two independent subsystems containing the agents $1, \dots, i$ and $i+1, \dots, n$ respectively.

Finally, the average opinion $\bar{x} = \frac{1}{n}\sum_{i=1}^n x_i$ is preserved as

$$\sum_i \dot{x}_i = \sum_{(i,j):|x_i - x_j|<1} (x_j - x_i) = 0.$$

Moreover, the variance $\sum_{i=1}^n (x_i - \bar{x})^2 = \left(\sum_{i=1}^n x_i^2\right) - n(\bar{x})^2$ of x is non-increasing. We do not prove this last fact here as it is proved later for a more general system and is not used in this section. We can now prove the convergence of the system.

Theorem 10.1. *Let $x_i^* = \lim_{t\to\infty} x_i(t)$. If x evolves according to (9.9), these limits are well defined, and for any i, j we either have $x_i^* = x_j^*$ or $\left|x_i^* - x_j^*\right| \geq 1$.*

Proof. Since the opinions are assumed to be sorted, the opinion x_1 is nondecreasing, and bounded by the initial opinion of x_n. As a result, it converges to a value x_1^*. Let p be the highest index for which x_p converges to x_1^*, with possibly $p = 1$. It follows from the order preservation that the opinion of any agent between 1 and p also converges to x_1^*. We prove that unless $p = n$, there is a time at which $x_{p+1} - x_p \geq 1$. This suffices to prove our result. After this time, the system can indeed be separated into two independent subsystems containing respectively the agents $1, \dots, p$ and $p+1, \dots, n$. In the first one, all opinions converge to x_1^* by definition, and we know that the opinions in the second one remain at a distance at least 1 from all opinions of the first one and in particular from x_1^*. We can then re-apply recursively our argument to this second system and the following ones. Since the number of agents is finite, this establishes the statement of the theorem.

Suppose now to obtain a contradiction that $x_{p+1} - x_p$ always remains smaller than 1. Since x_{p+1} does not converge to x_1^*, it is infinitely often larger than $x_1^* + 2\delta$ for some $\delta > 0$. Since it follows directly from the system definition

[1] For this reasoning to be formally correct, we should prove in addition that the differential equation describing the evolution of one agent opinion admits a unique solution.

(9.9) that all speeds are bounded by $nx_n(0) - nx_1(0)$, this implies that it is infinitely often larger than $x_1^* + \delta$, each time during a time interval larger than or equal to $\Delta t = \delta/(nx_n(0) - nx_1(0))$. Consider now an $\epsilon < \delta$, and a time after which all x_i are distant from x_1^* by less than ϵ for $i = 1, \dots, p$. During any of the intervals in which $x_p \geq x_1^* + \delta$, there would hold

$$\dot{x}_p \geq x_{p+1} - x_p + \sum_{i=1}^{p-1} (x_i - x_p) \geq (\delta - \epsilon) - 2p\epsilon,$$

which is larger than $2\epsilon/\Delta t$ if ϵ is chosen sufficiently small. Then x_p would increase by more than 2ϵ during this interval and becomes larger than $x_1^* + \epsilon$, contradicting our assumption. □

Note that the result above does not hold if the number of agents is infinite. Consider for example that at $t = 0$ and for each $p = 0, 1, 2, \dots,$ there are p agents with opinion $p/2$. Then each agent i constantly increases its opinion x_i with $\dot{x}_i = 1$ and the system does therefore not converge.

We call *clusters* the limiting values to which opinions converge. We also call cluster the set of agents whose opinions converge to a same value. It is stated in Theorem 10.1 that the opinions converge to clusters separated by at least one. However, the typical inter-cluster distances experimentally observed are significantly larger than 1, as shown in Figure 9.4(b). On Figure 10.1 we show the evolution with L of the cluster number and positions, for opinions initially equidistantly distributed on an interval $[-L/2, L/2]$. One can see that the number of clusters is approximately equal to $L/2$ and that they are separated by approximately a distance 2. Such incremental analysis appears in the literature for various similar systems [13, 59, 91, 92]. Ben-Naim et al. also distinguishes different classes of clusters (minor, central and major) [13].

We propose an explanation of this phenomenon based on a stability analysis of the equilibria. The intuition behind our notion of stability can be understood by considering the evolution of the system in Figure 10.2. The system first converges to a “meta-stable” situation where two clusters are separated by a distance slightly larger than one and do therefore not directly interact with each other. Both clusters are however slowly attracted by some isolated agents located between them, these isolated agents remain at the weighted average of the cluster positions as they are attracted by both clusters. Finally, the distance between the clusters becomes smaller than one, so that they attract each other and merge in a single cluster. Note that such phenomenon has already been observed in the literature [91] for the discrete time system. The convergence to the initial potential equilibrium containing two clusters is thus made impossible by the presence of a few agents between the clusters. Moreover, the

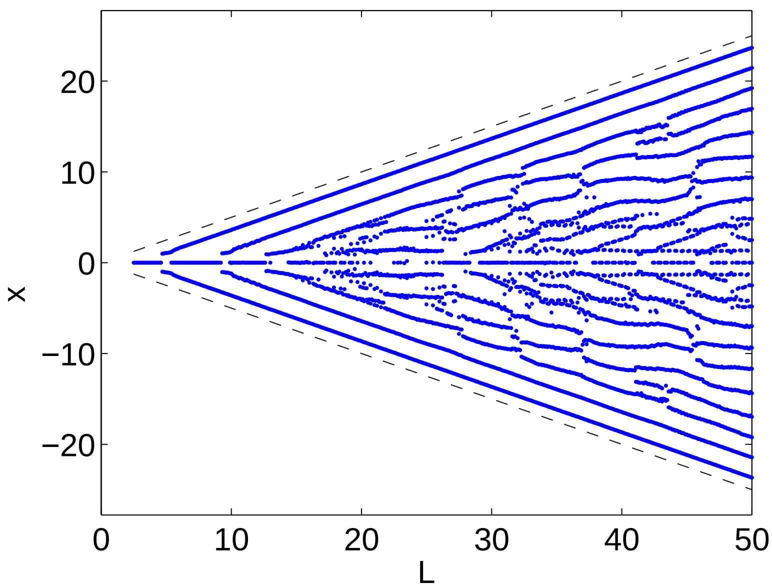

Figure 10.1: Location of the different clusters at equilibrium, as a function of L, for $100L$ agents with opinions initially equidistantly located on $[0, L]$. Clusters are represented in terms of their distance from $L/2$, and the dashed lines represent the endpoints 0 and L of the initial opinion distribution. We do not represent clusters containing less than 10 agents.

number of these isolated agents required to destabilize this equilibrium can be arbitrarily small as compared to the number of agents in the clusters.

The abstract idea of "stability" is that a system should not be too much affected by a small perturbation. There exist many instantiations of this general definition, the suitability of which depend on the systems considered. A usual notion of stability would be to check the robustness of the equilibria with respect to small perturbations δx_i of all the opinions. One can however easily be convinced that every equilibrium would be stable except if two clusters are separated by exactly 1. A notion of stochastic stability also exists, characterizing the robustness to small but repeated perturbations [82,98,108]. This notion is relevant in many situations where the equilibrium is constantly perturbed. In our case, we believe however that a constant perturbation would eventually make all clusters merge, independently of their initial positions, so that every nontrivial set of clusters would be unstable. Observe indeed that if a cluster's position changes due to a perturbation, there is no force driving it back to its initial position. The perturbation would thus cause random displacements of the clusters. As a consequence, there would be with a probability one a time at which two adjacent clusters would be separated by a distance smaller than one, resulting in a rapid merge, independently of their initial positions. By repetitions of this phenomenon, all clusters would eventually merge. Notions of stability of sets such as cyclic stability also exist in game theory [55]. Cyclic stability characterizes the fact that once the system state is in the set, it never leaves it anymore, and can reach any point of the set. This notion is thus related to the notion of invariant set. Nevertheless, it is interesting to observe that our system could be viewed from a game theory point of view, where each agent would greedily minimize its cost function

$$Cost_i(x) = \frac{1}{2}\sum_j \min\left((x_i - x - j)^2, 1\right) \tag{10.1}$$

in continuous time. The set of Nash equilibria of such a game would then correspond to the set of equilibria of the system, that is, the set of collections of clusters separated by at least one. Note that the discrete time system 9.8 can also be viewed as the effect of players greedily minimizing the same cost function 10.1, acting synchronously in discrete time.

The particular notion of stability that we introduce in the sequel is based on the introduction or the displacement at an arbitrary position of an agent of small "weight". It is thus a perturbation of large amplitude, but applied to a small part of the system. If functions are used to represent the systems as in Section 10.2, one could say that the perturbation may have a large ∞-norm, but a small 1-norm, or that it has a large amplitude on a set of small measure. It has the advantage of discriminating some unstable equilibria that are never experimentally observed from the stable equilibria, and it is motivated by the

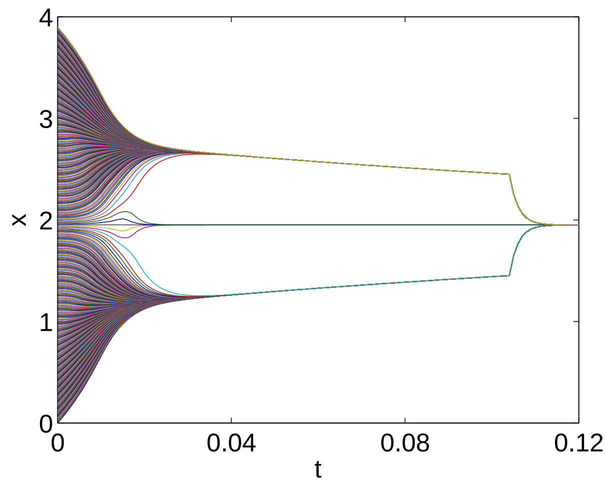

Figure 10.2: Example of temporary "meta-stable" equilibrium. Initially, two clusters are formed and do not interact with each other, but they both interact with a small number of agents lying between them. As a result, the distance separating them eventually becomes smaller than 1. The clusters then attract each other directly and merge into one larger cluster.

often experimentally observed presence of perturbing agents between emerging clusters as in Figure 10.2.

To formalize our notion of *stable equilibrium* we introduce a generalization of the system (9.9) in which each agent i has a weight w_i, and they evolves according to

$$\dot{x}_i(t) = \sum_{j:|x_i(t)-x_j(t)|<1} w_j \left(x_j(t) - x_i(t)\right). \tag{10.2}$$

The convergence result of Theorem 10.1 and the four properties of the system proved above can be generalized to this weighted case. We then call *weight of a cluster* the sum of the weights of all agents in this cluster. If all the agents of a cluster have exactly the same position or the same neighbors, the cluster behaves as one agent with this particular weight. Let $\bar{x}$ be a vector of agent opinions at equilibrium. Suppose that one adds a new agent 0 of weight δ and opinion initially set at $\tilde{x}_0$, and let the system re-converge to a perturbed equilibrium. One then removes the perturbing agent. The opinion vector $\bar{x}'$ so obtained still represents an equilibrium. We denote by $\Delta_{\tilde{x}_0,\delta} = \sum_i w_i \left|\bar{x}_i - x_i'\right|$ the distance between the initial and perturbed equilibria. We say that $\bar{x}$ is *stable* if $\max_{x_0} \Delta_{\tilde{x}_0,\delta}$, the largest distance between initial and perturbed equilibria for a perturbing agent of weight δ can be made arbitrarily small by choosing a sufficiently small δ. An equilibrium is thus unstable if some modification of fixed size can be achieved by adding an agent of arbitrarily small weight. Note that the same notion of stability could also be defined with respect to the displacement of an already present agent of weight δ, possibly obtained by splitting some agent into two parts.

Proposition 10.1. *An equilibrium of the system (10.2) is stable if and only if every two clusters A, B are distant by more than $1 + \frac{\min(W_A,W_B)}{\max(W_A,W_B)}$ or at least 2, where W_A, W_B are the respective weights of the clusters. Equivalently, it is stable if and only if the weighted average of any two clusters is distant by at least one from at least one of the clusters.*

Proof. Suppose that an agent 0 with initial position $\tilde{x}_0$ and weight δ is added to an equilibrium. If its position is the same as one of the clusters, the system remains at the equilibrium. If it is not, it can be connected to at most two clusters as clusters are distant by at least 1 from each other. If it is connected to no cluster, it trivially causes $\Delta_{\tilde{x}_0,\delta} = 0$. If it is connected to one cluster A, the cluster and the agent get attracted one to each other and asymptotically converge to a single position. Since the average position is preserved, this asymptotic position is $\frac{W_A x_A + \delta \tilde{x}_0}{W_A + \delta}$, so that $\Delta_{\tilde{x}_0,\delta} \simeq \frac{\delta}{W_A} |\tilde{x}_0 - x_A| \leq \frac{\delta}{W_A}$, tends to 0 with δ.

We now consider the case where it is connected to two clusters A, B of positions and weights x_A, x_B and W_A, W_B, assuming without loss of generality that $x_B > x_A$ and $W_B \geq W_A$. Let $\bar{x}_{AB}$ be the weighted average of the cluster positions $\frac{W_A x_A + w_B x_B}{W_A + W_B}$. As long as the perturbing agent is connected to both clusters, there holds

$$\dot{x}_0(t) = W_A \left(x_A(t) - x_0(t)\right) + W_b \left(x_B(t) - x_0(t)\right) = \left(W_A + W_B\right) \left(\bar{x}_{AB}(t) - x_0(t)\right),$$

and x_0 is thus attracted by $\bar{x}_{AB}$ the weighted average of x_A and x_B. $x_A(t)$ and $x_B(t)$ are respectively non-decreasing and non-increasing, so that as long as x_0 remains in $(x_B(0) - 1, x_A(0) + 1) \subseteq (x_B(t) - 1, x_A(t) + 1)$, it remains connected to both of them. Suppose first that $x_B(0) - x_A(0) < 1 + \frac{W_A}{W_B}$ (remember that $W_A \leq W_B$), which implies that $\bar{x}_{AB}(0) \in (x_B(0) - 1, x_A(0) + 1)$. $x_{AB}(t)$ differs by at most $O(\delta/(W_A + W_B))$ from the weighted average of x_A, x_B, x_0, which remains constant. If δ is sufficiently small, $\bar{x}_{AB}$ remains thus forever in $(x_B(0) - 1, x_A(0) + 1)$. This implies that x_0 remains also in this interval and remains also connected to both A and B forever. It follows then from Theorem 10.1 and the fact that the average position is preserved implies that x_0, x_A, x_B converge to

$$\frac{w_A x_A(0) + w_B x_B(0) + \delta \tilde{x}_0}{w_A + w_B + \delta} \simeq \frac{W_A x_A + W_B x_B}{W_A + W_B}.$$

The initial equilibrium is thus unstable as for any δ, we have

$$\Delta_{\tilde{x}_0,\delta} = W_A \left|\tilde{x}_0 - x_A\right| + W_B \left|\tilde{x}_0 - x_B\right| = \frac{2W_A W_B}{W_A + W_B} \left|x_B - x_A\right|.$$

To treat the case $x_B - x_A = 1 + \frac{W_A}{W_B}$, it suffices to see that if 0 is initially connected to the two clusters, there holds $x_B - x_A < 1 + \frac{W_A}{W_B}$ after any positive time. Note that the equality is never observed if $W_A = W_B$, as the clusters would then be separated by 2 it would be impossible for 0 to be connected to both of them initially. The condition of the theorem is thus necessary for stability.

Suppose now that $x_B(0) - x_A(0) > 1 + \frac{W_A}{W_B}$, which implies that $\bar{x}_{AB}(0) > x_A(0) + 1$. As long as A and B are disconnected from each other but connected to 0, it follows from (10.2) that

$$\begin{aligned} \dot{x}_0(t) &= (W_A + W_B)(\bar{x}_{AB}(t) - x_0(t)), \\ \dot{x}_A(t) &= \delta(x_0 - x_A) = O(\delta), \\ \dot{x}_B(t) &= \delta(x_0 - x_B) = O(\delta). \end{aligned}$$

0 moves thus toward $\bar{x}_{AB}$ with a speed that can be lower-bounded independently of δ or $\tilde{x}_0$, while x_A, x_B and therefore $\bar{x}_{AB}$ move with a speed proportional to δ. Since $\bar{x}_{AB}(0) > x_A(0) + 1$, this implies that for a sufficiently small δ, there is a time interval that can be upper-bounded independently of δ or $\tilde{x}_0$ after which x_0 becomes larger than $x_A + 1$ and is then not connected to A anymore. In this process, it has influenced x_A and x_B by an amount proportional to δ. As explained above, it then asymptotically merges with the cluster B, modifying x_B by again an amount proportional to δ, so that $\Delta_{\tilde{x}_0,\delta} = O(\delta)$.

□

This result characterizes the set of stable equilibria in terms of lower bounds on the inter-cluster distances. It allows inter-cluster distances smaller than 2 in a stable equilibrium when the clusters have different weight. This is consistent with what can be experimentally observed for some opinion distributions as shown in the example of Figure 10.3. It is however not guaranteed that the system (9.9) always converges to a stable equilibrium. A trivial counterexample is indeed obtained by taking an unstable equilibrium as initial condition. Experimentally tough, we observe that for a given distribution of opinions, convergence almost always occur to a stable equilibrium when the number of agents increases. We therefore make the following conjecture.

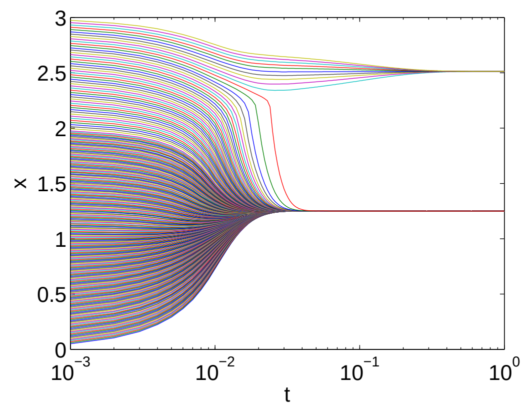

Figure 10.3: Example of convergence to a stable equilibrium where the clusters are separated by less than 2. The equilibrium is obtained by taking initially 201 equidistant opinions on $[0,2]$ and 50 equidistant opinions on $[2,3]$. The clusters contain respectively 241 and 50 agents and $1.2623 > 1.0415 = 1 + \frac{10}{241}$. Note that a logarithmic scale is used due to the presence of different time-scales.

Conjecture 10.1. *If agents evolving according to (9.9) are initially randomly distributed according to a particular continuous p.d.f.*[2]*, the probability that they converge to an equilibrium that is stable tends to 1 when the number of agent tends to infinity.*

In addition to extensive numerical evidences (see e.g. Figure 10.4), this conjecture is supported by the intuitive idea that if the number of agents is sufficiently large, convergence to an unstable equilibrium is made impossible by the presence of at least an agent connected to the two clusters. It is also supported by Proposition 10.2 obtained below for a system allowing a continuum of agents.

10.2 System on a continuum of agents

To further analyze the properties of (9.9) and its behavior when the number of agents increases, we now consider a modified version of the model, which involves a continuum of agents. We use the interval $I = [0,1]$ to index the agents, and suppose that the opinions are non-negative and bounded from above by a certain $L > 0$. We denote by X the set of measurable functions $x : I \to [0,L]$, attributing an opinion $x(\alpha) \in [0,L]$ to every agent in I. As an

[2] probability density function

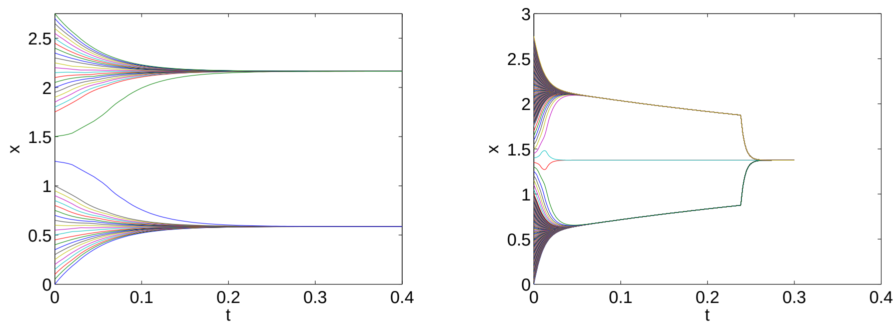

(a) (b)

Figure 10.4: Example of an initial opinion distribution leading to an unstable equilibrium (a): Agents are initially separated by a distance .05 between 0 and 1 and between 1.75 and 2.75, and by a a distance .25 between 1 and 1.75. The clusters at equilibrium have equal weight and are separated by 1.5793. Multiplying the number of agent by 5 leads to a system which converges to a stable equilibrium (b).

example, a uniform distribution of opinions is given by $x(\alpha) = L\alpha$. We use the function $\tilde{x} : I \times \Re^+ \to [0, L] : (\alpha, t) \to \tilde{x}(\alpha, t)$, to represent the evolution of all agent opinions with time. The opinion of an agent α at time t is thus $\tilde{x}(\alpha, t)$. To ease the reading, we define $x_t : I \to [0, L] : \alpha \to x_t(\alpha) = \tilde{x}(\alpha, t)$ to be the restriction of $\tilde{x}$ to any particular time t. Supposing that it is measurable, every x_t belongs thus to X. We also define $\dot{x}_t : I \to \Re : \alpha \to \dot{x}_t(\alpha) = \frac{\partial}{\partial t}\tilde{x}(\alpha, t)$. So, $x_t(\alpha)$ and $\dot{x}_t(\alpha)$ represent respectively the opinion and the opinion changing rate of the agent α at time t. We can now formally define the dynamics of the opinions by

$$\dot{x}_t(\alpha) = \int_{\beta:(\alpha,\beta)\in C_{x_t}} (x_t(\beta) - x_t(\alpha))\, d\beta, \tag{10.3}$$

where $C_x \subseteq I^2$ is defined for any $x \in X$ by

$$C_x := \{(\alpha, \beta) \in I^2 : |x(\alpha) - x(\beta)| < 1\}.$$

In the sequel, we denote by χ_x the indicator functions of C_x. Note that in this section, we do not treat the issues of the existence or uniqueness of the solutions of (10.3), but only characterize the behavior of possible solutions.

It can easily be shown that if x_0 takes a finite number of values, then a solution to (10.3) is provided by the evolution of a discrete system (10.2), where the weight of each agent is the measure of the set on which x takes the corresponding value. Moreover, observe that as in the system for discrete agents, if $x_t(\alpha) > x_t(\beta)$ holds for some t, it holds for all further time. The proof is similar to the one for discrete agents, replacing sums by integrals.

To analyze (10.3), it is convenient to introduce a few concepts. By analogy with interaction graphs in the discrete system, we define for $x \in X$ the *adjacency operator* A_x, which maps the set of bounded measurable functions on I into itself, by

$$A_x y(\alpha) = \int_{\beta \in I} \chi_x(\alpha, \beta) y(\beta) d\beta,$$

and the degree function $d_x(\alpha) : I \to \Re^+$ representing the measure of the set of agents to which it is connected by

$$d_x(\alpha) = \int_{\beta \in I} \chi_x(\alpha, \beta) d\beta = A_x \mathbf{1},$$

where $\mathbf{1} : I \to 1$ represents the constant function taking 1 as value for any $\alpha \in I$. Multiplying a function by a degree function can be viewed as applying the operator (defined on the same set of functions as A_x)

$$D_x y(\alpha) = d_x(\alpha) y(\alpha) = \int_{\beta \in I} a_x(\alpha, \beta) y(\alpha) d\beta.$$

Finally, we define the *Laplacian operator* $L_x = D_x - A_x$. It follows directly from these definitions that $L_x \mathbf{1} = 0$, similarly to what is known for the Laplacian matrix. In the sequel, we use the scalar product $\langle x, y \rangle = \int_{\alpha \in I} x(\alpha) y(\alpha) d\alpha$. We now introduce two lemmas to ease the manipulations of these operators.

Lemma 10.1. *The operators defined above are symmetric with respect to the scalar product: For any $x \in X$, and measurable functions y, z defined on I, there hold $\langle z, A_x y \rangle = \langle A_x z, y \rangle$, $\langle z, D_x y \rangle = \langle D_x z, y \rangle$ and $\langle z, L_x y \rangle = \langle L_x z, y \rangle$.*

Proof. The result is trivial for D_x. For A_x, there holds

$$\begin{aligned} \langle z, A_x y \rangle &= \int_{\alpha \in I} z(\alpha) \left(\int_{\beta \in I} \chi_x(\alpha, \beta) y(\beta) d\beta \right) d\alpha \\ &= \int_{\beta \in I} y(\beta) \left(\int_{\alpha \in I} \chi_x(\alpha, \beta) z(\alpha) d\alpha \right) d\beta \\ &= \langle A_x z, y \rangle \end{aligned}$$

where in the last equality we use the fact that $\chi_x(\alpha, \beta) = \chi_x(\beta, \alpha)$. By linearity the result also holds for L_x and any other linear combination of those operators. □

Lemma 10.2. *For any $x \in X$, and measurable functions y defined on I, there hold*

$$\langle y, (D_x \pm A_x) y \rangle = \frac{1}{2} \int_{(\alpha, \beta) \in I^2} \chi_x(\alpha, \beta) \left(y(\alpha) \pm y(\beta) \right)^2,$$

and as a result $L_x = D_x - A_x$ is positive semi-definite.

Proof. By definition of the operators, we have

$$\langle y, (D_x \pm A_x) y\rangle = \int_{(\alpha,\beta)\in I^2} \chi_x(\alpha,\beta) y(\alpha) \left(y(\alpha) \pm y(\beta)\right).$$

The second member of this equality can be rewritten as

$$\begin{array}{rl} & \frac{1}{2}\left(\int_{(\alpha,\beta)\in I^2} \chi_x(\alpha,\beta) y(\alpha) \left(y(\alpha) \pm y(\beta)\right)\right) \\ + & \frac{1}{2}\left(\int_{(\beta,\alpha)\in I^2} \chi_x(\beta,\alpha) y(\beta) \left(y(\beta) \pm y(\alpha)\right)\right). \end{array}$$

The symmetry of χ_x implies then that

$$\begin{array}{rcl} \langle y, (D_x \pm A_x) y\rangle & = & \frac{1}{2}\int_{(\alpha,\beta)\in I^2} \chi_x(\alpha,\beta) \left(y(\alpha)^2 \pm 2y(\alpha)y(\beta) + y(\beta)^2\right) \\ & = & \frac{1}{2}\int_{(\alpha,\beta)\in I^2} \chi_x(\alpha,\beta) \left(y(\alpha) \pm y(\beta)\right)^2. \end{array}$$

□

The constitutive equation (10.3) can be rewritten, more compactly, in the form

$$\dot{x}_t = -L_{x_t} x_t. \tag{10.4}$$

The set of fixed points of this system is thus characterized by $L_x x = 0$. One can easily see that this set contains the set $F := \{x \in X : x(\alpha) \neq x(\beta) \Rightarrow |x(\alpha) - x(\beta)| \geq 1\}$ of functions taking a discrete number of values, all different by at least one. We prove later that F is exactly the set of solutions to $L_x x = 0$ and thus of fixed points of (10.4), up to a zero-measure correction. By fixed points we mean here those for which $\dot{x}_t = 0$ holds everywhere except maybe on a zero-measure set.

Using the operators that we have defined, it is immediate to see that the system (10.4) preserves $\bar{x}_t = \langle \mathbf{1}, x_t\rangle$, the average value of x_t. The symmetry of L_x and the fact that $L_x \mathbf{1} = 0$ holds for all $x \in X$ imply indeed

$$\dot{\bar{x}}_t = \langle \mathbf{1}, \dot{x}_t\rangle = -\langle \mathbf{1}, L_{x_t} x_t\rangle = -\langle L_{x_t}\mathbf{1}, x_t\rangle = 0.$$

The system also never increases the variance of x, which can be expressed as $Var(x_t) = \langle x_t, x_t\rangle - (\bar{x}_t)^2$. Differentiating this relation indeed yields

$$\dot{Var}(x_t) = 2\langle x_t, \dot{x}_t\rangle - 2\dot{\bar{x}}_t \bar{x}_t = -2\langle x_t, L_{x_t} x_t\rangle,$$

which is non-positive since L_{x_t} is positive semi-definite. This also implies that $\int_{t^*}^{\infty} |\langle x_t, \dot{x}_t\rangle|\, dt < \infty$. We now show that the system variation speed decays to 0. The decay of the variance does indeed not imply that $\dot{x}_t \to 0$ as $\dot{x}_t = -L_{x_t} x_t$ could become arbitrary close to being perpendicular to x_t.

Theorem 10.2. *For any initial condition x_0 of the system (10.3), there holds*

$$\int_0^\infty ||\dot{x}_t)||_2^2 \, dt = \int_0^\infty ||L_{x_t} x_t||_2^2 \, dt < \infty.$$

As a result, the system does not produce cycles other than fixed points, and the variation speed decays to 0.

Proof. We consider the energy function $V : X \to \Re^+$ defined by

$$V(x) = \frac{1}{2} \int_{(\alpha,\beta)\in I^2} \min\left(1, (x(\alpha) - x(\beta))^2\right) d\alpha d\beta \geq 0, \tag{10.5}$$

and show that its derivative is bounded from above by $-2\,||L_x x||_2^2 = -2\,||\dot{x}||_2^2$. Since $\min\left(1, (y(\alpha) - y(\beta))^2\right)$ is not greater than neither 1 or $(y(\alpha) - y(\beta))^2$, there holds for any $y \in X$

$$V(y) \leq \frac{1}{2} \int_{(\alpha,\beta)\in C_x} (y(\alpha) - y(\beta))^2 + \frac{1}{2} \int_{(\alpha,\beta)\in I^2 \setminus C_x} 1 = \langle y, L_x y\rangle + \frac{1}{2} \left|I^2 \setminus C_x\right|, \tag{10.6}$$

where Lemma 10.2 is used to obtain the last equability. For $y = x$, it follows from the definition of C_x that the above inequality is tied. In particular, $V(x_t) = \langle x_t, L_{x_t} x_t\rangle + \frac{1}{2}\left|I^2 \setminus C_{x_t}\right|$ and $V(x_s) \leq \langle x_s, L_{x_t} x_s\rangle + \frac{1}{2}\left|I^2 \setminus C_{x_t}\right|$ for any s and t. Therefore, there holds

$$\begin{array}{rcl} \frac{d}{dt} V(x_t) & \leq & \frac{d}{ds}|_{s=t} \langle x_s, L_{x_t} x_s\rangle + \frac{d}{ds}|_{s=t} \frac{1}{2} \left|I^2 \setminus C_{x_t}\right| \\ & = & \frac{d}{ds}|_{s=t} \langle x_s, L_{x_t} x_s\rangle , \end{array} \tag{10.7}$$

because for any two functions $f, g : \Re \to \Re$, if $f \leq g$ and $f(t) = g(t)$, then $\frac{d}{dt} f(t) \leq \frac{d}{dt} g(t)$. Using $\dot{x}_t = -L_{x_t} x_t$ and the symmetry of L_{x_t}, (10.7) becomes

$$\frac{d}{dt} V(x_t) \leq -2 \langle L_{x_t} x_t, L_{x_t} x_t\rangle = -2\,||L_{x_t} x_t||_2^2 = -2\,||\dot{x}_t||_2^2 .$$

□

We do not use the dependence on x of C_x in the proof of Theorem 10.2, nor when proving that the variance of x is decreasing. These results are therefore valid for any system $\dot{x}_t = -L_t x_t$, with L_t defined with respect to a time evolving symmetric set $C_t \subseteq I^2$. We now show that for the dependence of system (10.3), $L_x x$ is small only if x is close to F, the set of functions taking discrete values separated by at least 1. We also show that F is exactly the set of fixed point of the system. The intuition behind the proof of these results is the following: Consider an agent α with one of the smallest opinions $x(\alpha)$. If the resulting attraction it gets is small, its attraction by those with a larger opinions must be small, as almost no agent has an opinion smaller than it. Therefore, there must be very few agents with an opinion significantly larger than $x(\alpha)$ that interact with α, even if there might be many of them who have an opinion close to α. In other words, possibly many agents have approximately the same opinion close to $x(\alpha)$, and very few agents have an opinion in $[x(\alpha)+\epsilon, x(\alpha)+1]$. So, an agent having an opinion close to $x(\alpha)+1+\epsilon$ interacts with very few agents having an opinion significantly smaller than its own. If its resulting attraction is small, this means thus that its attraction by those agents having a larger opinions is also small, and we can repeat the reasoning.

In order to provide a more formal proof, we need to introduce a measure formalism. For a function $x : I \to [0, L]$ (i.e., a function $x \in X$), and a set $S \subseteq [0, L]$ we let $\mu_x(S)$ be the Lebesgue measure of $\{\alpha : x(\alpha) \in S\}$. By convention, we let $\mu(S) = 0$ if $S \in \Re \setminus [0, L]$. For a given measure μ, we define on $[0, L]$ the function $\hat{L}_\mu$ by

$$\hat{L}_\mu(y) = \int_{z=y-1}^{y+1} (y-z)d\mu.$$

For any $\alpha \in I$, there holds $(L_x x)(\alpha) = \hat{L}_{\mu_x}(x(\alpha))$, so that $-\hat{L}_\mu(y)$ represents the derivative of an agent with opinion y for an opinion function of measure μ. In the sequel, we use $|S|$ to denote the standard Lebesgue measure of a set S, in order to avoid confusion with μ that we just defined.

It is convenient in the sequel to use a topology adapted for measure functions. We say that $x \leq_\mu \epsilon$ if $|\{\alpha : x(\alpha) > \epsilon\}| \leq \epsilon$. Similarly, $x =_\mu 0$ if $|\{\alpha : x(\alpha) \neq 0\}| = 0$, and we call $B_\mu(x, \epsilon)$ the set $\{y : |x-y| <_\mu \epsilon\}$. This allows us to define the corresponding notion of limit. We say that $x_t \to_\mu y$ if for all $\epsilon > 0$, there is a t' after which for all $t > t'$ there holds $x_t \in B_\mu(y, \epsilon)$. We write $x_t \to_\mu S$ for a set S if for any $\epsilon > 0$, there is a t' such that for all $t > t'$, there is a $y \in S$ for which $x_t \in B_\mu(y, \epsilon)$. Before proving our result, we need the following lemma:

Lemma 10.3. *For any real numbers $\epsilon > 0$, $M > 1$ and integer N, there exist $\Delta_1 > 0$ and a sequence $K_1, K_2, \dots K_N$ such that*
a) $K_i > M$ for all i
b) The sequence (Δ_i) defined by $\Delta_{i+1} = 3K_i\Delta_i + \frac{1}{K_i}$ satisfies $\Delta_i K_i < \epsilon$ for all i.

Proof. We prove this lemma by recurrence. It is obviously valid for $N = 1$. We suppose now that it holds for N and prove that it then holds for $N + 1$. Take a $K_{N+1} > M$. Using the recurrence hypothesis, take also a Δ_1 and a sequence $K_1, \dots, K_N$ such that for all $i = 1, \dots N$, there hold $K_i\Delta_i < \frac{\epsilon}{6K_{N+1}}$ and $K_i > \max\left(\frac{2K_{N+1}}{\epsilon}, M\right)$. The conditions on K_i are satisfied for $i = 1, \dots, N + 1$, and so are those on $K_i\Delta_i$ for $i = 1, \dots, N$. The result follows then from

$$K_{N+1}\Delta_{N+1} = K_{N+1}\left(3K_N\Delta_N + \frac{1}{K_N}\right) < K_{N+1}\left(3\frac{\epsilon}{6K_{N+1}} + \frac{\epsilon}{2K_{N+1}}\right) = \epsilon.$$

$\square$

Theorem 10.3. *The distance between $x \in X$ and the set $F \subset X$ of functions taking discrete values separated by at least 1 decays to 0 when $L_x x \to_\mu 0$. In other words, for all $\epsilon > 0$, there is a $\delta > 0$ such that $|L_x x| <_\mu \delta$ implies the existence of a $s \in F$ with $|x - s| <_\mu \epsilon$. As a consequence, $L_x x =_\mu 0$ if and only if $x \in F$.*

Proof. To ease the reading of this proof, we introduce some new notations. For a measure μ, we define on $[0, L]$ the functions

$$\hat{L}^+_\mu(y) = \int_{z\in(y,y+1)} (z-y)d\mu \geq 0, \qquad \text{and} \qquad \hat{L}^-_\mu(y) = \int_{z\in(y-1,y)} (y-z)d\mu \geq 0,$$

so that $\hat{L}_\mu = \hat{L}^+_\mu - \hat{L}^-_\mu$. Observe then that if $|L_x x| <_\mu \delta$, then the set $S \subseteq [0, L]$ on which $\left|\hat{L}^+_\mu(.) - \hat{L}^-_\mu(.)\right| > \delta$ satisfies $\mu_x(S) < \delta$. As a consequence, for any $z \in [0, L]$, there is a y such for which $\hat{L}^+_\mu(y) < \hat{L}^-_\mu(y) + \delta$ and such that $\mu\left((y, z)\right) \leq \delta$, unless of course $\mu(z, L] < \delta$ or $\mu[0, z) < \delta$.

Let us now be given an $\epsilon > 0$. Using Lemma 10.3, we take two sequences $K_1, \dots, K_{\lceil L\rceil}$ and $\Delta_1, \dots, \Delta_{\lceil L\rceil}$ such that for all i, $K_i > (\lceil L\rceil + 1)/\epsilon$ and $\Delta_i < K_i\Delta_i < \epsilon$, and thus $\Delta_i < \epsilon^2/(\lceil L\rceil + 1)$. We also take δ smaller than all $\Delta_i/3$. We prove by induction that if $|L_x x| <_\mu \delta$, then there exist two increasing sequences $x_1 \leq \dots \leq x_N \leq x_{n+1} = L+1$ and $-1 = y_0 \leq y_1 \leq \dots \leq y_N$ with $N \leq \lceil L\rceil$ and $\mu\left((y_n, L]\right) < \epsilon^2/(\lceil L\rceil + 1)$ such that the following conditions hold (for all i for which they make sense)

(a) $\hat{L}^+_\mu(x_i) < \Delta_i$
(b) $x_i \geq y_{i-1} + 1$
(c) $\mu\left([y_{i-1}, x_i)\right) \leq \Delta_i - \delta$
(d) $0 \leq y_i - x_i \leq K_i\Delta_i \leq \epsilon$

This implies that μ is close to a discrete measure taking values separated by at least 1. More precisely, each interval $[x_i, y_i]$ has a length at most ϵ, and the measure of $[0, L]\backslash\bigcup_i[x_i, y_i]$ is by condition (c) at most $\sum_i(\Delta_i - \delta) + \mu\left((y_n, L]\right) \leq \sum_i(\Delta_i) \leq \epsilon^2$. Moreover, Let $s \in F$ be a function which for every α take as value the closest x_i to $x(\alpha)$. Since x takes value differing from all x_i by more than ϵ on a set of measure at most ϵ, there will hold $|x - s| \leq_\mu \epsilon$. Finally, if $L_x x =_\mu 0$, then $|L_x x| <_\mu \delta$ for all positive δ. As a consequence, the distance between x and F is smaller than any positive ϵ and is thus 0. Since F is closed, it follows that $x \in F$.

We begin by initializing our induction. By hypothesis, there exists thus x_1 such that $\mu([0, x_1]) \leq \delta$ and $\hat{L}^+_\mu(x_1) \leq \hat{L}^-_\mu(x_1) + \delta$. Since $y_1 = -1$, this implies that conditions (b) and (c) are satisfied for $i = 1$. Moreover,

$$\hat{L}^-_\mu(x_1) = \int_{z \in (x_1 - 1, x_1)} (x_1 - z) d\mu \leq \int_{z \in (x_1 - 1, x_1)} d\mu \leq \mu((0, x_1)),$$

which we know is no greater than $\delta \leq \Delta_1 - \delta$, so that condition (a) is satisfied as $\hat{L}^+_\mu(x_1) \leq \hat{L}^-_\mu(x_1) + \delta$.

We now assume that there exist $x_1, \ldots, x_{i-1}$ and $y_0, \ldots, y_{i-1}$ satisfying the four conditions (a)-(d) and, x_i satisfying conditions (a)-(c). We then construct a y_i such that condition (d) is satisfied for i and an x_{i+1} satisfying conditions (a)-(c). By induction, this proves the result provided that one stops the construction once a y_i is obtained such that $\mu((y_i, L)) \leq \epsilon^2/(\lceil L \rceil + 1)$, and set then $x_N = x_{i+1} = L + 1$. Our construction guarantees indeed that a $y_i \leq L$ can be found for each $x_i \leq L$, and that a x_{i+1} can be found for each y_i as long as $\mu((y_i, L)) \leq \epsilon^2/(\lceil L \rceil + 1)$. Moreover, condition (b) together with $y_i \geq x_i$ guarantee that the construction stops after at most $\lceil L \rceil$ steps.

Suppose first that $\mu([x_i, x_i + 1)) \leq \delta + \frac{1}{K_i}$, which means that very few agent have opinions between x_i and $x_i + 1$. The following construction is represented in Figure 10.5(a). Let $y_i = x_i$ so that condition (d) is trivially satisfied for i. By hypothesis, we can take a $x_{i+1} \geq y_i + 1$ such that $\hat{L}^+_\mu(x_{i+1}) \leq \hat{L}^-_\mu(x_{i+1}) + \delta$ and $\mu[y_i + 1, x_{i+1}) < \delta$. Condition (b) then trivially holds for $i+1$. By construction, we also have

$$\mu([y_i, x_{i+1})) = \mu([x_i, x_i + 1)) + \mu([y_i + 1, x_{i+1})) \leq \delta + \frac{1}{K_i} + \delta \leq \Delta_{i+1} - \delta.$$

where the last inequality comes form the recurrence $\Delta_{i+1} = 3K_i\Delta_i + \frac{1}{K_i}$ such as defined in Lemma 10.3, and the fact that $2\delta < \Delta_i < K_i\Delta_i$ for all i. As a result, condition (c) holds for $i + 1$. To prove condition (a), observe that

$$\hat{L}^-_\mu(x_{i+1}) = \int_{z \in (x_{i+1} - 1, x_{i+1})} (x_{i+1} - z) d\mu \leq \mu([x_{i+1} - 1, x_{i+1})) \leq \mu([y_i, x_{i+1})).$$

where the last inequality follows from conditions (b) for $i + 1$. By condition (c) for $i + 1$, we have then $\hat{L}^-_\mu(x_{i+1}) \leq \Delta_{i+1} - \delta$, and condition (a) follows then from the fact that x_{i+1} has been chosen under the constraint that $\hat{L}^+_\mu(x_{i+1}) \leq \hat{L}^-_\mu(x_{i+1}) + \delta$.

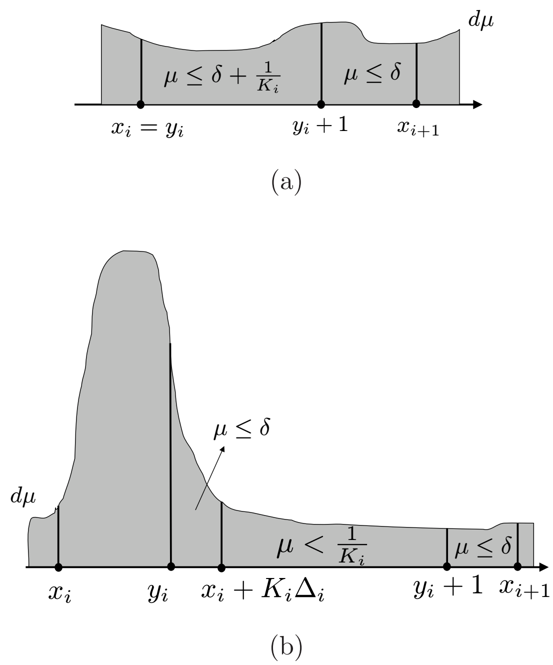

Figure 10.5: Representation of the iterative construction in the proof of Theorem 10.3, for $\mu\left([x_i, x_i + 1)\right) \leq \delta + \frac{1}{K_i}$ (a) and for $\mu\left([x_i, x_i + 1)\right) > \delta + \frac{1}{K_i}$ (b). Note that the density $d\mu$ is represented as continuous here for the sake of clarity but is not assumed to be continuous nor even to exist in the theorem.

We now treat the other case, in which $\mu\left([x_i, x_i+1)\right) > \delta + \frac{1}{K_i}$. Our construction is represented in Figure 10.5(b). It follows from condition (a) that $\mu\left([x_i + K_i\Delta_i, x_i+1)\right) \leq \frac{1}{K_i}$ holds, as otherwise we would have

$$\begin{array}{rcl}
\hat{L}^+_\mu(x_i) & = & \int_{z\in(x_i,x_i+1)}(z-x_i)d\mu \\
& \geq & \int_{z\in(x_i+K_i\Delta_i,x_i+1)}(z-x_i)d\mu \\
& \geq & K_i\Delta_i\mu\left([x_i+K_i\Delta_i, x_i+1)\right) \\
& > & \Delta_i.
\end{array}$$

Let us take a $y_i \leq x_i + K_i\Delta_i$ such that $\mu\left([y_i, x_i+K_i\Delta_i)\right) \leq \delta$ and $\hat{L}^+_\mu(y_i) \leq \hat{L}^-_\mu(y_i) + \delta$. As a result, $y_i \geq x_i$ and condition (d) is satisfied for i. We select then $x_{i+1} \geq y_i + 1$ such that $\hat{L}^+_\mu(x_{i+1}) \leq \hat{L}^-_\mu(x_{i+1}) + \delta$ and $\mu\left([y_i+1, x_{i+1})\right) < \delta$, which implies that condition (b) holds for $i+1$.

To prove the two remaining conditions (c) and (a) for $i+1$, we need an upper bound on $\mu\left([x_i+1, y_i+1)\right)$. Observe first that

$$\begin{array}{rcl}
\hat{L}^+_\mu(y_i) & = & \int_{z\in[y_i,y_i+1)}(z-y_i)d\mu \\
& \geq & \int_{z\in[x_i+1,y_i+1)}(z-y_i)d\mu \\
& \geq & (1+x_i-y_i)\,\mu\left([x_i+1, y_i+1)\right) \\
& \geq & \frac{1}{2}\mu\left([x_i+1, y_i+1)\right),
\end{array} \tag{10.8}$$

where the last inequality comes from condition (d) and the fact that ϵ is assumed to be smaller than $\frac{1}{2}$. We now give an upper bound on $\hat{L}^+_\mu(y_i) \leq \hat{L}^-_\mu(y_i) + \delta$. It follows from conditions (b) and (d) for i that $y_i - 1 \geq x_i - 1 \geq y_{i+1}$. Therefore, there holds

$$\begin{array}{rcl}
\hat{L}^-_\mu(y_i) & \leq & \int_{[y_{i-1},x_i)}(y_i-z)d\mu + \int_{[x_i,y_i)}(y_i-z)d\mu \\
& \leq & \mu\left([y_{i-1}, x_i)\right) + \mu\left([x_i, y_i)\right)(y_i - x_i) \\
& \leq & \Delta_i - \delta + K_i\Delta_i,
\end{array}$$

where the last inequality comes from conditions (c) and (d) for i, and from the fact that $\mu\left([x_i, y_i)\right) \leq \mu\left([0, L]\right) = 1$. Tying this with the lower bound (10.8) leads to

$$\mu\left([x_i+1, y_i+1)\right) \leq 2(K_i+1)\Delta_i. \tag{10.9}$$

We can now use this bound to prove conditions (a) and (c) for $i+1$. Observe that $\mu\left([y_i, x_{i+1})\right)$ can be expressed as

$$\mu\left([y_i, x_i+K_i\Delta_i)\right) + \mu\left([x_i+K_i\Delta_i, x_i+1)\right) + \mu\left([x_i+1, y_i+1)\right) + \mu\left([y_i+1, x_{i+1})\right).$$

y_i has been chosen in such a way that $\mu\left([y_i, x_i+K_i\Delta_i)\right) \leq \delta$, and x_{i+1} such that $\mu\left([y_i+1, x_{i+1})\right) \leq \delta$. Moreover, $\mu\left([x_i+K_i\Delta_i, x_i+1)\right)$ has been proved to be no greater then $\frac{1}{K_i}$. It follows then from (10.9) that

$$\mu\left([y_i, x_{i+1})\right) \leq 2\delta + \frac{1}{K_i} + 2(K_i+1)\Delta_i \leq 3K_i\Delta_i + \frac{1}{K_i} - \delta,$$

where we have used the facts that $3\delta \leq \Delta_i$ and $K_i \geq 3$. Condition (c) for $i+1$ follows then from the definition of the sequence (Δ_i) in Lemma 10.3: $\Delta_{i+1} = 3K_i\Delta_i + \frac{1}{K_i}$. To prove condition (a), there remains to see that

$$\hat{L}_\mu^-(x_{i+1}) = \int_{z\in(x_{i+1}-1,x_{i+1})} (x_{i+1}-z)d\mu \leq \mu\left((x_{i+1}-1, x_{i+1})\right),$$

holds, which by conditions (b) and (c) for $i+1$ implies

$$\hat{L}_\mu^-(x_{i+1}) \leq \mu\left([y_i, x_{i+1})\right) \leq \Delta_i - \delta,$$

completing the proof of Theorem 10.3. □

The following theorem summarizes our results on convergence:

Theorem 10.4. *Let $\tilde{x}$ be a function solution of (10.3), and $F := \{x \in X : x(\alpha) \neq x(\beta) \Rightarrow |x(\alpha) - x(\beta)| \geq 1\}$ be the set of functions of X taking discrete values separated by at least 1. Then $\dot{x}_t =_\mu 0$ if and only if $x_t \in F$. Moreover, $\dot{x}_t \to_\mu 0$ and $x_t \to_\mu F$. As a result all limiting points of μ_{x_t} are discrete measures taking values separated by at least 1.*

Proof. If $x_t \in F$, then it is trivial that $\dot{x}_t =_\mu 0$. It is proved in Theorem 10.3 that this condition is also necessary.
It follows from Theorem 10.2 that $||\dot{x}_t||_2 \to 0$. This implies that $\dot{x}_t \to_\mu 0$, as if $|\dot{x}_t| > \epsilon$ on a set of measure larger than ϵ, then $||\dot{x}_t||_2 \geq \epsilon^{3/2}$. Since $\dot{x}_t = -L_{x_t}x_t$, it follows then from Theorem 10.3 that $x_t \to_\mu F$.
Finally, let M be the set of limiting points of μ_{x_t}, which existence follows from the semi-compactness of the set of measures. Since $x_t \to_\mu F$, and since the measure of any function in F is a discrete one taking value separated by at least one, M contains only such measure. □

Motivated by this theorem, we make the following conjecture:

Conjecture 10.2. *Let $\tilde{x}$ be a function solution of (10.3). Then there is a function $x^* \in F$ such that $x_t \to_\mu x^*$.*

We call clusters the discrete values taken on a positive measure set by a function $x \in F$: c is a cluster of x if $\mu_x c > 0$. We now attempt to characterize the stability of the set F of fixed points, and show that a condition on the inter-cluster distance similar to the one of Proposition 10.1 is necessary for stability. We say that $s \in F$ is *stable* if for any $\epsilon > 0$, there is a $\delta > 0$ such that for any x_0 satisfying $|x_0 - s| \leq_\mu \delta$, there holds $|x_t - s| \leq_\mu \epsilon$ for all time if x evolves according to (10.3).

Proposition 10.2. *Let $s \in F$ be a fixed point of (10.3), and a, b two values taken by s with $\mu_s(a), \mu_s(b) > 0$. If s is stable, then there holds*

$$|b-a| \geq 1 + \frac{\min\left(\mu_s(a), \mu_s(b)\right)}{\max\left(\mu_s(a), \mu_s(b)\right)}. \tag{10.10}$$

Proof. The proof of this Proposition is similar to its discrete counterpart Proposition 10.1. Suppose s does not satisfy this condition for $a < b$. Then since $s \in F$, there holds $\mu_s((a,b)) = 0$. Let $S_a, S_b \subset I$ be sets on which s takes respectively a and b as values, and such that the measure of $S_a \cup S_b$ is δ, and the ratio of their measures is $\frac{\mu_s(a)}{\mu_s(b)}$. We take $x_0 = s$ on $I \setminus (S_a \cup S_b)$ and $x_0 = \frac{\mu_s(a)a + \mu_s(b)b}{\mu_s(a)+\mu_s(b)}$ on $S_a \cup S_b$. Since x_0 takes a finite number of discrete values, a solution to (10.3) $\tilde{x}$ representing its evolution is obtained by considering the corresponding discrete system. Moreover, $\frac{\mu_{x_0}(a)}{\mu_{x_0}(b)} = \frac{\mu_s(a)}{\mu_s(b)}$. Explicit computations similar to those in Proposition 10.1 show then that eventually, x_t tends to $\frac{\mu_{x_0}(a)a+\mu_{x_0}(b)b}{\mu_{x_0}(a)+\mu_{x_0}(b)} = \frac{\mu_s(a)a+\mu_s(b)b}{\mu_s(a)+\mu_s(b)}$ on the set S_{ab} on which x_0 takes a value either a or b. Since this can be done for any $\delta > 0$, s is unstable. $\square$

Unlike in the case of discrete agents, this condition is only necessary.However, a stronger result can be obtained if we assume that $\mu_{x_t}(S)$ remains positive on any subset of $[\inf_{\alpha \in I} x_t(\alpha), \sup_{\alpha\in I} x_t(\alpha)]$ for all t, which is guaranteed if x_t is continuous for all t.

Proposition 10.3. *Let $s \in F$ be a fixed point of (10.3) such that two of its values a and b (with $\mu_s(a), \mu_s(b) > 0$) do not satisfy the condition (10.10) that $|b-a| \geq 1 + \frac{\min(\mu_s(a),\mu_s(b))}{\max(\mu_s(a),\mu_s(b))}$. Then, for every solution $\tilde{x}$ of (10.3) such that $x_t \to_\mu s$, there is a finite t' and an interval $J \subseteq (a,b)$ of positive length such that for all $t \geq t'$, x_t is discontinuous and $\mu_{x_t}(J) = 0$.*

Proof. Consider such a s, and let c be the weighted average $\frac{\mu_s(a)a+\mu_s(b)b}{\mu_s(a)+\mu_s(b)}$, assuming without loss of generality that $a < b$. Take a $\delta > 0$ such that $[c-\delta, c+\delta] \subseteq (b-1, a-1)$. We show if $x \in B_\mu(s, \epsilon)$ for a sufficiently small $\epsilon > 0$, then $-\hat{L}_{\mu_x}(c-\delta) > 0$ and $-\hat{L}_{\mu_x}(c+\delta) < 0$. So any agent having a position $c - \delta$ has a positive derivative, and any agent having a position $c + \delta$ has a negative one. As a consequence, if x_t converges to s, there is a time after which all agents having positions in $J := [c-\delta, c+\delta] \subseteq [a,b]$ remain in this interval forever. It follows then from the convergence to s that $\mu_{x_t}(J) = 0$. This implies that x_t is not continuous on I.
Since $|x - s| \leq_\mu \epsilon$, there hold

$$\begin{aligned}
\mu_x([a-\epsilon, a+\epsilon]) &\leq \mu_s(a) + \epsilon, \\
\mu_x([b-\epsilon, b+\epsilon]) &\geq \mu_s(b) - \epsilon, \\
\mu_x\left((a-1, b+1) \setminus \left([a-\epsilon, a+\epsilon] \cup [b-\epsilon, b+\epsilon]\right)\right) &\leq \epsilon,
\end{aligned}$$

where we have used the fact that all values taken by s are separated by at least 1. If ϵ is sufficiently small so that $a+\epsilon < c-\delta < b-\epsilon$ holds, this implies

$$\begin{array}{lcl} \int_{z\in[a-\epsilon,a+\epsilon]}(z-c+\delta)d\mu \geq (a-\epsilon-c+\delta)(\mu_s(a)+\epsilon) & \leq & 0, \\ \int_{z\in[b-\epsilon,b+\epsilon]}(z-c+\delta)d\mu \geq (b-\epsilon-c+\delta)(\mu_s(b)-\epsilon) & \geq & 0, \\ \int_{z\in(c-\delta-1,c-\delta+1)\setminus([a-\epsilon,a+\epsilon]\cup[b-\epsilon,b+\epsilon])}(z-c+\delta)d\mu & \geq & -\epsilon. \end{array}$$

Therefore, since $[a-\epsilon, b+\epsilon] \subset (c-\delta-1, c-\delta+1) \subset\subset (a-1, b+1)$, there holds

$$\begin{array}{rcl} \hat{L}_{\mu_x}(c-\delta) & \geq & (a-c+\delta-\epsilon)(\mu_s(a)+\epsilon)+(b-c+\delta-\epsilon)(\mu_s(b)-\epsilon)-\epsilon \\ & = & (\mu_s(a)+\mu_s(b))\,\delta - O(\epsilon), \end{array}$$

which is positive if ϵ is sufficiently small. A symmetric argument can be applied to prove that $\hat{L}_{\mu_x}(c+\delta)$ for sufficiently small ϵ. $\square$

Chapter 11

Discrete-Time Opinion Dynamics

11.1 Finite number of discrete agents

We now consider the discrete-time system (9.8) for n discrete agents. We first present some properties of the system, and expose the main differences and similarities with the continuous-time system (9.9). Some of these properties have already been proved in the literature [60, 78, 90].

Unlike in the continuous-time model the strict inequalities between agents' opinions are not preserved. As a simple example, consider two agents with $x_1(0) = -1/4$ and $x_2(0) = 1/4$. We obtain after one iteration $x_1(1) = x_2(1) = 0$.

The order of agents is preserved, but the proof is different from the one for the continuous system: Suppose that $x_i(t) \geq x_j(t)$, and call $N_i(t)$ the set of agents connected to i and not to j, $N_j(t)$ the set of those connected to j and not to i, and $N_{ij}(t)$ the set of those connected to both i and j. We assume here that these sets are non-empty, but our argument can easily be adapted if some of them are empty. For any $k_1 \in N_i(t), k_2 \in N_{ij}(t), k_3 \in N_j(t)$, there holds $x_{k_1}(t) \geq x_{k_2}(t) \geq x_{k_3}(t)$. Therefore, $\bar{x}_{N_i} \geq \bar{x}_{N_{ij}} \geq \bar{x}_{N_j}$, where $\bar{x}_{N_i}, \bar{x}_{N_{ij}}, \bar{x}_{N_j}$ are the average of $x(t)$ on the corresponding set. It follows from (9.8) that

$$x_j(t+1) = \frac{|N_{ij}|\, \bar{x}_{N_{ij}} + |N_j|\, \bar{x}_{N_j}}{|N_{ij}| + |N_j|} \leq \frac{|N_{ij}|\, \bar{x}_{N_{ij}} + |N_i|\, \bar{x}_{N_i}}{|N_{ij}| + |N_i|} = x_i(t+1).$$

We therefore assume in the sequel that the agents are sorted: If $i > j$ then $x_i \geq x_j$.

For exactly the same reasons as in the continuous-time system, the first opinion is nondecreasing and the last one is non-increasing. And, if at some time the distance between two consecutive agent opinions x_i and x_{i+1} is larger than or equal to 1 it remains so for all further time, so that system can then be decomposed into two independent subsystems containing the agents $1,\dots,i$ and $i+1,\dots,n$ respectively.

A difference with the continuous-time system is that the average opinion is not necessarily preserved, and that the variance may increase at some iterations. Consider for example agents with initial opinions $(0, 0.8, 1.2)$. After one iteration, the opinions are $(0.4, \frac{2}{3}, 1)$ so that the average opinion moves from $\frac{2}{3}$ to $\frac{31}{45}$. An example where the variance increases is obtained by taking initial opinions $(-11, -11, -10.5, -10, 10, 10.5, 11, 11)$. The variance then increases from 113.0625 to 113.1688. It has to be noted that examples of increasing variances are hard to find, and this increase is here a small one. One possible way to bound the increase is to use the update matrix norm. To avoid confusion with the adjacency operator, let us temporarily call M the $n \times n$ stochastic matrix representing the iteration (9.8) at some time t, with $x(t+1) = Mx(t)$. Since the variance of $x(t)$ is unaffected by the addition of a constant value to all $x_i(t)$, let us assume that $\sum x_i(t) = 0$. We have then $Var\,(x(t)) = ||x(t)||_2^2$. The variance of $x(t+1)$ on the other hand is

$$\left|\left|x(t+1) - \frac{1}{n}\mathbf{1}\mathbf{1}^T x(t+1)\right|\right|_2^2 = \left|\left|(M - \frac{1}{n}\mathbf{1}\mathbf{1}^T)x(t)\right|\right|_2^2,$$

where we have use the fact that $\mathbf{1}^T M = \mathbf{1}^T$ for any stochastic matrix M. The proportional increase is thus bounded by

$$\frac{Var\,(x(t+1))}{Var\,(x(t))} = \frac{\left|\left|(M - \frac{1}{n}\mathbf{1}\mathbf{1}^T)x(t)\right|\right|_2^2}{||x(t)||_2^2} \leq \left|\left|M - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right|\right|_2^2.$$

Note that since $\mathbf{1}^T M = \mathbf{1}^T$ and $\left|\left|I - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right|\right|_2 = 1$, there holds $\left|\left|M - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right|\right|_2 \leq ||M||_2$. Bounding the increase of the variance could thus be achieved by studying the 2-norm of the matrix M, which is related to the adjacency matrix of an interval graph. Such bound would however not take the nonlinearity of the system into account, i.e. the fact that the update matrix depends on $x(t)$.

We have seen in Section 9.3 that the convergence of (9.8) to a set of opinion clusters separated by at least 1 is a consequence of Theorem 9.4. We now provide a simpler convergence proof similar to the one of Theorem 10.1, and taking advantage of the particular system dynamics, that is, of the way the neighborhood topology depends on x.

Theorem 11.1. *Let $x_i^* = \lim_{t\to\infty} x_i(t)$. If x evolves according to (9.8), these limits are well defined, and there is a t' such that $x_i(t) = x_i^*$ holds for any i and $t' \geq t$. For any i, j we either have $x_i^* = x_j^*$ or $\left|x_i^* - x_j^*\right| \geq 1$. As a consequence, there is a t'' such that for any $t \geq t$" and any i, j, there holds $|x_i(t) - x_j(t)| \geq 1$ if $x_i^* \neq x_j*$ and $|x_i(t) - x_j(t)| < 1$ otherwise.*

Proof. Since the opinions are assumed to be sorted, the opinion x_1 is nondecreasing, and bounded by the initial opinion of x_n. As a result, it converges to a value x_1^*. Let p be the highest index for which x_p converges to x_1^*. For the same reasons as in the proof of Theorem 10.1, convergence is obtained by proving that if $p \neq n$, there is a time at which $x_{p+1} - x_p \geq 1$. Suppose to obtain a contradiction that this is not the case, and take an ϵ and a time after which all x_i are distant from x_1^* by less than ϵ for $i = 1, \ldots, p$. Since x_{p+1} does not converge to x_1^*, there is always a further time at which it is larger than $x_1^* + \delta$ for some $\delta > 0$. For such time t^*, there holds

$$x_{p+1}(t^* + 1) \geq \frac{1}{p+1}\left(\sum_{i=1}^{p+1} x_i(t^*)\right) \geq \frac{1}{p+1}\left((p+1)x_1^*(t) + \delta - p\epsilon\right),$$

which is larger than $x_1^* + \epsilon$ is ϵ is chosen sufficiently small. This however contradicts the fact that x_p remains distant from x_1^* by less than ϵ.

There is thus a time after which $x_{p+1} \geq x_p + 1$, which implies that the agents $p+1, \ldots, n$ do not influence the agents $1, \ldots, p$. Since those have opinion converging to x_1^*, they eventually become sufficiently close to each other to be all connected to each other. When this happens, they all compute the same average and reach the same opinion at the next time-step, and keep then this opinion for all further time. They converge thus in finite time. □

Although Theorem 11.1 only states that inter-cluster distances are no smaller than 1, it has been observed in the literature that these distance are usually significantly larger than 1 [78,91], similarly to what is observed for the continuous-time system. This can for example be seen in Figure 9.4(b). This phenomenon is further represented in Figure 11.1. As in the continuous-time system, this can be attributed to the fact that clusters too close to each other are forced to merge by the presence of isolated agents between them. We now show stability results similar to Proposition 10.1. For this purpose, we again consider a weighted variant of the system, in which each agent has a weight w_i and evolves according to

$$x_i(t+1) = \frac{\sum_{j:|x_i(t)-x_j(t)|<1} w_j x_j(t)}{\sum_{j:|x_i(t)-x_j(t)|<1} w_j}. \tag{11.1}$$

The convergence result of Theorem 11.1 and the other properties of the system detailed at the beginning of this section are also valid for this weighted system.

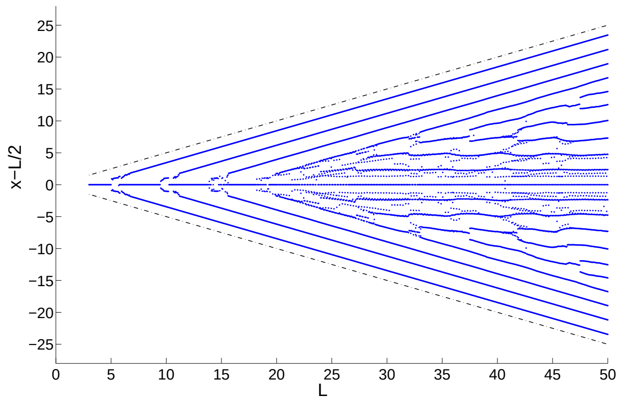

Figure 11.1: Location of the different clusters at equilibrium, as a function of L, for $5000L$ agents with opinions initially equidistantly located on $[0, L]$. Clusters are represented in terms of their distance from $L/2$, and the dashed lines represent the endpoints 0 and L of the initial opinion distribution.

We call again *weight of a cluster* the total weight of all the agents converging to this cluster. We use the same notion of stability with respect to the addition of an agent as in Section 10.1.

Proposition 11.1. *An equilibrium is stable if and only if the distance between any two clusters A and B is larger than $1 + \frac{\min(W_A,W_B)}{\max(W_A,W_B)}$. In this expression, W_A and W_B are the weights of the clusters.*

Proof. Consider an equilibrium $\bar{x}$ and an additional agent 0 of initial opinion $\tilde{x}_0$ and weight δ. If this agent is disconnected from all clusters, it has no influence and $\Delta_{\tilde{x}_0,\delta} = 0$. If it is connected to one cluster A of position x_A and weight W_A, the system reaches a new equilibrium after one time step, where both the additional agent and the cluster have an opinion $(\tilde{x}_0\delta + x_A W_A)/(\delta + W_A)$. So $\Delta_{\tilde{x}_0,\delta} \leq \delta\,|\tilde{x}_0 - x_A|$. Suppose now that the perturbing agent 0 is connected to two clusters A, B (it is never connected to more than two clusters). For a sufficiently small δ, its position after one time step is approximately

$$x_0' = x_A + \frac{x_B - x_A}{1 + \frac{W_A}{W_B}} = \frac{x_A - x_B}{1 + \frac{W_B}{W_A}} + x_B, \tag{11.2}$$

while the new positions of the clusters are $(\tilde{x}_0\delta + x_A W_A)/(\delta + W_A)$ and $(\tilde{\delta} + x_B W_B)/(\delta + W_B)$. If $|x_A - x_B| > 1 + \frac{\min(W_A,W_B)}{\max(W_A,W_B)}$, it follows from (11.2) that for small δ the agent is then connected to only one cluster and that equilibrium is thus reached at the next time step, with a $\Delta_{\tilde{x}_0,\delta}$ proportional to δ. The condition of this theorem is thus sufficient for stability of the equilibrium as $\Delta_{\tilde{x}_0,\delta}$ is proportional to δ when it is satisfied.

If the condition is not satisfied, the agent is still connected to both clusters. An explicit recursive computation shows that in the sequel its opinion remains approximately at the weighted average of the two clusters (11.2), while these get steadily closer one to each other. Note that their weighted average moves at each iteration in the direction of the largest cluster by a distance bounded by $\delta/(W_A + W_B)$. Once the distance separating the clusters becomes smaller than or equal to 1, they merge in one central cluster of opinion x_0'. Thus, in this case, the addition of a perturbing agent of arbitrary small weight δ connected to both A and B results in the merging of the clusters. Let x_{AB} be the position of the new cluster, there holds

$$\Delta_{\tilde{x}_0,\delta} = W_A \left|x_{AB} - x_A\right| + W_B \left|x_{AB} - X_B\right| \geq \min(W_A, W_B) \left|x_B - x_A\right|,$$

independently of the weight δ of the agent. □

A system does not necessarily always converge to a stable equilibrium, but we make the following conjecture:

Conjecture 11.1. *If agents evolving according to (9.8) are initially randomly distributed according to a continuous p.d.f., the probability that they converge to a stable equilibrium tends to 1 when the number of agent tends to infinity.*

This conjecture is supported by the same intuitive reasons as Conjecture 10.1, and by extensive numerical experiments such as shown in Figure 11.2. Moreover, it is consistent with similar results obtained in Section 11.2 on system defined on an agent continuum, and results obtained in Section 11.3 linking the systems defined on continuum with those on discrete agents.

To close this section, we show that the relative continuity in the cluster positions observed in Figure 11.1 and the fact that the external clusters positions become constant with respect to the initial distribution edge when L is sufficiently large can both be understood by analyzing the "information" propagation.

During an iteration, an agent is only influenced by those opinions within distance 1 of its own, and its opinion is modified by less than 1. So information is propagated by at most a distance 2 at every iteration. In the case of an initial uniform distribution on $[0, L]$ for a large L, during the first iterations the agents with initial opinions close to 0 behave as if opinions were initially

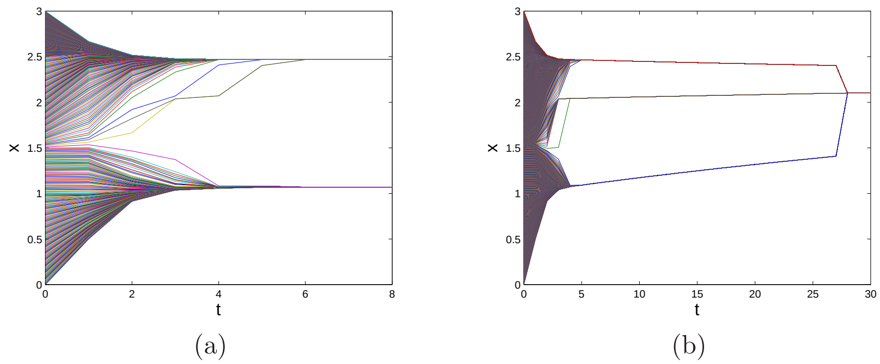

Figure 11.2: Evolution with time of agent opinions initially approximating the same opinion function: the density of opinions between on (2.5,3) is five time as high as the density on (0,2.5). In (a), the 501 agents converge to an unstable equilibrium: the clusters have respective weights 152 and 349, and are separated by a distance $1.399 < 1 + \frac{152}{349} \simeq 1.436$. In (b), the 5001 agents converge to a stable equilibrium.

distributed uniformly on $[0, +\infty)$. Moreover, once a group of opinions is separated from other opinions by more than 1, they are not influenced by them at any subsequent iteration. Therefore, agents with initial opinions close to 0 and getting separated from the other opinions after some finite time follow exactly the same trajectories when the initial uniform distribution is on $[0, +\infty)$ or on $[0, L]$ for a sufficiently large L, and its final position would thus be constant with respect to the initial distribution edge.

We performed simulations with an initial semi-infinite interval, i.e. opinions equidistantly distributed between 0 and $+\infty$. It appears that every agent eventually gets disconnected from the semi-infinite set but remains connected with some other agents. Each group behaves then independently of the rest of the system and converges to a single cluster. As shown in Figure 11.3, the distance between two consecutive clusters converges to approximately 2.2. The same observation was made by Lorenz for his "interactive Markov chain model" which approximates this system [91]. If it were proved, this asymptotic inter-cluster distance would partially explain the precise evolution of the number of clusters (as a function of L) shown in Figure 11.1.

Such semi-infinite simulation could not be performed in continuous-time. In discrete time, one can verify that if initial opinions are equidistantly distributed, an agent opinion x_i does not vary during the $\lfloor x_i(0) \rfloor$ first time-steps. This property allows us to simulate a semi-infinite distribution of opinions while only performing a finite number of operations. Unfortunately, it cannot be applied in continuous time.

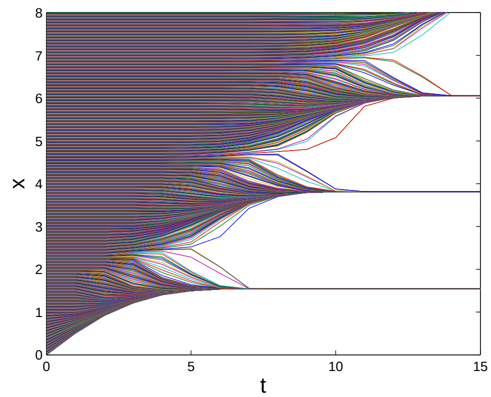

Figure 11.3: Evolution with time of the opinions for an initial semi-infinite equidistant distribution of opinions (initially, there are 100 agents within each unit of the x axis).

11.2 System on a continuum of agents

In order to analyze the behavior of (9.8) for large numbers of agents, we now adapt it to allow a continuum of agents. We first adapt the results of Section 10.2 to the discrete-time case. We show then that in the discrete case it is possible to analyze more formally the notion that the system on a continuum is "the limit" of (9.8) for large number of agents, as we show in Section 11.3. We use again the interval $I = [0, 1]$ to index the agents and denote by $x_t(\alpha)$ the opinion at time t of the agent $\alpha \in I$. The evolution of the opinions is then described by

$$x_{t+1}(\alpha) = \frac{\int_{\beta:(\alpha,\beta)\in C_{x_t}} x_t(\beta)d\beta}{\int_{\beta:(\alpha,\beta)\in C_{x_t}} d\beta} \tag{11.3}$$

where $C_x \subseteq I^2$ is defined for any $x \in X$ by

$$C_x = \{(\alpha, \beta) \in I^2 : |x(\alpha) - x(\beta)| < 1\},$$

as in Section 10.2. To fully define our iteration, we let $x_{t+1}(\alpha) = x_t(\alpha)$ if $\int_{\beta:(\alpha,\beta)\in C_{x_t}} d\beta = 0$. However, since the set of α for which this convention needs to be used has zero measure, we do not consider them in the sequel. Note that since this system is defined in discrete-time, there is no issue of existence nor of uniqueness of its solutions. Besides, for the same reasons as for the system (9.8), if $x_t(\alpha) \geq x_t(\beta)$ holds for some t, it holds for any further time. And, if $x_t(\alpha) = x_t(\beta)$ holds for some time, it also holds for any further time. As a consequence, one can see that if x_0 only takes a finite number of values, the system simulates a discrete system where agents initially have these values, and where each agent weight is the measure of the set on which x_0 takes the corresponding values.

Using the operators formalism, (11.3) can be rewritten as

$$\Delta x_t := x_{t+1} - x_t = -d_{x_t}^{-1} L_{x_t} x_t \tag{11.4}$$

Since d_x is bounded by $|I| = 1$, this implies that $\Delta x = 0$ if and only if $L_x x =_\mu 0$. It follows then from Theorem 10.3 that the set of fixed points of this system is the set F of functions taking discrete values separated by at least one, as for the system (10.3). We now prove that Δx_t decays to 0.

Theorem 11.2. *For any initial condition of the system (11.4) there holds*

$$\sum_{t=0}^{\infty} \int_{(\alpha,\beta)\in C_{x_t}} \left(\Delta x_t(\alpha) + \Delta x_t(\beta)\right)^2 < \infty.$$

As a result, the system does not produce cycles other than fixed points.

Proof. We consider the non-negative energy function defined in (10.5) as in the proof of Theorem 10.5, and show that

$$V(x_{t+1}) - V(x_t) \leq - \left\langle \Delta x_t, (A_{x_t} + D_{x_t})\Delta x_t \right\rangle,$$

which by Lemma 10.2 implies the desired result. For reasons explained in the proof of Theorem 10.2, there holds $V(x_t) = \langle x_t, L_{x_t} x_t \rangle + \frac{1}{2}\left|I^2 \setminus C_{x_t}\right|$, and for all time s, $V(x_s) \leq \langle x_s, L_{x_t} x_s \rangle + \frac{1}{2}\left|I^2 \setminus C_{x_t}\right|$. Taking $s = t+1$, we obtain from these two relations

$$V(x_{t+1}) - V(x_t) \leq \langle x_{t+1}, L_{x_t} x_{t+1} \rangle - \langle x_t, L_{x_t} x_t \rangle = 2\left\langle \Delta x_t, L x_t \right\rangle + \left\langle \Delta x_t, L_{x_t} \Delta x_t \right\rangle,$$

where we have used the symmetry of L_{x_t}. It follows from (11.4) that $L_{x_t} x_t = -D_x \Delta x$, so that we have

$$V(x_{t+1}) - V(x_t) \leq -2\left\langle \Delta x_t, D_{x_t} x_t \right\rangle + \left\langle \Delta x_t, L_{x_t} x_t \right\rangle = \left\langle \Delta x_t, (A_{x_t} + D_{x_t})\Delta x_t \right\rangle,$$

since $L_x = D_x - A_x$. □

Note again that the above proof does not use the dependence of the topology C_x on x, and is therefore valid for any dependence. This is not the case though for the next results.

Theorem 11.3. *Let (x_t) be a sequence of functions of X evolving according to the model (10.3), and F be the set of functions taking discrete values separated by at least 1. Then $(x_{t+1} - x_t) \to_\mu 0$ and $x_t \to_\mu F$. As a result all limiting points of μ_{x_t} are discrete measures taking values separated by at least 1. Moreover, x is a fixed point of (10.3) if and only if $x \in F$.*

Proof. We begin by proving the convergence of Δx_t. Suppose that $\Delta x_t = (x_{t+1} - x_t) \rightarrow_\mu 0$ does not hold. Then there is an $\epsilon > 0$ such that for arbitrarily large t, there is a set of measure at least ϵ on which $|\Delta x_t| > \epsilon$. Consider such a time t Let us suppose that there is then a set $S \subseteq I$ of measure at least $\epsilon/2$ on which $\Delta x_t > \epsilon$. If this does not hold, a similar reasoning can be done on a set on which $\Delta x_t < -\epsilon$. For each $i \in \{1, \ldots, 2\lceil L \rceil\}$, let $A_i \subset I$ be the set on which $x_t \in [(i-1)/2, i/2]$. For any i and for any $\alpha, \beta \in A_i$, there holds $|x_t(\alpha) - x_t(\beta)| < 1$ and thus $(\alpha, \beta) \in C_{x_t}$. Therefore, $A_i^2 \subseteq C_{x_t}$ for all i. Moreover, the sets A_i cover $[0,1]$, so that $\sum_{i=1}^{2\lceil L \rceil} |A_i \cap S| \geq |S| \geq \epsilon/2$. There is thus at least a i^* such that $|A_i \cap S| \geq \epsilon/(4\lceil L \rceil)$. We have then

$$\begin{array}{rcl} \int_{(\alpha,\beta)\in C_{x_t}} \left(\Delta x_t(\alpha) + \Delta x_t(\beta)\right)^2 & \geq & \int_{(\alpha,\beta)\in (A_{i^*}\cap S)^2} \left(\Delta x_t(\alpha) + \Delta x_t(\beta)\right)^2 \\ & \geq & 4\epsilon^2 |A_{i^*} \cap S|^2 \\ & \geq & \frac{\epsilon^4}{4\lceil L \rceil^2}. \end{array}$$

So if $\Delta x_t \rightarrow_\mu 0$ does not hold, $\int_{(\alpha,\beta)\in C_{x_t}} \left(\Delta x_t(\alpha) + \Delta x_t(\beta)\right)^2$ does not decay to 0, which contradicts Theorem 11.2.

For the same reasons as in Theorem 10.4, all limiting points of μ_{x_t} are discrete measures taking values separated by at least 1. Finally, we know that x is a fixed point of (11.3) if and only if $L_x x =_\mu 0$, which is equivalent to $x \in F$ as proved in Theorem 10.4. □

Motivated by this theorem we now make the following conjecture:

Conjecture 11.2. *Let (x_t) be a sequence of functions of X evolving according to the model (10.3). Then there is a function $x^* \in F$ such that $x_t \rightarrow_\mu x^*$.*

As in the continuous-time case, F is the set of fixed points of the systems, but some of these fixed points are unstable. We say that $s \in F$ is stable if for any $\epsilon > 0$, there is a $\delta > 0$ such that for any $x_0 \in B_\mu(s, \delta)$, there holds $x_t \in B_\mu(x, \epsilon)$ for all time t if the sequence.

Proposition 11.2. *Let $s \in F$ be a fixed point of (11.3), and a, b two values taken by s. If s is stable, then there holds $|b - a| \geq 1 + \frac{\min(\mu_s(a), \mu_s(b))}{\max(\mu_s(a), \mu_s(b))}$.*

Proof. The proof can be adapted from the proof of Proposition 11.1 exactly as the proof of Proposition 10.2 is adapted from the proof of Proposition 10.1. □

To further analyze the properties of the opinion function sequences, we introduce a notion of framed opinion function. We say that a function $x \in X$ is *framed* if there exist $M \geq m > 0$ such that for any interval $J \subseteq [\inf_\alpha x, \sup_\alpha x]$, there holds $m\,|J| \leq \mu_x(J) \leq M\,|J|$. Intuitively, a function is framed if the set of opinions taken is connected, and if the density of agents on any interval of opinion is positively bounded from above and from below. As a consequence, no value is taken by a set of positive measure. For example, any piecewise differentiable $x \in X$ with positive lower and upper bound on its derivative is framed. We show in the rest of this section that if x_0 is framed and if (x_t) converges, then it asymptotically converges to an equilibrium satisfying the condition (10.10) on the minimal distance between opinions at the equilibrium, provided that $\sup_\alpha x_t - \inf_\alpha x_t$ remains always larger than 2. To conveniently express this, we define on X the nonlinear update operator U which takes its values in X by $U(x) = x - d_x^{-1} L_x x = d^{-1} A_x x$, so that the recurrence (11.3) can be rewritten $x_{t+1} = U(x_t)$. For a function $x \in X$, let now $u_x : [0, L] \to [0, L]$ be the function sending an opinion on its update opinion defined by

$$u_x(a) = \frac{\int_{z\in(a-1,a+1)} z d\mu}{\mu_x\left((a-1, a+1)\right)}.$$

As a consequence $(U(x))(\alpha) = u_x\left(x(\alpha)\right)$ and $x_{t+1}(\alpha) = u_{x_t}\left(x_t(\alpha)\right)$.

Proposition 11.3. *Let $x \in X$ be a framed function such that $\sup_\alpha x - \inf_\alpha x > 2$. Then $U(x)$ is framed.*

Proof. By hypothesis, there exists $0 < m < M$ such that for any $[a, b] \subseteq [\inf_\alpha x, \sup_\alpha x]$ there holds $m(b-a) \leq \mu_x([a, b]) \leq M(b-a)$. Let $\delta = \min\{\frac{1}{2}, \sup_\alpha x - \inf_\alpha x - 2, \frac{M}{m}\}$. We first prove the existence of $M', m' > 0$ such that if $[a, b] \subseteq [\inf_\alpha x, \sup_\alpha x]$ and $b-a < \delta$, then $m'(b-a) \leq u_x(b) - u_x(a) \leq M'(b-a)$.

Due to the value of δ, either $a \geq \inf_\alpha x + 1$ or $b \leq \sup_\alpha x - 1$. We consider here the second case, but the first one can be treated exactly in the same way. There holds therefore $(a, b+1) \subseteq [\inf_\alpha x, \sup_\alpha x]$. Let $\bar{\mu}_{ab} = \mu_x\left((b-1, a+1)\right)$, $\bar{\mu}_{a\backslash b} = \mu_x\left((a-1, b-1]\right)$ and $\bar{\mu}_{b\backslash a} = \mu_x\left([a+1, b+1)\right)$. It follows then from the framed character of x and the value of δ that

$$\begin{array}{l} \bar{\mu}_{ab} = \mu_x\left((b-1, a+1)\right) \geq \mu_x\left((a, a+1)\right) \geq m, \\ 0 \leq \bar{\mu}_{a\backslash b} = \mu_x\left((a-1, b-1]\right) \leq M(b-a) \leq M\delta \leq m \\ m(b-a) \leq \bar{\mu}_{b\backslash a} = \mu_x\left([a+1, b+1)\right) \leq M(b-a) \leq M\delta \leq m. \end{array} \tag{11.5}$$

Let now $\bar{x}_{ab}, \bar{x}_{a\backslash b}, \bar{x}_{b\backslash a}$ be the average value of x (weighted by μ) on respectively $(b-1, a+1)$, $(a-1, b-1]$ and $[a+1, b+1)$. In case $\bar{\mu}_{a\backslash b} = 0$, we let $x_{a\backslash b} = b-1$. There holds $\bar{x}_{a\backslash b} \leq \bar{x}_{ab} \leq \bar{x}_{b\backslash a}$, and $1 \leq 2-(b-a) \leq \bar{x}_{b\backslash a} - \bar{x}_{a\backslash b} \leq 2+(b-a) \leq 3$. Moreover, by definition of u_x, there holds

$$u_x(a) = \frac{\bar{\mu}_{ab}\bar{x}_{ab} + \bar{\mu}_{a\backslash b}\bar{x}_{a\backslash b}}{\bar{\mu}_{ab} + \bar{\mu}_{a\backslash b}} = \bar{x}_{ab} - \frac{\bar{\mu}_{a\backslash b}(\bar{x}_{ab} - \bar{x}_{a\backslash b})}{\bar{\mu}_{ab} + \bar{\mu}_{a\backslash b}},$$

and

$$u_x(b) = \frac{\bar{\mu}_{ab}\bar{x}_{ab} + \bar{\mu}_{b\backslash a}\bar{x}_{b\backslash a}}{\bar{\mu}_{ab} + \bar{\mu}_{b\backslash a}} = \bar{x}_{ab} + \frac{\bar{\mu}_{b\backslash a}(\bar{x}_{b\backslash a} - \bar{x}_{ab})}{\bar{\mu}_{ab} + \bar{\mu}_{b\backslash a}}.$$

Using the bounds (11.5), we have then

$$u_x(b) - u_x(a) \leq \frac{\bar{\mu}_{b\backslash a}}{\bar{\mu}_{ab}}(\bar{x}_{b\backslash a} - \bar{x}_{ab}) + \frac{\bar{\mu}_{a\backslash b}}{\bar{\mu}_{ab}}(\bar{x}_{ab} - \bar{x}_{a\backslash b}) \leq \frac{M(b-a)}{m}(\bar{x}_{b\backslash a} - \bar{x}_{a\backslash b}),$$

and thus $u_x(b) - u_x(a) \leq 3\frac{M}{m}(b-a)$, which proves our upper bound for $M' := 3\frac{M}{m}$. For the lower bound, observe first that $\bar{x}_{b\backslash a} - \bar{x}_{ab} \geq a + 1 - \bar{x}_{ab} \geq \frac{1}{2}\sqrt{\frac{m}{M}}$, where the last inequality is obtained by considering the worst case in which $\mu_x([\bar{x}_{ab}, a+1)) = M(a+1-\bar{x}_{ab})$ and $\mu_x((b-1, \bar{x}_{ab})) = m\mu_x((b-1, \bar{x}_{ab}))$. Besides, since $\bar{x}_{a\backslash b} \leq \bar{x}_{ab}$, it follows from the bounds (11.5) that

$$u_x(b) - u_x(a) \geq \frac{\bar{\mu}_{b\backslash a}}{2\bar{\mu}_{ab}}(\bar{x}_{b\backslash a} - \bar{x}_{ab}) \geq \frac{m}{6M}\frac{1}{2}\sqrt{\frac{m}{M}}(b-a),$$

where we have used the fact that $\bar{\mu}_{ab} \leq M(a+1-b-1) \leq 3M$. The lower bound is thus proved for with $m' = \frac{m^{3/2}}{12M^{3/2}}$

By finite unions and intersections of intervals, this result holds for all intervals $[a,b] \subseteq [\inf_\alpha x, \sup_\alpha x]$. Consider now an interval $[a', b'] \in [\inf_\alpha U(x), \sup_\alpha U(x)]$, and let $a = \inf\{z \in [0,L] : u_x(z) \in [a', b']\}$ and $b = \sup\{z \in [0,L] : u_x(z) \in [a', b']\}$. As a consequence of the order preservation property, $u_x((a,b)) \subseteq [a', b']$, and $[a', b'] \subseteq [u_x(a), u_x(b)]$. Since by hypothesis $\mu_x(a) = \mu_x(b) = 0$, this implies that $\mu_{U(x)}([a', b']) = \mu_x([a,b]) \in [m(b-a), M(b-a)]$. Using the bounds on $\frac{u_x(b)-u_x(a)}{b-a}$, we finally obtain

$$mm'(b'-a') \leq \mu_{U(x)}([a', b']) \leq MM'(b'-a')$$

$\square$

A consequence of this result and of Theorem 11.3 is that if x_0 is framed, then (x_t) does not converge in finite time unless $\sup_\alpha x_t - \inf_\alpha x_t$ becomes smaller than or equal to 2. We can now show that a sequence of framed functions never converges to an equilibrium that does not satisfy the condition (10.10).

Proposition 11.4. *Let (x_t) be a sequence of functions of X evolving according to (11.3) such that x_0 is framed and $\sup_\alpha x_t - \inf_\alpha x_t > 2$ for all t. If (x_t) converges, then it converge to a function $s \in F$ such that*

$$|b-a| \geq 1 + \frac{\min(\mu_s(a), \mu_s(b))}{\max(\mu_s(a), \mu_s(b))},$$

hold for any values a, b taken by s with $\mu_s(a), \mu_s(b) > 0$. In particular, if $\mu_s(a) = \mu_s(b)$, then $|b-a| \geq 2$.

Proof. It follows from Theorem 11.3 that $s \in F$, and from Proposition 11.3 that all x_t are framed. Suppose now that s does not satisfy the condition of the proposition for some $a < b$. Since $b - a < 2$, $\mu_s((a, b)) = 0$, as all discrete values taken by s must differ by at least 1. We now show the existence of a positive length interval $J \subseteq (a, b)$ such that $\mu_{x_{t+1}}(J) \geq \mu_{x_t}(J)$ if $x_t \in B_\mu(s, \epsilon)$ for a sufficiently small $\epsilon > 0$. Since x_t converge to s, there is thus a finite time t^* after which $\mu_{x_t}(J)$ is nondecreasing. But since this value converges to $\mu_s(J) \leq \mu_s((a, b)) = 0$, there must hold $\mu_{x_{t^*}}(J) = 0$. This however contradicts the fact that x_t is framed for all finite t.

Let $c = \frac{\mu_s(a)a+\mu_s(b)b}{\mu_s(a)+\mu_s(b)}$ be the weighted average between a and b, and take $J := [c-\delta, c+\delta]$ for a positive δ such that $c-\delta+1 > b$ and $c+\delta-1 < a$. For any $x \in B_\mu(s, \epsilon)$ for some $\epsilon > 0$, there hold

$$\begin{aligned}
\mu_x([a-\epsilon, a+\epsilon]) &\in [\mu_s(a)-\epsilon, \mu_s(a)+\epsilon], \\
\mu_x([b-\epsilon, b+\epsilon]) &\in [\mu_s(b)-\epsilon, \mu_s(b)+\epsilon], \\
\mu_x\left((a-1, b+1) \setminus ([a-\epsilon, a+\epsilon] \cup [b-\epsilon, b+\epsilon])\right) &\leq \epsilon,
\end{aligned}$$

where we have used the fact that all values taken by s are separated by at least 1. Suppose that ϵ is sufficiently small so that $c-\delta+1 > b+\epsilon$ and $c+\delta-1 < a-\epsilon$, which implies that for every $y \in J$, $(a-\epsilon, b+\epsilon) \subseteq (y-1, y+1)$. For any such y, a lower bound on $u_x(y)$ is obtained by considering the situation in which $\mu_x(a-\epsilon) = \mu_s(a)+\epsilon$, $\mu_x(b-\epsilon) = \mu_s(b)-\epsilon$, and a set of measure at most ϵ distant from y by at most 1. This leads to

$$\begin{aligned}
u_x(y) &\geq \frac{(\mu_s(a)+\epsilon)(a-\epsilon)+(\mu_s(b)-\epsilon)(b-\epsilon)+\epsilon(y-1)}{\mu_s(a)+\mu_s(b)+\epsilon} \\
&= c-\epsilon+\epsilon\frac{(a-b+y-1-c+\epsilon)}{\mu_s(a)+\mu_s(b)+\epsilon} \\
&\geq c-\epsilon-\frac{5\epsilon}{\mu_s(a)+\mu_s(b)},
\end{aligned}$$

which for a sufficiently small ϵ is larger than $c-\delta$. Similarly, we have

$$\begin{aligned}
u_x(y) &\leq \frac{(\mu_s(a)-\epsilon)(a+\epsilon)+(\mu_s(b)+\epsilon)(b+\epsilon)+\epsilon(y+1)}{\mu_s(a)+\mu_s(b)+\epsilon} \\
&= c+\epsilon+\epsilon\frac{(b-a+y+1-c+\epsilon)}{\mu_s(a)+\mu_s(b)+\epsilon} \\
&\leq c+\epsilon+\frac{5\epsilon}{\mu_s(a)+\mu_s(b)},
\end{aligned}$$

which again for a sufficiently small ϵ is smaller than $c+\delta$. Therefore, there is an ϵ such that for any $x \in B_\mu(s, \epsilon)$ and any $y \in J$, $u_x(y) \in J$. This implies then that $\mu_{U(x)}(J) \geq \mu_x(J)$, which proves our result. □

The conditions under which a sequence of framed function maintain $\sup_\alpha x_t - \inf_\alpha x_t \geq 2$ need however still to be determined.

11.3 Relations between systems on discrete agents and on agent continuum

We now analyze to which extent the system (11.3) defined on an agent continuum can be viewed as a limiting case of the system (9.8) defined for discrete agents when the number of agents tends to infinity. A discrete system can be simulated by a system involving a continuum of agents. To represent a vector of discrete opinions $\hat{x} \in \Re^n$, take n subsets $I_1, \dots, I_n$ of measure $\frac{1}{n}$ defining a partition of I, and define a function $x \in X$ taking a value $\hat{x}(i)$ on I_i for each i. It follows then from (11.3) that all x_t are constant on these sets, and their value corresponds to the discrete opinions $\hat{x}_t$ obtained by the discrete system (9.8). This can be expressed using a function $f : I \to \{1, \dots, n\}$ assigning an agent number to each $\alpha \in I$. We have then $x_t(\alpha) = \hat{x}\left(f(\alpha)\right)$ for all t and $\alpha \in I$. Different weights can also be given to the discrete agents by varying the measures of the sets I_i. To further analyze the link between discrete and continuous system, we need the following results on the continuity of the update operator.

Proposition 11.5. *Let $x \in X$ be a framed function. Then the update operator U is continuous at x with respect to the norm $||.||_\infty$. In other words, for all $\epsilon > 0$ there is a δ such that $||y - x||_\infty \leq \delta$ implies $||U(y) - U(x)||_\infty \leq \epsilon$.*

Proof. Consider a framed function $x \in X$, and an arbitrary ϵ. Let δ smaller than $\frac{m\epsilon}{25M}$, where $m \leq M$ are the bounds coming from the definition of framed opinion functions applied to x. We show that if a function $y \in X$ satisfies $||x - y||_\infty < \delta$, then $||U(y) - U(x)||_\infty = ||u_x - u_y|| < \epsilon$ holds.
Take $\alpha \in I$, and call $S_x, S_y \subseteq I$ the set of agents connected to α according to the topologies C_x and C_y defined by x and y respectively. We let $S_{xy} = S_x \cap S_y$, $S_{x\setminus y} = S_x \setminus S_{xy}$ and $S_{y\setminus x} = S_y \setminus S_{xy}$. Since $||x - y||_\infty < \delta$ the values $|x(\alpha) - x(\beta)|$ and $|y(\alpha) - y(\beta)|$ differ by at most 2δ. Therefore, there hold

$$\begin{aligned} &[x(\alpha) - 1 + 2\delta, x(\alpha) + 1 - 2\delta] \subseteq x(S_{xy}) \subseteq [x(\alpha) - 1 - 2\delta, x(\alpha) + 1 + 2\delta], \\ &x(S_{x\setminus y}) \subseteq [x(\alpha) - 1, x(\alpha) - 1 + 2\delta] \cup [x(\alpha) + 1 - 2\delta, x(\alpha) + 1], \\ &x(S_{y\setminus x}) \subseteq [x(\alpha) - 1 - 2\delta, x(\alpha) - 1] \cup [x(\alpha) + 1, x(\alpha) + 1 + 2\delta]. \end{aligned}$$

Since x is framed we have then $|S_{xy}| \geq m(2 - 4\delta) \geq m$ and $\left|S_{x\setminus y}\right|, \left|S_{y\setminus x}\right| \leq 4M\delta$, for some $M \geq m \geq 0$ independent of α. Let now $\bar{x}_{xy}, \bar{x}_{x\setminus y}$ be the average value of x on S_{xy} and $S_{x\setminus y}$ respectively, and $\bar{y}_{xy}, \bar{y}_{y\setminus x}$ be the average value of y on S_{xy} and $S_{y\setminus x}$. Since $||x - y||_\infty < \delta$, $\bar{x}_{xy}$ and $\bar{y}_{xy}$ differ by at most δ. It follows from the definition of the model (11.3) that

$$u_x(x(\alpha)) = \bar{x}_{xy} + \frac{\left|S_{x\setminus y}\right|}{\left|S_{xy}\right| + \left|S_{x\setminus y}\right|}(\bar{x}_{x\setminus y} - \bar{x}_{xy}),$$

and

$$u_y(y(\alpha)) = \bar{y}_{xy} + \frac{\left|S_{y\backslash x}\right|}{\left|S_{xy}\right| + \left|S_{y\backslash x}\right|}(\bar{y}_{y\backslash x} - \bar{y}_{xy}).$$

Since there trivially holds $\left|\bar{x}_{x\backslash y} - \bar{x}_{xy}\right| \leq 3$ and $\left|\bar{y}_{y\backslash x} - \bar{y}_{xy}\right| \leq 3$, we have then the following bound

$$\begin{array}{rcl} |u_y(y(\alpha)) - u_x(x(\alpha))| & \leq & |\bar{x}_{xy} - \bar{y}_{xy}| + 3\frac{\left|S_{y\backslash x}\right|}{\left|S_{xy}\right|} + 3\frac{\left|S_{x\backslash y}\right|}{\left|S_{xy}\right|} \\ & \leq & \delta + 6\frac{4M\delta}{m} \\ & \leq & \delta(1 + 24\frac{M}{m}) \leq \epsilon. \end{array}$$

where we have used the fact that $|\bar{x}_{xy} - \bar{y}_{xy}| \leq \delta$. Since this is true for any $\alpha \in I$, we have $||y - x||_\infty \leq \epsilon$. □

Let $U^t : X \to X$ be the composition of the update operator defined by $U^t(x) = U\left(U^{t-1}(x)\right)$, so that $U^t(x_0) = x_t$.

Corollary 11.1. *Let $x \in X$ be a framed function such that $\sup_\alpha U^t(x) - \inf_\alpha U^t(x) > 2$ for every $t \geq 0$. Then for any finite t, U^t is continuous at x with respect to the $||.||_\infty$ norm.*

Proof. Since x is framed and since for all t there holds $\sup_\alpha U^t(x) - \inf_\alpha U^t(x) > 2$, it follows from Proposition 11.3 that all $U^t(x)$ are framed. Proposition 11.5 implies then that U is continuous at all $U^t(x)$ so that their composition U^t is continuous. □

This result allows us to prove that for any finite time, the system on an agent continuum is the limit of the one defined for discrete agents when the number of agents grows.

Theorem 11.4. *Let $x \in X$ be a framed function such that $\sup_\alpha U^t(x) - \inf_\alpha U^t(x) > 2$ for every $t \geq 0$. Then the sequence $(U^t(x))$ is approximated with an arbitrary accuracy until any finite time t^* by a sequence $(\hat{x}_t)$ of opinion vectors evolving according to (9.8). In other words, for any t^* and $\epsilon > 0$, there is a vector $x_0 \in [0, L]^n$ and a function $f : I \to \{1, \dots, n\}$ taking constant values on set of measures $1/n$ such that for all $t \leq t^*$, there holds $||U^t(x) - \hat{x}_t(f)||_\infty \leq \epsilon$, where the sequence (x_t) satisfies (9.8).*

Proof. Let us be given a positive ϵ. Since all U^t are continuous at x, it there is a $\delta > 0$ such that for any $y \in X \cap B_\infty(x, \delta)$, $||U^t(y) - U^t(x)||_\infty \leq \epsilon$ for all $t \leq t^*$. Since x is framed, it is possible to divide $[0, L]$ in disjoint intervals $J_1, J_2, \dots, J_n$ for some finite n in such a way that $\mu_x(J_i) = 1/n$ and $|J_i| \leq \delta$ for each i. Define then $f : I \to \{1, \dots, n\}$ by $f(\alpha) = i$ if $x(\alpha) \in J_i$, and $\hat{x} \in [0, L]^n$ by taking a each value $\hat{x}_i$ in the corresponding set J_i. For any α there holds $|\hat{x}(f(\alpha)) - x(\alpha)| < \delta$. This proves our result as it implies $\hat{x}(f) \in B_\infty(x, \delta)$.

□

This Theorem supports the intuition that for large values of n, the continuous systems behaves approximatively as the discrete one for a certain number of time-steps. In view of Proposition 11.2, this suggests that the discrete system should always converge to a stable equilibrium (in the sense defined in Section 11.1) when n is sufficiently large, as stated in Conjecture 11.1. For example, Figure 11.2(a) and (b) show the evolution of respectively 501 and 5001 agent opinions initially approximating the same opinion function: $x(\alpha) = 5\alpha$ for $\alpha \in [0, \frac{1}{2}]$ and $x(\alpha) = 2.5 + \alpha$ for $\alpha \in [\frac{1}{2}, 1]$. In Figure 11.2(a), the system converges to an unstable equilibrium. The respective weight of the clusters are indeed 152 and 349. They are separated by a distance 1.399 while the stability condition of Theorem 11.1 requires here a distance larger than $1 + \frac{152}{349} \simeq 1.436$. As shown in Figure 11.2(b), this unstable equilibrium disappears when the number of agents increases. Indeed, some agents lying between the two temporary clusters cause them to merge. However the continuity argument above is not rigorous, because the continuity of U^t for all t does not imply the continuity of $U^\infty := \lim_{t\to\infty} U^t$.

Theorem 11.4 provides also a new insight on Hegselmann's conjecture [59]. This conjecture states that for every $L > 0$ there is a n' such that for any initial equidistant distribution of $n > n'$ discrete opinions, the system converge to one single cluster. Although numerical evidences seems to contradict it, the conjecture has not been proved or disproved yet. In view of Theorem 11.4, a sufficient condition for the conjecture to hold is that for any linear initial function $x \in X$, the sequence $U^t x$ converges to a function that is constant.

Finally, using the comparison between discrete and continuous systems, we can obtain a new result about the discrete one. Consider a discrete distribution $\hat{x}_0$ of n agents approximating a continuous distribution x_0 as above. Until any time step t, x_t is approximated arbitrarily well by $\hat{x}_t$ if n is sufficiently large, but x_t never reaches the equilibrium. For any t, there is thus a n above which $\hat{x}_t$ has not yet reached equilibrium. Therefore, by increasing the number n of agents in a discrete system (in a way that approximates a continuous function x_0), the convergence time will increase to infinity, even though it is finite for any particular finite n.

Chapter 12

Extensions

12.1 Higher dimensions and distance depending weights

A first natural extension of the model considered in Chapters 10 and 11 is to consider multi-dimensional opinions, as often proposed in the literature for various continuous opinion dynamics system (see [92] for a survey). The formulation of Krause's model and its continuous-time counterpart 9.9) remain formally the same, but the distance used to define[1] the neighborhood relation has to be specified. We use here the usual Euclidean distance. In discrete time, opinions are thus updated by

$$x_i(t+1) = \frac{\sum_{j:||x_i(t)-x_j(t)||<1} x_j(t)}{|j : ||x_i(t) - x_j(t)|| < 1\}|},$$

and in continuous time their evolution is described by

$$\dot{x}_i(t) = \sum_{j:|x_i(t)-x_j(t)|<1} (x_j(t) - x_i(t)) \,.$$

An example of two-dimensional opinions evolving according to Krause's discrete time model is presented in Figure 12.1.

Let us mention that simulating such multi-dimensional systems is much more time-consuming than the one-dimensional ones. This is first due to the curse of dimensionality, requiring the number of agents to grow as L^D to reach a constant discretization in D dimensions. Moreover, whereas a naive implementation of Krause's model requires $O(n)$ operations for *each* agent's value

[1]The use of "pseudo distances" that do not satisfy the triangular inequality is also acceptable.

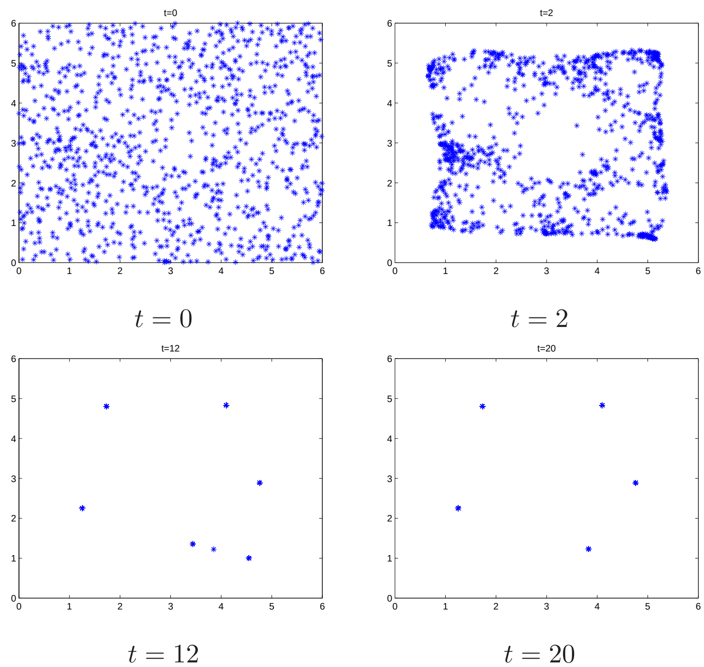

$t = 0$ $t = 2$

$t = 12$ $t = 20$

Figure 12.1: Representation of the evolution of a set of 1000 two-dimensional initially random opinions according to Krause's model (with the Euclidean distance). The opinions converge to clusters which are separated by more than 1. At $t = 12$, two clusters and some isolated agents in the lower right part of the figure are in a meta-stable situation. They eventually merge, as seen at $t = 20$.

update, the presence of an order in one dimension and its preservation at each iteration allow the computation of *all* agent updates in $O(n)$. This is however not generalizable to higher dimensions. As a result, the number of operations for one iteration grows for example as $O(n^4)$ for two-dimensional opinions, as compared to linearly for one-dimensional opinions. This quartic number can however be reduced in practice by the use of many heuristic methods, as partly explored by Kushagra Nagaish during his internship at the UCL in 2007 [102].

One could also use spaces with different topologies such as the circle or the torus. This could indeed be interesting as it avoids the effects caused by the edges of the initial distribution. Preliminary simulations on the circle had results comparable to those on an interval, with a few notable differences. The set of equilibria is indeed richer in the cyclic case. One can for example verify that opinions equidistantly distributed on the whole circle are at equilibrium, independently of the distance separating them. This does not contradict our convergence results, as the latter assumes (and uses the fact) that the opinions are real. This equidistant distribution equilibrium appears to have a nontrivial, though small, attraction basin. By applying a perturbation of sufficiently small amplitude, one does not change the interaction graph (unless the distance separating the opinions is exactly $1/m$ for some natural number m). The system may then re-converge to an equidistant distribution. Such re-convergence has been experimentally observed, but it is not known whether it always takes place when the interaction graph is not affected by the perturbation. When larger perturbations are applied, they appear to propagate on the whole domain, and the system converges then to a set of clusters, which satisfy our stability condition with a high probability. Note that all this discussion relies on numerical experiments, and the observations remain thus to be formally proved. We do however not treat the cyclic domain in the rest of this Chapter.

The second extension we consider is the introduction of distance-depending weights, or more generally relative position-depending weights, on the agent influence. Whereas in the initial model, the opinion of an agent j is taken into account by i if $|x_i - x_j| < 1$ and not taken into account otherwise, one can imagine a smoother way to weight the influences, giving for example more importance to the agents close to i than to the others, or having a continuous evolution of the importance with the distance. As suggested in [93], this can be generically represented by a nonnegative influence function $f : \Re \to \Re^+$, whose support [2] is supposed to be connected, to contain 0 and to have a positive length. The update rule in discrete time then becomes

$$x_i(t+1) = \frac{\sum_j f\left(x_j(t) - x_i(t)\right) x_j(t)}{\sum_j f\left(x_j(t) - x_i(t)\right)}, \tag{12.1}$$

[2] The support of a non-negative function is the set on which it takes positive values.

and the evolution in continuous time is described by

$$\dot{x}_i = \sum_j (x_j - x_i) f(x_j - x_i). \tag{12.2}$$

The initial models correspond to the situation where f is the indicator function of the interval $(-1, 1)$. Figure 12.2 shows examples of opinions initially identically distributed but evolving according to different influence functions. Note that when the support of f is not symmetric as in Figure 12.2(b), the communication graph is directed, as an agent may be influenced by another agent without necessarily influencing it. When f is not symmetric but has symmetric support, the communication graph remains symmetric, but some properties such as preservation of the average in continuous time are in general lost.

We believe it worth to explore how our results can or cannot be generalized to these more general situations. In addition to leading to more powerful results, this could also allow distinguishing the real reasons behind some phenomena from other conditions that happen to be equivalent in one dimension. Taking an example from the first part of this thesis, when studying the deformation of graph representations in one-dimension, a trivial necessary and sufficient condition for rigidity is connectivity. It is thus only by studying higher-dimensional systems that one is led to introduce powerful tools such as rigidity matrices, and Laman-type counting conditions.

We consider thus in Section 12.2 the convergence properties of these extended systems, and in Section 12.3 the stability of their equilibria. We finish in Section 12.4 by giving necessary and sufficient conditions for the order of one-dimensional opinions to be preserved.

Note that we do not consider here systems with heterogenous agents. In such systems, different agents may have different interaction radius, leading to interesting complex behaviors [93]. The communication graphs are then inherently directed.

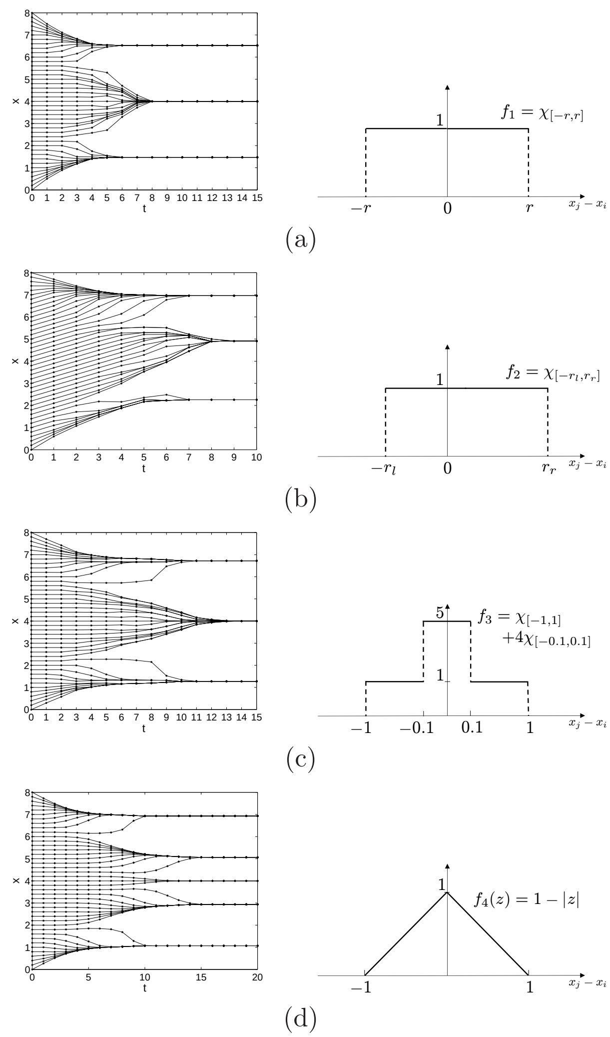

Figure 12.2: Evolution with time of 41 opinions initially equidistantly distributed on $[0, 8]$ according to the discrete-time model (12.1), for three different influence functions also represented. The model (b) represents asymmetric agent behaviors, as proposed in [60]. Unsurprisingly, opinions converge to clusters with higher value than in the symmetric represented in (a). In (c), the agents give a larger weight to those that are very close to them than to the others. One can see that convergence is then slower than with the usual model. Finally, (d) exhibits similar convergence properties as the other systems, although the influence decays continuously to 0.

12.2 Convergence

All experiments conducted on the two extended models presented above show that the opinions converge to clusters sufficiently distant so that no interaction takes place between them. Proving this convergence is however difficult in some cases.

12.2.1 Generic results

In discrete time, convergence to clusters that do not interact is guaranteed by the generic Theorem 9.4 as long as the interactions are symmetric, and that there is a positive lower bound on all positive coefficients involved in the average. These two conditions are clearly satisfied by the multi-dimensional extension of Krause's model, as the argument provided in Section 9.3 can directly be generalized. When distance depending functions are used as in (12.1), all interactions are symmetric provided that the function support is symmetric around 0, i.e., provided that $f(x) > 0$ if and only if $f(-x) > 0$. Observe now that the coefficient given to x_j in the computation of $x_i(t+1)$ is

$$a_{ij}(t) = \frac{f(x_j - x_i)}{\sum_{k=1}^{n} f(x_k - x_i)}.$$

There is thus a positive lower bound on all positive values that it can take provided that there exist positive upper and lower bound on the positive values of f. The presence of an upper bound is a rather natural assumption. The lower bound implies the presence of an influence threshold, so that the importance given by an agent to another cannot decay to 0, forbidding for example dependence inversely proportional to the squared distance. To summarize, Theorem 9.4 guarantees the convergence of systems with distance-depending influence provided that the influence function has a symmetric support, and admits positive lower and upper bound on its positive values. This result is also valid if distance-depending influence function are used for multidimensional opinions. The convergence does however not necessarily happen in finite time, as all-to-all communications do not necessarily imply that all agents compute the same value, due to the possible different importance that they give to each other.

No generic convergence result is however known proving convergence of continuous-time systems in the possible absence of consensus. As in the one-dimensional case, there is thus no simple way to prove a priori the convergence of the multidimensional and the distance-depending importance systems.

12.2.2 Results based on the topology dependence

In Theorems 11.1 and 10.1, we have proved the convergence of Krause's model and its continuous-time counterpart based on the systems' particular dynamics and position-dependence interaction topology. The arguments can be generalized to systems with distance-depending influence, provided that the influence function is such that the order of opinions is preserved, that it admits positive lower and upper bound on its positive values and that its support is connected.

Upper and lower bounds are indeed needed to ensure that if the $a_{ij}(x_j - x_i)$ decays to 0, either the distance between x_i and x_j decays to 0, or i and j become disconnected and do not interact anymore after a finite time. We hope however to be soon able to extend the results to function admitting no positive lower bound.

The order preservation is needed in the convergence proof to be able to designate a first agent whose opinion can never decrease. We do however not know if this is an essential condition for convergence or for this type of proof to work, or if it just allows a simpler writing of a slightly more general argument. The conditions under which the order of opinions is preserved are detailed in Section 12.4.

Note that the symmetry of the influence function or of its support is not required for our proof to be extended.

These proofs can however not immediately be generalized to multi-dimensional opinions, as they appear to strongly rely on the existence of an order. Two key arguments are indeed that the smallest opinion is nondecreasing and is therefore convergent, and that the system can be decomposed into two subsystems when two consecutive opinions are separated by more than 1. It should however be explored if the smallest opinion can be replaced by the vertex of the opinions convex hull for multi-dimensional systems, and if the separation idea can be applied to separate a group of agents whose opinions are sufficiently far from the convex hull of all other agent opinions.

To summarize, our convergence results based on the topology evolution easily extend to one-dimensional systems with relative-position-depending influence provided that the influence function admits positive lower and upper bound on its positive values, and that the order of opinions is preserved. This is valid for both continuous and discrete-time systems. Moreover, the generic Theorem 9.4 allows us to prove the convergence of multi-dimensional systems in discrete time. This is also valid if relative-position dependent influence functions are used, provided that they have symmetric support, and again admit positive lower and upper bound on their positive values.

12.2.3 Agent continuum

Theorems 10.2 and 11.2 on the decay of $\dot{x}$ and $x(t+1) - x(t)$ for agent continuum can both be generalized to multi-dimensional opinions and to distance-depending influence functions, provided that the latter functions are symmetric and are decreasing with the distance. This is done by extending the operator formalism for these new systems, and performing then the same formal computations.

For multi-dimensional opinions, one just needs to use the scalar product $\langle x, y\rangle = \int_{\alpha\in I} x(\alpha)^T y(\alpha)d\alpha$ instead of the usual one $\langle x, y\rangle = \int_{\alpha\in I} x(\alpha)y(\alpha)d\alpha$. Moreover, the energy functions used in both proofs should be

$$V(x) = \frac{1}{2}\int_{(\alpha,\beta)\in I^2} \min\left(1, ||x(\alpha) - x(\beta)||^2\right).$$

Moreover, one can also prove that the variance is non-increasing in continuous time, as there holds $Var(x) = \langle x, x\rangle - ||\bar{x}||^2$, which leads to

$$\dot{Var}(x) = 2\langle x, \dot{x}\rangle - 2\dot{\bar{x}}^T\bar{x} = -2\langle x, L_x x\rangle \leq 0$$

where the last inequality comes from the fact that L_x is positive semi-definite. Note that it may be necessary to use a multi-dimensional set I to index the agents. The dimension of this set does however not necessarily need to be the same as the dimension of the opinions. One can consider for example opinions forming a two-dimensional surface embedded in a three-dimensional space.

When distance depending influence functions are used, the same formalism holds, but one needs to redefine the function $\chi_x : I \times I \to \Re$ introduced after equation (10.3) to model the interaction topology. In the initial model, this function defined on pair of agent indices takes a value 1 if the agents are neighbors influencing each other, and 0 else. When continuous evolution of the interaction strength is considered, one can verify that it suffices to take the function representing this evolution: $\chi_x : I \times I \to \Re : (\alpha, \beta) \to \chi_x(\alpha, \beta) = f\left(|x(\alpha) - x(\beta)|\right)$. For the symmetry property of the different operators to holds, f needs however to be symmetric. One can then for example prove the non-increase of the opinion variance formally exactly in the same way as in Section 10.2. As already observed, the initial model corresponds thus to the case where f is the indicator function of the interval $(-1, 1)$.

Finally, one should use in Theorem 10.2 and 11.2 the following general energy function

$$V(x) = \int_{(\alpha,\beta)\in I^2} F\left(x(\alpha) - x(\beta)\right) \geq 0, \tag{12.3}$$

where $F(w) = \int_{v=0}^{w} f(v)vdv$. One can again verify that this expression reduces to (10.5) when the indicator function of $(-1, 1)$ is used for f. For the reasoning of the two theorems, f needs to be nondecreasing, so that for any v there holds $F(v) \leq F(w) + \frac{1}{2}(v^2 - w^2)f(w)$. Using $\chi_x(\alpha, \beta) = f\left(x(\alpha) - x(\beta)\right)$ and (12.3) we can then indeed bound $V(y)$ by

$$\begin{aligned} & V(x)+\tfrac{1}{2}\int_{(\alpha,\beta)\in I^2}\left(\left(y(\alpha)-y(\beta)\right)^2-\left(x(\alpha)-x(\beta)\right)^2\right)f\left(x(\alpha)-x(\beta)\right) \\ = \quad & \langle y, L_x y\rangle - \langle x, L_x x\rangle + V(x), \end{aligned}$$

a relation extending (10.6) on which both theorems are based. Note that these ideas can also be extended to multi-dimensional opinions with distance depending influence, taking in the definition of the energy function a F such that $\nabla F(v) = vf(v)$, if such a F exists.

12.3 Equilibrium stability

Experiments conducted for the extended models also show that the clusters to which the opinions converge are usually separated by distances significantly larger than the interaction radius. We have proposed an explanation of this phenomenon for the initial models based on a notion of equilibrium stability with respect to the addition of an agent. We now consider the possible extension of this stability notion to systems with multi-dimensional opinions and/or our notion of equilibrium stability. This analysis is exploratory, and several approximations are made. We believe that it leads to the correct characterization of stable equilibrium up to some border and zero-measure set effects, but this characterization and the related results should nevertheless not be considered as proved. For the sake of simplicity, we only consider the continuous-time system. We expect the discrete-time model to present similar equilibrium properties, although some more complex phenomena could be caused by the discontinuity of the opinion evolutions.

12.3.1 Generic analysis

Remember that an equilibrium is stable if its perturbation by addition of an arbitrary small agent cannot cause a fixed size deformation after re-convergence. The motivation behind this definition was that in the presence of sufficiently many smoothly distributed agents, there will most likely be some agent playing the perturbing role during the convergence process, preventing the system to converge to an unstable equilibrium. Intuitively, the presence of these agents comes from the fact that if the function describing the opinion distribution is smooth, agent opinions are likely to be found everywhere on the interval of opinions, including at the position where they could destroy the unstable equilibrium. While this intuitive behavior can still be applied when distance-depending influence functions are used, the multi-dimensional case presents more difficulties. In a multi-dimensional space, connectivity and convexity are not equivalent. The smoothness of the initial opinion function ensures the connectivity of the set where opinions can be found, but not necessarily its

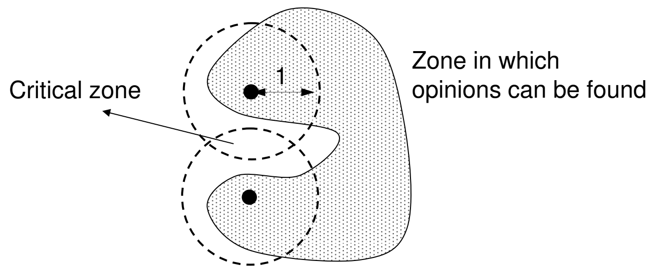

Figure 12.3: Representation of a possible (generic) convergence to two clusters in an unstable configuration. The presence of any agent in the critical zone at the intersection between the two cluster's influence disc would destroy this possible equilibrium by making the cluster merge. If the zone in which the opinions can be found in not convex however, there could never be any agent in that critical zone, independently of the number of agent. Such phenomenon cannot happen in one dimension, as the connectivity of the zone with opinions implies its connectivity.

convexity. As a result, we may not a priori exclude systems converging to equilibrium where two clusters are so close that the addition of any perturbing agent in a certain zone would make them merge, because it might be that no agent is to be found in that zone during the convergence process, even if the number of agents if very large, as represented in Figure 12.3.

Let us (re-)define more formally the notion of equilibrium, let $\bar{x}$ be a vector of agent opinions at equilibrium. Suppose that one adds a new agent 0 of weight δ and opinion initially set at $\tilde{x}_0$, and let the system re-converge to a perturbed equilibrium. One then removes the perturbing agent. The opinion vector $\bar{x}'$ so obtained represents still an equilibrium. We denote by $\Delta_{\tilde{x}_0,\delta} = w^T \left|\bar{x} - \bar{x}'\right|$ the distance between the initial and perturbed equilibria. The equilibrium $\bar{x}$ is stable if $\max_{x_0} \Delta_{\tilde{x}_0,\delta}$, can be made arbitrarily small by choosing a sufficiently small δ.

Since the perturbing agent 0 has an arbitrary small weight, its influence on the other agents is also arbitrary small during any fixed time-interval. The influence that the other agents have on 0 is however approximately independent of its weight δ. The perturbation of the equilibrium involves thus two distinct phenomena with different time-scales. First the agent added moves alone while all other agents are approximately fixed. The way it moves in this phase is as a first order approximation independent of its weight. If it is not isolated and does not converge to one cluster, it eventually converges to a position where it

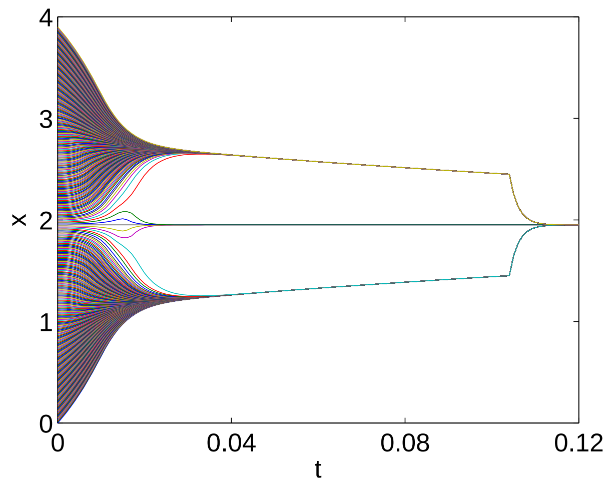

Figure 12.4: Example of temporary "meta-stable" equilibrium. Initially, two clusters are formed and do not interact with each other, but they both interact with a small number of agents lying between them. As a result, the distance separating them eventually becomes smaller than 1. The clusters then attract each other directly and merge into one larger cluster. This example is also shown in Figure 10.2.

interacts with two or more clusters, whose attractions on it cancel each other. A second phase takes then place on a much slower time-scale. The agent remains at equilibrium[3] as the result of different cancelling attractions, and it slowly attracts the clusters. The strength of this attraction is of the order of δ, but it takes place during an arbitrary large interval, so that large deformations are eventually caused. An example of this second phase is shown in Figure 12.4.

To the exception of some border and zero-measure set effects, large variations of the equilibrium are guaranteed if the second phase begins. In that phase indeed, the perturbing agent is at equilibrium and constantly attracts some clusters while no other interaction takes place. Supposing that no large variation occurs, this constant attraction would take place forever and therefore eventually causes a large variation, contradicting our hypothesis. The stability of the equilibrium depends thus only on the possibility for the perturbing agent to reach in the first phase an equilibrium position or more generally an

[3]In general (for distance depending influence), the agent does not remain exactly at the same position, but constantly adapts it due to the slow move of the clusters. Since this adaptation takes place at a faster time-scale, we can suppose that the agent remains at an equilibrium point which evolves at the same speed as the clusters.

invariant set, from which it can apply a positive attraction on two or more clusters. Note that this condition implicitly assumes that large variation of the attractions by clusters cannot be caused by arbitrarily small variation of the positions, a condition which may not always be verified.

To represent the first phase, we consider the evolution of a virtual agent with weight 0. More formally, let x be a set of n $D-$dimensional opinions at equilibrium for a continuous time system of opinion dynamics where the influence depends on the relative position of the agents via an influence function $f : \Re^D \to \Re^+$. We define the vector field

$$v_x : \Re^D \to \Re^D : y \to v_x(y) = \sum_{j=1}^{n} f(x_j - y)(x_j - y).$$

Observe that the evolution of a virtual perturbing agent with weight 0 added to the equilibrium x is described by $\dot{y} = v_x(y)$. According to our approximate two-phases analysis, the equilibrium x is unstable if and only if there is a non-trivial set invariant for the system $\dot{y} = v_x(y)$, and included in the interaction zone of at least two clusters. In other words, a set from which the virtual agent would never escape, and in which it would constantly be interacting with at least two (same) clusters. The existence of a non-trivial invariant set is required instead of a simple equilibrium point to avoid unstable equilibrium points to which no trajectory converges but the one starting on it.

12.3.2 One-dimensional opinion stability

Let us consider an influence function f and one-dimensional opinions whose evolution is described by (12.2). We scale the system so that the support of f is $[-1, 1]$. In one dimension, the perturbing agent can at most interact with two clusters as clusters are separated by at least 1, so that the stability condition can be expressed with respect to the positions and weights of each pair of clusters. Consider thus two clusters A, B of weight w_A, w_B and positions $x_A \leq x_B$. According the criterion derived above, two such clusters are in an unstable configuration if there is a set $S \subseteq [x_B - 1, x_A + 1]$ invariant for the system $\dot{y} = v_x(y)$. One can verify that this is equivalent to the existence of two points $m < p$ in $[x_B - 1, x_A + 1]$ such that $v_x(m) \geq 0$ and $v_x(p) \leq 0$, where $v_x(y)$ is on $[x_B - 1, x_A + 1]$ equal to $w_A(x_A - y)f(x_A - y) + w_B(x_B - y)f(x_B - y)$. Note that this implies the stability of any equilibrium in which all clusters are separated by at least 2, independently of f.

As a first example, suppose that f is equal to 1 on all its support, corresponding to our continuous-time model studied in Chapter 10. One can verify that on $[x_B - 1, x_A + 1]$, v_x takes its smallest value $-w_A + w_B(x_B - x_A - 1)$

on $x_A + 1$ and its largest one $w_B - w_A(x_A - x_B + 1)$ on $x_B - 1$. Since $x_B - 1 \leq x_A + 1$, it is thus stable if and only if $-w_A + w_B(x_B - x_A - 1) \leq 0$ and $w_B - w_A(x_A - x_B + 1) \geq 0$ hold, that is, if $x_B - x_A \geq 1 + \frac{\min\{w_A, w_B\}}{\max\{w_A, w_B\}}$, the condition obtained in Theorem 10.1.

Consider now that the influence decays linearly to 0 with the distance, that is, $f(z) = 1 - |z|$. For $y \in [x_B - 1, x_A + 1]$, there holds $v_x = w_A(x_A - y)(1 - y + x_A) + w_B(x_b - y)(1 + x_B - y)$. Observe that $v_x(x_B - 1) \leq 0$ and $v_x(x_A + 1) \geq 0$. Since v_x is a quadratic function of y, it cancels exactly once between $x_B - 1$ and $x_A + 1$. It contains thus an equilibrium point, but this equilibrium is a repulsive one, as $\dot{y}$ is positive for values of y above it and negative for values below it, as represented in Figure 12.5. Every configuration of clusters defining an equilibrium is thus stable, as it is impossible to find $m, p \in [x_B - 1, x_A + 1]$ with $m < p$ and $v_x(m) \geq 0$ and $v_x(p) \leq 0$. Similarly, there is no unstable equilibrium for any function f such that $|z|\, f(z)$ is decreasing with $|z|$, as $v_x(y)$ is then increasing with y on $[x_B - 1, x_A + 1]$. When such functions are used, the stability condition does thus not forbid convergence to clusters that are separated by distances close to 1. Figure 12.6 shows the evolution of opinions according to the 12.2 with $f(z) = 1 - |z|$. All inter-cluster distances are significantly larger than 1. This suggests that equilibrium stability is not the only reason explaining the distance observed between clusters, and that other dynamical phenomena are probably involved. The existence of such other phenomena could also be deduced from the fact that in the usual model, the stability condition explains why the distance between clusters of same size should be larger than 2, but not why it is often close to 2.1 or 2.2.

Consider now an influence function taking the values $f(z) = |z|^k$ for some $k \geq -1$. If k ≥ 0, The importance given to an agent's opinion grows thus with the difference of opinion, while it decays if $k \leq 0$. If $k < -1$, every equilibrium is stable as $|z|\, f(z)$ is then decreasing. For $y \in [x_B - 1, x_A + 1]$, there holds $v_x(y) = w_B(y - x_B)^k - w_A(y - x_A)^k$, and v_x is thus decreasing with y. So it takes its smaller value at $x_B - 1$ and its larger one at $x_A + 1$. If both values have the same sign, the equilibrium is clearly stable as the instability condition requires two points $m, p \in [x_B - 1, x_A + 1]$ at which v_x takes alternative signs. On the other hand, if the values have different signs, there necessarily hold $v_x(x_B - 1) \geq 0$ and $v_x(x_A + 1) \leq 0$. The interval $[x_B - 1, x_A + 1]$ is thus an invariant set for $\dot{y} = v_x(y)$, so that the equilibrium x is unstable. Simple algebraic manipulations show then that the equilibrium is stable if and only if there holds $x_B - x_A \geq 1 + \left(\frac{\min\{w_A, w_B\}}{\max\{w_A, w_B\}}\right)^{\frac{1}{k+1}}$, generalizing Theorem 10.1. Observe that the strength of the stability condition increases with k.

The stability analysis appears to be similar to the one presented for the

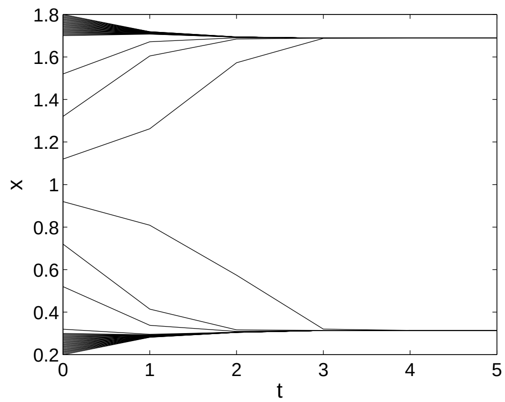

Figure 12.5: Representation of the stability of all equilibrium when the influence function $f(z) = 1 - |z|$ is used. Two clusters of 20 agents are formed at a distance 1.5 from each other. The presence of isolated agent between them does not prevent the convergence to these two clusters, as each of these agents is attracted by one of the clusters, and none remains at equilibrium between them. This could however have happened with a probability 0 if one of these agent was exactly at the average of the clusters.

initial model, but the stability criterion varies with the influence functions. Conditions for stability are for example more restrictive when $|z|\, f(z)$ increases strongly for z close to 1, that is, when the attraction of an agent on another is very strong when they are at the limit of the interaction zone. For such systems, it is indeed hard for the perturbing agent to escape from one cluster's attraction, so that it can easily remain connected to two clusters at equilibrium. On the other extreme, all equilibrium are stable when $|z|\, f(z)$ is decreasing. Experiments with such functions show however distances between consecutive clusters that are significantly larger than 1 although smaller than with the usual model, suggesting that equilibrium stability is not the only reason explaining the inter-cluster distances observed.

12.3.3 Multidimensional opinion stability

Our stability criterion formulated in terms of invariant sets can be applied to multi-dimensional opinions, but it does not appear to lead to so such conditions as in one dimension. In addition to the fact that an invariant set's boundary is not reduced anymore to two points, one has to take into account possible interactions with more than two clusters. With two-dimensional opinions for example, a perturbing agent can be connected to up to five clusters, and the

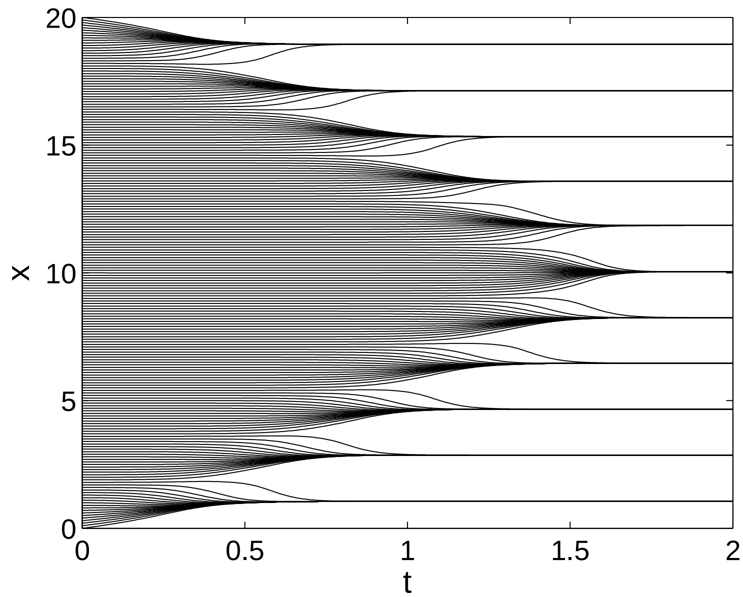

Figure 12.6: Evolution of 201 opinions initially equidistantly distributed between 0 and 20, according the discrete time system (12.1) with $f(z) = 1 - |z|$. All observed inter-cluster distances lie between 1.74 and 1.82. So although the equilibrium condition imposes no lower bound on the inter-cluster distances when this f is used, the observed distances are significantly larger than 1, even if they are sometimes closer to 1 than to 2. Note that a small random perturbations was added to all initial opinion to avoid nongeneric artifacts that happen with a probability 0 but can be observed when the system is totally symmetric.

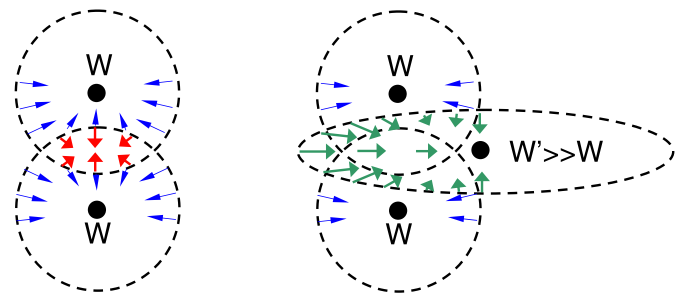

(a) (b)

Figure 12.7: The two clusters of equal weight in (a) are in an unstable configuration as the intersection of their influence disk is an invariant set. The arrows represent v_x. The configuration (b) obtained from (a) by adding a cluster of much higher weight is however stable. The contribution of this new cluster to v_x dominates indeed all other contributions, so that every agent in its influence disk is attracted by it. As a result, there is no more invariant set included in the influence disks of two or more clusters. The arrows represent the field v_x. The curvature of the third cluster disk of influence is exaggerated for presentation purpose.

consequent interaction may be nontrivial even with the usual model without influence function.

Consider for example an equilibrium consisting of two clusters of equal weight W. Such equilibrium is stable if and only if the distance separating the clusters is no smaller than 2. In that case indeed, there is no set and a fortiori invariant set of points at distance smaller than 1 from both cluster. If the distance is smaller than 2 on the other hand, one can verify that the intersection of the disc centered on the clusters constitutes such an invariant set, as shown in Figure 12.7(a). Suppose now that the clusters are separated by a distance smaller than but close to 2, and that a third cluster is added, equidistant to them, out of the invariant set, but at a distance smaller than 1 from all points in it, as represented in Figure 12.7(b). If the weight of this cluster is sufficiently larger than W, its contribution to the field v_x is larger than all other contributions in its disc of influence, which includes the former invariant set. As a result, the system $\dot{y} = v_x(y)$ does not contain any invariant set that cause instability, and any initially non-isolated perturbing agent converges to one of the three clusters, depending on its initial condition. So the addition of a large cluster to an unstable equilibrium stabilizes it, a phenomenon that cannot happen in one dimension.

12.3.4 Stability quantification

As a final remark, our stability study has so far be binary. An equilibrium is either stable or unstable. One might however expect that two clusters separated by a distance close to the stability limit might be "less unstable" than clusters separated by a distance close to 1. Intuitively, if the distance is small, the zone in which a perturbing agent can destroy the equilibrium is large, so that the addition of such an agent at a random position leads to a large variation with a high probability. On the other hand, if the distance is close to the stability condition, the probability for a perturbing agent with random initial position to be in the critical zone is smaller. The convergence to a highly unstable system thus seems improbable as the probability that some agents are in a zone where they prevent this convergence is high, while one may expect to see convergence to slightly unstable equilibria having escaped destruction due to the small size of their critical zones and a number of agent that is not too large. An example of such convergence is shown in Figure 10.4(a). A possible formalization of this stability quantification could be an interesting research opportunity.

12.4 Order preservation

Several proofs in Chapters 10 and 11 use the fact that the order of the opinions is preserved. Whether this preservation is essential for the result to hold or just allows a simpler writing still needs to be determined, and might depend on the results It is in any way consistent with the intuition that agent opinions that evolve unconstrained in a one-dimensional space according to the same rules should not cross each other. But, systems with distance depending influence as (12.1) do not necessarily preserve the opinion order, as shown in Figure 12.8. We therefore provide in this section necessary and sufficient condition for the order to be preserved.

Observe first that the continuous-time system (12.2) preserves the opinion order for any influence function f. Suppose indeed that $x_i(0)$ is smaller than $x_j(0)$. If at some time there holds $x_i(t) < x_j(t)$, there exists by continuity a time t^* at which both opinions are equal. Since the evolution of the opinions only depend on their value (under some uniqueness of solution assumption), this implies that x_i and x_j remain then equal for all further time, contradicting the fact that they cross each other. This continuity-based argument cannot be applied in discrete time, but we have the following necessary and sufficient condition for order-preservation.

Theorem 12.1. *Let $f : \Re \to \Re^+$ be a function with connected support. The iteration (12.1) preserves the order of opinions for any initial condition if and*

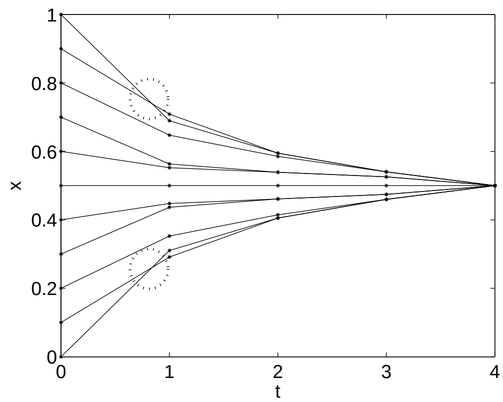

Figure 12.8: Evolution with time of 11 opinions initially equidistantly distributed on $[0,1]$ and following the model (12.1) with $f(x_i - x_j) = 5$ if $|x_i - x_j| < \frac{1}{10}$, $f(x_i - x_j) = 1$ if $\frac{1}{10} < |x_i - x_j| < 1$ and 0 else. The order of opinions is not preserved between $t = 0$ and $t = 1$.

only if

$$\frac{f(a+c)}{f(a)} \geq \frac{f(b+c)}{f(b)} \tag{12.4}$$

holds for all $a \leq b$, $c \geq 0$ *such that* $f(a), f(b), f(a+c), f(b+c) > 0$.

Proof. For the sake of conciseness, we denote $x(t)$ by x and $x(t+1)$ by x'. We first prove the necessity of the condition by constructing for any f not satisfying it an example where the order is not preserved.

Consider a system with $2n+2$ agents, where the agents 1 and 2 have opinions $x_1 = a$ and $x_2 = b$ respectively, for some $b > a$. Suppose also that among the remaining $2n$ agents, n have an opinion $a+b$ and n others an opinion $a+b+c$ for some $c > 0$. We suppose that 1 and 2 are both influenced by all other agents, that is, $f(a), f(b), f(a+c), f(b+c) > 0$. If n is sufficiently large, we can neglect the agents x_1 and x_2 in the computation of x_1' and x_2', which according to (12.1) are given by

$$\begin{aligned} x_1' &\simeq \tfrac{nf(b)(a+b)+nf(b+c)(a+b+c)}{nf(b)+nf(b+c)} = a+b+\tfrac{c}{1+f(b)/f(b+c)}, \\ x_2' &\simeq \tfrac{nf(a)(a+b)+nf(a+c)(a+b+c)}{nf(a)+nf(a+c)} = a+b+\tfrac{c}{1+f(a)/f(a+c)}. \end{aligned}$$

So if $\frac{f(a)}{f(a+c)} > \frac{f(b)}{f(b+c)}$, then $x_1' > x_2'$ although $x_1 \leq x_2$, and the order of opinions is thus not preserved.

To prove the sufficiency of the condition, we now consider system of n agents among which we select two agents a, b such that $x_b \geq x_a$ (we may possibly choose $a = b$). We suppose that $f(x_i - x_a) > 0$ and $f(x_i - x_b) > 0$ for all agents i. Since the system (12.1) is invariant under translation we assume that all x_i are nonnegative, and we relabel the agents in such a way that $x_1 \leq x_2 \leq \cdots \leq x_n$. The updated values of a and b are

$$x'_a = \frac{\sum_{i=1}^n f(x_i - x_a)x_i}{\sum_{i=1}^n f(x_i - x_a)}, \qquad \text{and} \qquad x'_b = \frac{\sum_{i=1}^n f(x_i - x_b)x_i}{\sum_{i=1}^n f(x_i - x_b)}.$$

As a consequence, $x'_b \geq x'_a$ holds if

$$\left(\sum_{i=1}^n f(x_i - x_b)x_i\right)\left(\sum_{i=1}^n f(x_i - x_a)\right) \geq \left(\sum_{i=1}^n f(x_i - x_a)x_i\right)\left(\sum_{i=1}^n f(x_i - x_b)\right)$$

holds. This can be rewritten as

$$\begin{array}{llll}
 & \left(\sum_{i=1}^{n-1} f(x_i - x_b)x_i\right)\left(\sum_{i=1}^{n-1} f(x_i - x_a)\right) & + & f(x_n - x_b)x_n\left(\sum_{i=1}^{n-1} f(x_i - x_a)\right) \\
+ & f(x_n - x_a)\left(\sum_{i=1}^{n-1} f(x_i - x_b)x_i\right) & + & f(x_n - x_b)f(x_n - x_a)x_n \\
\geq & \left(\sum_{i=1}^{n-1} f(x_i - x_a)x_i\right)\left(\sum_{i=1}^{n-1} f(x_i - x_b)\right) & + & f(x_n - x_a)x_n\left(\sum_{i=1}^{n-1} f(x_i - x_b)\right) \\
+ & f(x_n - x_b)\left(\sum_{i=1}^{n-1} f(x_i - x_a)x_i\right) & + & f(x_n - x_a)f(x_n - x_b)x_n.
\end{array}$$

For $n = 1$ (and thus $a = b = 1$), this relation reduces to $f(0)^2 x_1 \geq f(0)^2 x_1$ and is trivially satisfied. Suppose now that it holds for $n-1$, then it also holds for n provided that

$$\begin{array}{lll}
 & f(x_n - x_b)x_n\left(\sum_{i=1}^{n-1} f(x_i - x_a)\right) & +f(x_n - x_a)\left(\sum_{i=1}^{n-1} f(x_i - x_b)x_i\right) \\
\geq & f(x_n - x_a)x_n\left(\sum_{i=1}^{n-1} f(x_i - x_b)\right) & +f(x_n - x_b)\left(\sum_{i=1}^{n-1} f(x_i - x_a)x_i\right).
\end{array} \tag{12.5}$$

holds. Reorganizing the terms of (12.5) and dividing them by $f(x_n - x_a)f(x_n - x_b)x_n > 0$ yields

$$\sum_{i=1}^{n-1}\left(\frac{f(x_i - x_a)}{f(x_n - x_a)} - \frac{f(x_i - x_b)}{f(x_n - x_b)}\right) \geq \sum_{i=1}^{n-1}\frac{x_i}{x_n}\left(\frac{f(x_i - x_a)}{f(x_n - x_a)} - \frac{f(x_i - x_b)}{f(x_i - x_n)}\right).$$

Since all x_i are nonnegative and no greater than x_n, it is sufficient for this relation to hold that $\frac{f(x_i - x_a)}{f(x_n - x_a)} \geq \frac{f(x_i - x_b)}{f(x_n - x_b)}$ holds for all i. Since $x_a \leq x_b$, the latter is always true if f is such that $\frac{f(a+c)}{f(a)} \geq \frac{f(b+c)}{f(b)}$ holds for any $b \geq a$ and $c \geq 0$ for which $f(a), f(b), f(a+c), f(b+c) > 0$. It suffices indeed to take $a = x_i - x_b$, $b = x_i - x_a$ and $c = x_n - x_i$.

Suppose now that there is some i for which $f(x_i - x_a) > 0$ and/or $f(x_i - x_b) > 0$ does not hold. Let J_a be the set of agents i such that $f(x_i - x_a) > 0$,

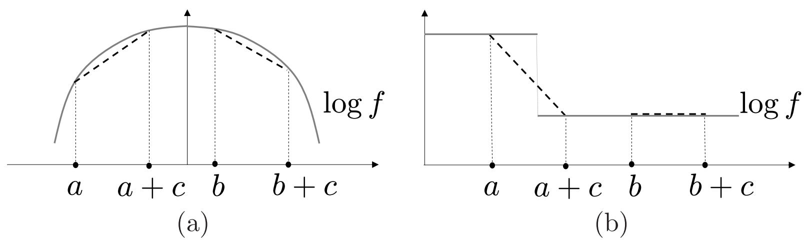

Figure 12.9: Illustration of the condition of Theorem 12.1. The function whose logarithm is presented in (a) preserves the order of opinion, as for any $a \leq b$ and $c \geq 0$, there holds $\log f(a+c) - \log f(a) \geq \log f(b+c) - \log f(b)$. On the other hand no piecewise constant function taking more than one value as in (b) preserves the opinion order, as there exists $a \leq b$ and $c \geq 0$ such that $\log f(a+c) - \log f(a) < \log f(b+c) - \log f(b)$ holds. In particular, $f = 4\chi_{[-\frac{1}{10},\frac{1}{10}]} + \chi_{[1,1]}$ used in Figure 12.8 does not preserve the order of opinions.

J_b the corresponding set for x_b and $I = J_a \cap J_b$. If $I = \varnothing$, then any value of J_b is larger than all values of J_a as the support of f is connected, so that $x'_b \geq x'_a$ trivially holds. If $I = J_a \cup J_b$, we have seen that the condition (12.4) is for $x'_b \geq x'_a$ to hold. Finally, observe that the presence of agents in $J_b \setminus I$ or in $J_a \setminus I$ only increases x'_b or decreases x'_a, so that this condition is still sufficient for $x'_b \geq x'_a$ to hold. $\square$

A similar proof shows that the same condition is necessary and sufficient for the order of opinions to be preserved by the following iteration defined on an agent continuum.

$$x_{t+1}(\alpha) = \frac{\int_{\beta \in I} x_t(\beta) f\left(x_t(\beta) - x_t(\alpha)\right) d\beta}{\int_{\beta \in I} f\left(x_t(\beta) - x_t(\alpha)\right) d\beta} \tag{12.6}$$

The condition of Theorem 12.1 can be re-expressed in term of $\log f$, as $\log f(a+c) - \log f(a) \geq \log f(b+c) - \log f(b)$. It thus means that the opinion order is preserved if and only if the increase in the value of $\log f$ by performing a step of fixed size c is non-increasing. This is represented on two examples in Figure 12.9, one of them encapsulating the function used in Figure 12.8 where we have seen that the opinion order was not preserved.

We now show that under some very weak and natural smoothness assumptions, the condition of Theorem 12.1 is equivalent to the concavity of $\log f$. This equivalence can easily be obtained if we assume that $\log f$ is differentiable but, this would be a strong restriction, as some of the functions that we have considered are not continuous on their support. It would for example not allow us to treat the functions used in Figures 12.8 and 12.2(c). Note that in the sequel we always implicitly assume that the points at which $\log f$ is evaluated belong to the support of f.

Proposition 12.1. *If* $\log f$ *is concave, then it satisfies the condition (12.4) that* $\frac{f(a+c)}{f(a)} \geq \frac{f(b+c)}{f(b)}$ *holds for all* $a \leq b$, $c \geq 0$ *such that* $f(a), f(b), f(a+c), f(b+c) > 0$.
If the condition (12.4) is satisfied and if $\log f$ *is continuous, then it is concave.*

Proof. Suppose first that $\log f$ is concave. Then for $b > a$ and $c > 0$, there hold

$$\begin{array}{lll} \log f(a+c) & \geq \frac{b-a}{b-a+c} \log f(a) & + \frac{c}{b-a+c} \log f(b+c), \\ \log f(b) & \geq \frac{c}{b-a+c} \log f(a) & + \frac{b-a}{b-a+c} \log f(b+c). \end{array}$$

This implies that $\log f(a+c) - \log f(a) \geq \log f(b+c) - \log f(b)$ holds, or equivalently that $\frac{f(a+c)}{f(a)} \geq \frac{f(b+c)}{f(b)}$.

To prove the second part of this result, consider now an arbitrary function f and let then $x < y < z$ be arbitrary points of its support such that $(y-x)/(z-x)$ is rational. There exists two integers m, n such that

$$\frac{z-y}{n} = \frac{y-x}{m} =: c > 0.$$

If f satisfies condition (12.4), then

$$\log f(a+c) - \log f(a) \geq \log f(b+c) - \log f(b).$$

holds for any $a < b$ in its support of f. So we have

$$\begin{array}{lll} \log f(y) - \log f(x) & = & \sum_{j=1}^{m} \left(\log f\left(x+jc\right) - \log f\left(x+(j-1)c\right)\right) \\ & \geq & m \left(\log f\left(y+c\right) - \log f\left(y\right)\right), \end{array}$$

and

$$\begin{array}{lll} \log f(z) - \log f(y) & = & \sum_{j=1}^{n} \left(\log f\left(y+jc\right) - \log f\left(y+(j-1)c\right)\right) \\ & \leq & n \left(\log f\left(y+c\right) - \log f\left(y\right)\right), \end{array}$$

which implies that

$$\log f(y) \geq \frac{n}{n+m} \log f(x) + \frac{m}{n+m} \log f(z) = \frac{z-y}{z-x} \log f(x) + \frac{z-y}{z-x} \log f(z). \tag{12.7}$$

Since this holds for any x, y, z for which $(y-x)/(z-x)$ is rational, the concavity of $\log f$ follows from its continuity. $\square$

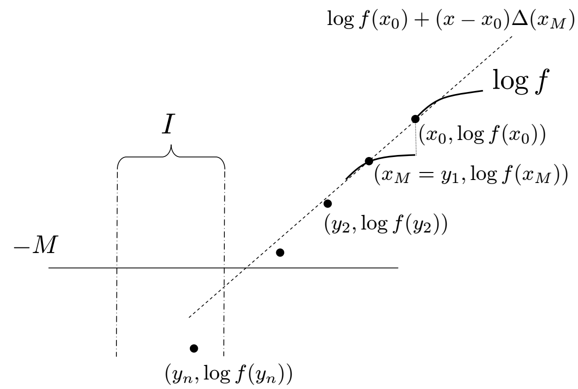

Figure 12.10: Illustration of the construction in the proof of Proposition 12.2. All $\log f(y_i)$ must be below $\log f(x_0) + (y_i - x_0)\Delta(x_M)$.

Functions that preserve the order of opinion but are not continuous have not been proved yet not to exist, but the next proposition shows that they would have so special properties that one should not expect to meet them in any practical situation.

Proposition 12.2. *Let f be a function satisfying the condition (12.4) that $\frac{f(a+c)}{f(a)} \geq \frac{f(b+c)}{f(b)}$ holds for all $a \leq b$, $c \geq 0$ such that $f(a), f(b), f(a+c), f(b+c) > 0$. If f not continuous on its support's interior, then it admits a positive lower bound on no positive length interval, and is as a consequence discontinuous everywhere on its support.*

Proof. Let S_f be the support of f, and suppose that f is not continuous at some x_0 in the interior of S_f. We prove that this implies the unboundedness of $\log f$ on all positive length intervals in S_f. In particular, $\log f$ is unbounded on $[x - \epsilon, x + \epsilon] \cap S_f$ for any $\epsilon > 0$ and $x \in S_f$, and therefore continuous at no point of S_f. This implies that f is also continuous nowhere on S_f, and admits a positive lower bound on no positive length interval.

Let $\Delta(x) = \frac{\log f(x) - \log f(x_0)}{x - x_0}$. The discontinuity of $\log f$ at x_0 implies that $\Delta(x)$ is unbounded on any open interval containing x_0. We suppose that it takes arbitrary large positive values any such interval. If it is not the case, then it necessarily takes arbitrary large negative values, and a similar argument can be applied. Consider an arbitrary large M and a interval $I \subset S_f$ of positive length $|I|$ with $\sup I < x_0$. The following construction is illustrated in Figure 12.10. There is a $x_M \in (x_0 - |I|, x_0 + |I|)$ such that $\left|\frac{M + \log f(x_0)}{\sup I - x_0}\right| < \Delta(x_M)$. Consider now the sequence of points defined by $y_0 = x_0$ and $y_i = y_{i-1} - |x_M - x_0|$.

Since all y_i are smaller than or equal to x_0, it follows from condition (12.4) that $\frac{f(y_{i-1})}{f(y_i)} \geq \frac{f(x_M)}{f(x_0)}$ holds if $x_M > x_0$ and $\frac{f(y_{i-1})}{f(y_i)} \geq \frac{f(x_0)}{f(x_M)}$ holds if $x_M < x_0$. In both cases, this implies that $\log f(y_{i-1}) - \log f(y_i) \geq |\log f(x_M) - \log f(x_0)|$ and thus that

$$\log f(y_i) \leq \log f(x_0) - i\,|\log f(x_M) - \log f(x_0)| = \log f(x_0) - (x_0 - y_i)\Delta(x_M).$$

Since $|x_M - x_0| < |I|$, there is a n such that $y_n \in I$. For this y_n, there holds $x_0 - y_n \geq |x_0 - \sup I| \geq \frac{M + \log f(x_0)}{\Delta(x_M)}$. It follows then from the inequality above that

$$\log f(y_n) \leq \log f(x_0) - (x_0 - y_n)\Delta(x_M) \leq \log f(x_0) - |x_0 - \sup I|\,\Delta(x_M) < -M.$$

Therefore, $\log f$ takes arbitrary large negative values on any positive length interval I with $\sup I < x_0$.

Consider now a $x_1 < x_0$. For any δ, $\log f$ takes arbitrary large negative values on $[x_1, x_1 + \delta]$, and therefore so does $\Delta_1(x) := \frac{\log f(x) - \log f(x_1)}{x - x_1}$. It follows then from a similar argument as above that $\log f$ admits no lower bound on any positive length interval I with $\inf I > x_1$, and thus that it does not admit any lower bound on any positive length interval contained in S_f since every such interval contains at least a subinterval I with $\inf I > x_1$ or with $\sup I < x_0$.

$\square$

Based on Propositions 12.1 and 12.1, we now have the following theorem, showing the almost equivalence between preservation of the opinion order and log-concavity.

Theorem 12.2. *Let $f : \Re \to \Re^+$ be a function with connected support. If $\log f$ is concave, then the iteration (12.1) preserves the order of opinions for any initial condition. Conversely, if this iteration preserves the order of opinions for any initial condition admits a positive lower bound on at least one positive-length interval or is continuous at one point of its support, then $\log f$ is concave*

From a purely theoretical point of view, one can still wonder if the concavity of $\log f$ is in general necessary and sufficient for the order of opinions to be preserved, or if there exists a function for which the order of opinions is preserved but that is discontinuous at every point and admits a positive lower bound on no positive-length interval. In view of Theorem 12.1, this leads to the following open question.

Open question 14. *Does there exists a function $f : \Re \to \Re^+$ with a connected support, that admits a positive lower bound on no positive-length interval but for which $\frac{f(a+c)}{f(a)} \geq \frac{f(b+c)}{f(b)}$ holds for all $a \leq b$, $c \geq 0$ such that $f(a), f(b), f(a+c), f(b+c) > 0$.*

Finally, whereas log-concavity of is well defined in any dimension, the usual notion of order only makes sense in one-dimensional spaces. Log-concave influence functions might however have a more generic property in which would imply order-preservation for one-dimensional spaces.

Chapter 13

Conclusions

13.1 On our analysis of opinion models

We have analyzed two models of opinion dynamics with endogenously changing topologies, one defined in continuous time and the other in discrete time. Our motivation was to study very simple systems with endogenously changing topologies taking explicitly into account the topology dynamics, which is usually not done for such systems. We particulary focused on understanding a phenomenon appearing in those systems that cannot be explained without taking the topology dynamics into account: The distance between clusters at equilibrium are usually significantly larger than 1, and typically close to 2. Strangely, this phenomenon appears in different models, such as the one of Deffuant. There could thus be a general reason for this, depending on some assumptions satisfied by all these opinion dynamics models. Our analysis is however limited to Krause's model and its continuous-time counterpart, and cannot be directly generalized to Deffuant's model.

We have proposed an explanation of the observed distances between clusters based on a notion of stability with respect to the addition of a perturbing agent. We have shown that stability is equivalent to the presence of a lower bound on the inter-cluster distances, which equals 2 when the clusters have identical weights. We conjecture that when the number of agents is sufficiently large, the system almost surely converges to a stable equilibrium. The intuition behind this conjecture is that when the number of agents is large, if two clusters are converging to an unstable configuration there is almost surely an agent in the region where it prevents convergence to the unstable equilibrium. Moreover, this conjecture and the intuition behind are experimentally confirmed.

It appears however that the stability of the equilibria cannot be the only reason for the observed distances between clusters. First, the value of the lower bound equals its maximum 2 only when consecutive clusters have the same weight. Experiments show that the clusters tend to have the same weight when the initial distribution is uniform, but this has not been proved nor theoretically analyzed. More importantly, the distances observed between the clusters are close to 2.2, while the stability can only explain why it should be larger than 2. When considering extensions of the models in Section 12.3.2, we have seen a system for which every equilibrium is stable, independently of the inter-cluster distance. For this model, the distances experimentally observed were between 1.35 and 2.35. Finally, the destruction of meta-stable situation where clusters are too close to each other is not always observed in practice, as the system often converges directly to a stable equilibrium. It thus seems that the system dynamics first rapidly drives almost all opinions to clusters, many of which are already separated by at least 2. Then some agents that are not yet in the clusters that have formed make some clusters merge when these are in an unstable configuration. Our stability analysis explains this second part. Nevertheless, it leads to the only nontrivial lower bound on the inter-cluster distances that has been provided till now.

To avoid granularity problems linked with the presence or absence of an agent in a particular region, we have introduced new opinion dynamics models allowing for a continuum of agents. For such models we could prove that, under some continuity assumptions, there is always a finite density of agents between any two clusters during the convergence process. As a result, we could prove that these systems never converge to an unstable equilibrium. For the discrete-time system, we have also proved that the behavior of the system defined on a continuum on any finite time-interval is indeed the limiting behavior of the system for discrete agents when $n \to \infty$. These two results support our conjecture that the discrete system tends to converge to a stable equilibrium when $n \to \infty$. In continuous time, we were unfortunately not able to obtain such a result comparing systems defined for discrete agents and continuum of agents. Besides, the continuity assumptions preventing the system to converge to an unstable equilibrium are stronger in continuous-time than in discrete-time.

The introduction of the system on agent continuum brings several interesting open questions and problems. The most important of them is probably the convergence of the system, for which we have only obtained partial results. We were indeed only able to prove that the variation speed of the system decays to 0, and that so does its distance to the set of fixed points. Moreover, we did not treat the issue of existence and uniqueness of a solution to the equation defining the continuous-time system for an agent continuum. A main difficulty for those system is precisely this absence of granularity, which was helpful to treat stability issues. Sets of agents converging to different opinions may indeed never be totally disconnected due to the (possible) continuity of the opinion distribution.

These systems on agent continuum led us to a broader question. The convergence of the discrete-time system on discrete agents can for example also be obtained from results on convergence of inhomogeneous stochastic product, i.e., convergence of $\lim_{t\to\infty} A_t A_{t-1} \dots A_1$ where all matrices are stochastic. To the best of our knowledge, no such result has been established for inhomogeneous compositions of stochastic operators, a particular case of which would be the convergence of our discrete-time system for agent continuum.

Finally, our analysis in this thesis was restricted to systems behaving autonomously, in the absence of control actions. Due to their simplicity, it could however be interesting to study control issues on these paradigm systems. One could for example study the cost of obtaining or preventing consensus by modifying some connections between agents. Another possibility would be to influence the via a small set of "controllable" agents, the opinion evolution of which would be described by

$$\dot{x}_i(t) = u_i(t) \sum_{j:|x_i(t)-x_j(t)|<1} (x_j(t) - x_i(t)) +,$$

where $u_i(t)$ is a control input.

13.2 On the mathematical multi-agent and consensus issues

Consensus systems are nice mathematical objects that are interesting and challenging to study. Their understanding is still very limited today. The models that we have considered in this thesis are among the simplest non trivial ones, and their analysis is far from being complete. Predicting theoretically the number of clusters and their positions based on the initial distribution remains for example generally open. Similarly, nothing is known on the convergence speed, or on the robustness of the system evolution with respect to modification of the initial conditions.

My personal belief is that we currently lack appropriate mathematical tools and formalism to efficiently analyze them. One possible way to obtain such tools could be to develop an approach based on a continuum of agents approximating a discrete distribution, similarly as what is done in Sections 10.2 and 11.2. This step could be similar to the one made between Newton's point-mass mechanics and the tools of fluid and solid-mechanics. In such a framework, our way of indexing the agents with a continuous variable would correspond to a

Lagrangian approach, while the formulations based on agent densities would be an Eulerian one. In a Lagrangian approach, the independent variable is indeed the particle, while in an Eulerian approach it is the position in space. It might also be that these multi-agent systems will be an opportunity to develop a formalism with which one can efficiently analyze large but finite number of discrete elements, a problem that is notoriously difficult in mathematics.

A more pessimistic, or more realistic, perspective is that these systems present so many different and complex behaviors, that they perhaps cannot be analyzed in a significantly nontrivial way with a general method. It would not be so surprising if some systems similar to Krause's model were Turing-equivalent. The latter model is however not Turing equivalent as it always converge to an equilibrium in finite time. Proving such equivalence, or the NP-hardness of questions related to the consensus and multi-agent models could be another promising, even though less application oriented, research direction.

13.3 On practical applications

The research in the domain of multi-agent systems can be separated in two parts, the analysis of multi-agent models, and the design of multi-agent systems or decentralized controllers.

Many of the multi-agent and consensus seeking models attempt to represent animal or human behaviors. One should obviously not consider that they have the same accuracy as those describing better understood physical phenomena, even if the rules proposed can be approximately observed under "laboratory conditions" [94]. These models can however provide some insight on the represented behavior, as for example the flocking models and their further generalization. Moreover, for large number of agents, some simple models can provide relatively efficient approximations of an average behavior, which in the absence of anything else can be very useful. Some simple discrete opinion models are for example used in marketing to understand the behavior of customers.

The design of systems reaching consensus, and more generally of decentralized controllers, is motivated by important perspectives of applications. The idea of accomplishing a task without having an identifiable entity that is "in charge" is moreover an attractive one. This idea actually only corresponds to the highest level of decentralization. Less decentralized system can involve for example a central controller using several sensors that are already preprocessing the information they send.

Decentralized techniques are already used or available for immediate use in applications such as automatic calibration and clock synchronization of networked sensors. Communication and peer-to-peer networks are also managed in a decentralized way, with different "players" negotiating to share the communication load.

We believe however that many of the methods developed now are not ready for practical application, and seem sometimes to concern "ideal laboratory conditions" not necessarily based on realistic practical constraints. Besides, the advantage of performing particular tasks with the presently available in a decentralized way as opposed to a centralized way is in our opinion not enough measured. To exemplify this, we now use three classical decentralized problems to raise different issues that we believe should more often be considered. They concern the initial conditions, the final use of the value computed, and the computational and communication capabilities of the agents.

When analyzing the rendezvous problem, one usually takes as initial condition a group of agents lying in the plane. Based on their capabilities and limitations, an algorithm is then designed to have all agents gather at one point. Some constraints appearing in possible practical applications may however render this approach less relevant. In real life, agents do not appear ex-nihilo in the plane with the purpose of gathering at one point. One should thus consider how the agents were evolving before beginning the rendezvous process, and how they were told to start it. If some entity has to tell the agents to start the process, it could as well specify a target point. Supposing that these agents were initially evolving on their own, they should most likely indeed have a navigation system allowing them to go to a specified position. Moreover, if the agents have to gather after accomplishing a mission, it is not unreasonable to suppose that the meeting point should not be any point produced by an algorithm, but some specified point at which it is relevant to gather all agents.

One could however of course imagine situations where those algorithms make sense as such. Consider for example agents that are required to explore some region or to accomplish a certain task on their own, and that only have a local navigation system unable to treat global coordinates. At some pre-specified time, or after receiving a message from some leading agents, they would then start to apply a rendezvous process. But we are not presently aware of any application were those conditions are met. Moreover, if such an application exists, one could then also consider taking the future rendezvous process into account in the first phase, making sure that the visibility graph remains connected, and that the agents do not go into configurations where the rendezvous process will be very slow. Finally, we could also imagine permanently applying a rendezvous method during a formation displacement. Suppose for example that all agents have a specified remote place where they should go and a global coordinated system, but that this system is relatively inaccurate[1]. We could then superpose a rendezvous process to the vehicle move to their remote

[1] This assumption a priori excludes vehicles moving in open air and using a GPS device.

target, to maintain the cohesion of the group during the displacement. This would probably lead to several practical issue needing to be solved.

To analyze the final use of the results, consider now a sensor network where the average of the sensed values need to be computed as described in Section 9.1. Since it is computed, we can suppose that this average value is either recorded or used to take some control actions. Unless each sensor is also a decentralized controller that takes actions based on the average value, this average value needs thus to be eventually communicated to some external device(s), which we may assume to have basic computational capabilities. One may then wonder what would be the purpose of involving all agents in an iterative decentralized average computation, while it could be done directly by those more evolved external units. From a practical point of view, the sensors should just broadcast their sensed values to these units, possibly together with some identification numbers.

Usual arguments in favor of such decentralized computations are the repartition of the computational load, and the improved system robustness. Suppose now for example that there are 1000 sensors, and that 100-dimensional vectors are sensed. The external unit needs to perform about 100.000 operations and store at most one Megabyte, which is much less than what a usual cell-phone or MP3-player can do. Moreover, it is likely that the time needed for the whole transmission and computation process would be smaller than the time needed by a normal consensus algorithm to converge. Study of efficient low-energy transmissions in large scale networks are for example available in [83, 107]. Besides, suppose that the computation is done in a decentralized way, and that the communications topology is fixed and known by the external unit, which we suppose again to have some basic computational capabilities. After n observations, it is often possible for the latter to compute the exact average or even each sensor's initial value by solving a system of linear equations. Considering now the robustness argument, it is true that a decentralized average computation can still be performed when some agents are missing, while the centralized computation cannot if the entity performing the computation is not functioning. But again, the external controlling or recording unit are those using the computed value. So if they are not functioning, there seems to be little sense in still computing the average in the network. In other words, although the robustness of the computation is improved by the use of decentralized computations, the improvement of the total robustness is in many case more questionable.

There may however be situations involving totally decentralized sensors and controllers in which such decentralized computation is the best option for robustness and load balance purpose. For these, there should be objective measure of robustness and efficiency showing that the performances are better than with a global controller. As as semi-fictional example, think for example to a swarm of nano-robots, some of them having sensors which might be not really accurate, and who all need to know some global value. Since their number and position may vary with time, and since communication with devices different from nano-robots may be costly, it could be relevant to perform some computations in a totally decentralized way. The decentralized average computation can also be relevant if the sensed value needs to be used by many other agents moving among the sensors, or by one agent whose presence time is small as compared to the time needed for all agents to communicate their information. As an example of the first situation, think for example to the satellite of the GPS system who have to agree on a common time and to broadcast it to the GPS devices.

Besides, when the information is used by an external unit, an intermediate an maybe more efficient way for the sensors to transmit the values and compute their average could be to pre-compute some averages during the transmission process. When a sensor receives values from different other sensors, it computes and broadcasts the averages of these values to some next agent, eventually reaching the external unit. This process would present some similarities with the work of Delvenne et al. [36] on finite time consensus.

Our last point concerns the assumptions on agent capabilities. In many of the designed algorithms the agents store one value that they update by communicating with one or several other agents. Much more efficient algorithms could however be designed if the agents were able to keep more values in memory, such as previous time values, identifying number of some agents, values of some neighbors, etc. The efficiency could also be improved by allowing more elaborate mathematical operations and communication of more values.

If the values are digitally stored numbers that the agent update after receiving messages from other agents, the tacit assumption that only one number can be stored does not appear to be realistic. It is indeed notorious that communication of information, whether wired or wireless, requires non negligible computational capabilities. There is thus no appearing reason for which the agent could not have some additional memory registers. Moreover, the cost of a communication is much higher than the cost involved by a mathematical operation. It could be much more efficient to dedicate a small additional amount memory and computational capabilities to the agent if it allows dramatic reduction in the number of communications.

The single-value assumption makes however much sense if the value and its storage are materialized by the position of the agent. But an efficient algorithm designed for practical applications should then also take into account some inherent constraints. Position modifications are for example not instantaneous, and they may take an amount of time and energy proportional to their amplitude. Collisions also need to be avoided. Finally, when algorithms for moving agents use a fixed communication topology, agents should be able to identify each other and to communicate even in the possible presence of obstacles. This might pose problems, especially if the communication take for the form of distance measurements. Moreover, the communication topology would also need to be pre-specified.

To summarize the discussion above, I believe that both multi-agent models and algorithms lead to rich and interesting mathematical problems, from the study of which much can be learned. They also open huge perspectives of applications for the future, but many of those are not ready for immediate use. This implies by no mean that research should not be done on multi-agent systems. On the contrary, the potential future applications and the mathematical interest make them very relevant to study. But I believe that the research in this domain should thus be focussed on better understanding the behavior of those systems, and on building tools that could potentially be used in the future. In this approach, exploratory studies bringing new insight on some phenomena, and theoretical results on the behavior of classes of systems or on the conditions under which a method provides the desired outcome may be relevant, even when they cannot be immediately applied to practical situations. Whenever possible, such results should characterize classes of systems as large as possible, and rely on as few hypotheses as possible. In other words, when no direct application is in view, new studies should not be dedicated to some particular case of system but be as general as possible, unless the particular case in question brings some new insight or present some unexplained features, susceptible to open new perspectives. Particular practical applications would of course be very relevant, provided that they are realistic applications solving a problem that could not be solved as efficiently otherwise.

Bibliography

[1] D. Aeyels and F. De Smet. A mathematical model for the dynamics of clustering. *Preprint*, 2007.

[2] D. Aeyels and F. De Smet. A model for the dynamical process of cluster formation. In *7th IFAC Symposium on Nonlinear Control Systems (NOLCOS'2007)*, pages 260–264, Pretoria, South Africa, August 2007.

[3] B.D.O. Anderson, C. Yu, S. Dasgupta, and A.S. Morse. Control of a three-coleader formation in the plane. *Systems & Control Letters*, 56:573–578, 2007.

[4] H. Ando, Y. Oasa, I. Suzuku, and M. Yamashita. Distributed memoryless point convergence algorithm for mobile robots with limited visibility. *IEEE Transactions on Robotics and Automation*, 15:818–828, October 1999.

[5] D. Angeli and P.-A. Bliman. Extension of a result by moreau on stability of leaderless multi-agent systems. In *Proceedings of the 44th IEEE Conference on Decision and Control (CDC'2005)*, Seville, Spain, December 2005.

[6] L. Azimow and B. Roth. The rigidity of graphs. *Transactions of the American Mathematical Society*, 245:279–289, November 1978.

[7] L. Azimow and B. Roth. The rigidity of graphs II. *Journal of Mathematical Analysis and Applications*, 68:171–190, March 1979.

[8] J. Baillieul. *The Geometry of Sensor Information Utilization in Nonlinear Feedback Control of Vehicle Formations*, volume 309, pages 1–24. Springer, 2005.

[9] J. Baillieul. Remarks on a simple control law for point robot formations with exponential complexity. In *Proceedings of the 45th IEEE Confenerence on Decision and Control (CDC'2006)*, pages 3357–3362, San Diego (CA), USA, December 2006.

[10] J. Baillieul and McCoy L. The combinatorial graph theory of structured formations. In *Proceedings of the 46th IEEE Conference on Decision and Control (CDC'2007)*, New-Orleans (LA), USA, December 2007.

[11] J. Baillieul and A. Suri. Information patterns and hedging brockett's theorem in controlling vehicle formations. In *Proceedings of the 42nd IEEE Conference on Decision and Control (CDC'2003)*, volume 1, pages 556–563, Hawaii (HA), USA, December 2003.

[12] J. Bang-Jensen and T. Jordán. On persistent directed graphs. *To appear in Networks*, 2007.

[13] E. Ben-Naim, P.L. Krapivsky, and S. Redner. Bifurcations and patterns in compromise processes. *Physica D*, 183(3):190–204, 2003.

[14] E. Ben-Naim, P.L. Krapivsky, F. Vasquez, and S. Redner. Unity and discord in opinion dynamics. *Physica A*, 330:99–106, 2003.

[15] V. D. Blondel, J.-L. Guillaume, J. M. Hendrickx, C. de Kerchove, and R. Lambiotte. Local leaders in random networks. *To appear in Physical Review E, arXiv:0707.4064v1 [physics.soc-ph]*, 2008.

[16] V. D. Blondel, J.-L. Guillaume, J. M. Hendrickx, and R. M. Jungers. Distance distribution in random graphs and application to complex networks exploration. *Physical Review E, arXiv:0706.3322v2 [cond-mat.stat-mech]*, 76:066101, 2007.

[17] V. D. Blondel, J. M. Hendrickx, and R. M. Jungers. Linear time algorithms for Clobber. *Submitted, arXiv:cs/0703054v1 [cs.GT]*, 2007.

[18] V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *Proceedings of the 44th IEEE Conference on Decision and Control (CDC'2005)*, pages 2996–3000, Seville, Spain, December 2005.

[19] V.D. Blondel, J.M. Hendrickx, and J.N. Tsitsiklis. On the 2r conjecture for multi-agent systems. In *Proceedings of the European control conference 2007 (ECC'2007)*, pages 874–881, Kos, Greece, July 2007.

[20] V.D. Blondel, J.M. Hendrickx, and J.N. Tsitsiklis. Opinion dynamics systems with position dependent communication topologies. *Preprint*, 2007.

[21] J.M. Borky. Payload technologies and applications for uninhabited air vehicles (UAVs). In *Proceedings of the IEEE Aerospace Conference*, volume 3, pages 267–283, Aspen (CO), USA, February 1997.

[22] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: Design, analysis and applications. In *Proceedings of the 24th IEEE conference on Computer Communications (Infocomm'2005)*, volume 3, pages 1653–1664, Miami (FL), USA, March 2005.

[23] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52:2508–2530, June 2006.

[24] C.R. Calladine. Buckminster Fuller's "tensegrity" structures and clerck maxwell's rules for the construction of stiff frames. *International Journal of Solids Structures*, 13:161–172.

[25] M. Cao, A. S. Morse, and B. D. O. Anderson. Reaching a consensus in a dynamically changing environment: A graphical approach. *To appear in SIAM Journal on Control and Optimization*, 2007.

[26] M. Cao, A. S. Morse, and B. D. O. Anderson. Reaching a consensus in a dynamically changing environment: Convergence rates, measurement delays and asynchronous events. *To appear in SIAM Journal on Control and Optimization*, 2007.

[27] M. Cieliebak and P. Flocchini. Gathering autonomous mobile robots. In *Proceedings of the 9th International Colloquium on Structural Information and Communication Complexity (SIROCCO'2002)*, pages 57–72, Andros, Greece, June 2002.

[28] M. Cieliebak, P. Flocchini, G. Prencipe, and N. Santoro. Solving the robots gathering problem. In *Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP'2003)*, pages 1181–1196, June 2003.

[29] R. Connelly and W. Whiteley. Second-order rigidity and prestress stability for tensegrity frameworks. *SIAM Journal on Discrete Mathematics*, 9:453–491, August 1996.

[30] H.R. Cor. On the synthesis and analysis of simply-stiff frameworks. *Proceedings of the London Mathematical Society*, 40:203–216, 1936.

[31] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. Introduction to algorithms. pages 549–552. MIT Press, second edition, 2001.

[32] J. Cortes, S. Martinez, and F. Bullo. Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions. *IEEE Transactions on Automatic Control*, 51(8):1289–1298, 2006.

[33] M.H. De Groot. Reaching a consensus. *Journal of the American Statistical Association*, 69:118–121, 1974.

[34] F. De Smedt. *Clustering Behavior in Systems of Coupled Oscillators and Networks of Mutually Attracting Agents*. PhD thesis, Ghent University, http://users.ugent.be/˜fidesmet/thesis/thesis.pdf, 2007.

[35] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3:87–98, 2000.

[36] J.-C. Delvenne, R. Carli, and S. Zampieri. Optimal strategies in the average consensus problem. In *Proceedings of the 46st IEEE Conference on Decision and Control (CDC'2007), arXiv:0708.3220v1 [cs.MA]*, New Orleans (LA), USA, 2007.

[37] J. Desai, V. Kumar, and J. Ostrowski. Controlling formations of multiple mobile robots. In *Proceedings of 1998 IEEE International Conference on Robotics and Automation*, pages 2864 – 2869, Leuven, Belgium, May 1998.

[38] J. Desai, V. Kumar, and J. Ostrowski. Control of changes in formation for a team of mobile robots. In *Proceedings of 1999 IEEE International Conference on Robotics and Automation*, pages 1556–1561, Detroit (MI), USA, May 1999.

[39] J. Desai, J. Ostrowski, and V. Kumar. Modeling and control of formations of nonholonomic mobile robots. *IEEE Transactions on Robotics and Automation*, 17:905–908, 2001.

[40] K. Dogancay. Optimal receiver trajectories for scan-based radar localization. In *Proceedings Information, Decision and Control Conference 2007*, pages 88–93, Adelaide (SA), Australia, February 2007.

[41] T. Eren, B.D.O. Anderson, W. Whiteley, A.S. Morse, and P.N. Belhumeur. Operations on rigid formations of autonomous agents. *Communications in Informations and Systems*, 3(4):223–258, 2004.

[42] T. Eren, P.N. Belhumeur, and A.S. Morse. Closing ranks in vehicle formations based on rigidity. In *Proceedings of the 41st IEEE Conference on Decision and Control (CDC'2002)*, Las Vegas (NV), USA, December 2002.

[43] T. Eren, P.N. Belhumeur, W. Whiteley, B.D.O. Anderson, and A.S. Morse. A framework for maintaining formations based on rigidity. In *Proceedings of the 15th IFAC World Congress*, Barcelona, Spain, July 2002.

[44] T. Eren, W. Whiteley, B.D.O. Anderson, A.S. Morse, and P.N. Belhumeur. Information structures to secure control of rigid formations with leader-follower structure. In *Proceedings of the American Control Conference (ACC'2005)*, pages 2966–2971, Portland (OR), USA, June 2005.

[45] T. Eren, W. Whiteley, A.S. Morse, P.N. Belhumeur, and B.D.O. Anderson. Sensor and network topologies of formations with direction, bearing and angle information between agents. In *Proceedings of the 42nd IEEE Conference on Decision and Control (CDC'2003)*, pages 3064–3069, Maui (HA), USA, December 2003.

[46] T. Eren, W. Whiteley, A.S. Morse, P.N. Belhumeur, and B.D.O. Anderson. Information structures to secure control of globally rigid formations. In *Proceedings of the American Control Conference (ACC'2004)*, pages 4945–4950, Boston (MA), USA, July 2004.

[47] L. Euler. *Opera Postuma*, volume 1, pages 494–496. 1862.

[48] H.R Everett. *Sensors for Mobile Robots: Theory and Application*. A.K. Peters, 1995.

[49] F. Fagnani, K. H. Johansson, A. Speranzon, and S. Zampieri. On multi-vehicle rendezvous under quantized communication. In *Proceedings of the 16th International Symposium on Mathematical Theory of Networks and Systems (MTNS'2004)*, Leuven, Belgium, July 2004.

[50] Jia Fang. Personnal communications. 2006.

[51] B. Fidan and B.D.O. Anderson. Switching control for robust autonomous robot and vehicle platoon formation maintenance. In *Proceedings of the 15th Mediterranean Conference on Control and Automation*, pages T33–002, Athens, Greece, June 2007.

[52] B. Fidan, C. Yu, and B.D.O. Anderson. Acquiring and maintaining persistence of autonomous multi-vehicle formations. *IET Control Theory and Applications*, 1(2):452–460, March 2007.

[53] S. Fortunato, V. Latora, A. Pluchino, and Rapisarda R. Vector opinion dynamics in a bounded confidence consenus model. *International Journal of Modern Physics C*, 16:1535–1551, October 2005.

[54] R.B. Fuller. *Synergetics: Explorations in the geometry of thinking*. Macmillan, New York (NY) ,USA, 1975.

[55] I. Gilboa and A. Matsui. Social stability and equilibrium. *Econometrica*, 59:859–867, May 1991.

[56] H. Gluck. Almost all simply connected closed surfaces are rigid. In *Springer Verlag Lecture Notes No 438, Geometric Topology*, pages 225–239. Berlin, 1975.

[57] J. Graver, B. Servatius, and H. Servatius. Combinatorial rigidity. In *Graduate Studies in Mathematics, volume 2*. American Mathematical Society, Providence (RI), USA, 1993.

[58] J.-L. Guillaume and M. Latapy. Complex network metrology. *Complex Systems*, 16:83–94, 2005.

[59] R. Hegselmann. Opinion dynamics: Insights by radically simplifying models. In D. Gillies, editor, *Laws and Models in Science*, pages 1–29, London, 2004.

[60] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulations. *Journal of Artificial Societies and Social Simulation*, 5(3):http://jasss.soc.surrey.ac.uk/9/1/8.html, 2002.

[61] B. Hendrickson. Conditions for unique graph realizations. *Siam Journal on Computing*, 1(21):65–84, 1992.

[62] J.M. Hendrickx. Order preservation in a generalized version of Krause's opinion dynamics model. *arXiv:0708.4343v1 [physics.soc-ph]*, 2007.

[63] J.M. Hendrickx, B.D.O. Anderson, and V.D. Blondel. Rigidity and persistence of directed graphs. In *Proceedings of the 44th IEEE Conference on Decision and Control (CDC'2005)*, pages 2176–2181, Seville, Spain, December 2005.

[64] J.M. Hendrickx, B.D.O. Anderson, J.-C. Delvenne, and V.D. Blondel. Directed graphs for the analysis of rigidity and persistence in autonomous agents systems. *International Journal of Robust and Nonlinear Control*, 17:960–981, 2007.

[65] J.M. Hendrickx and V.D. Blondel. Convergence of different linear and non-linear Vicsek models. In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS'2006)*, pages 1229–1240, Kyoto, Japan, July 2006.

[66] J.M. Hendrickx, B. Fidan, C. Yu, B.D.O. Anderson, and V.D. Blondel. Primitive operations for the construction and reorganization of minimally persistent formations. *arXiv:cs/0609041v1 [cs.MA]*.

[67] J.M. Hendrickx, B. Fidan, C. Yu, B.D.O. Anderson, and V.D. Blondel. Rigidity and persistence of three and higher dimensional formations. In *Proceedings of the First International Workshop on Multi-Agent Robotic Systems (MARS'2005)*, pages 39–46, Barcelona, Spain, September 2005.

[68] J.M. Hendrickx, B. Fidan, C. Yu, B.D.O. Anderson, and V.D. Blondel. Elementary operations for the reorganization of minimally persistent formations. In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS'2006)*, pages 859–873, Kyoto, Japan, July 2006.

[69] J.M. Hendrickx, B. Fidan, C. Yu, B.D.O. Anderson, and V.D. Blondel. Formation reorganization by primitive operations on directed graphs. *to appear in IEEE Transactions on Automatic Control*, 2008.

[70] J.M. Hendrickx, C. Yu, B. Fidan, and B.D.O. Anderson. Rigidity and persistence of meta-formations. In *Proceedings of the 45th IEEE Conference on Decision and Control (CDC'2006)*, pages 5980–5985, San Diego, (CA) USA, December 2006.

[71] J.M. Hendrickx, C. Yu, B. Fidan, and B.D.O. Anderson. Rigidity and persistence for ensuring shape maintenance in multiagent meta-formations. *to appear in Asian Journal of control's special issue on Collective Behavior and Control of Multi-Agent Systems*, 2008.

[72] L. Henneberg. Die graphische Statik der starren Systeme. Leipzig, 1911.

[73] B. Jackson and T. Jordán. Connected rigidity matroids and unique realization of graphs. *Journal of Combinatorial Theory, Series B*, 94:1–29, 2005.

[74] D.J. Jacobs and B. Hendrickson. An algorithm for two-dimensional rigidity percolation: the pebble game. *Journal of Computational Physics*, 137(2):346–365, 1997.

[75] 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.

[76] A. Kashyap, T. Basar, and R. Srikant. Consensus with quantized information updates. In *Proceedings of the 45th IEEE Confenerence on Decision and Control (CDC'2006)*, pages 2728–2733, San Diego (CA), USA, December 2006.

[77] U. Krause. Soziale Dynamiken mit vielen Interakteuren. Eine Problemskizze. In *Modellierung und Simulation von Dynamiken mit vielen interagierenden Akteuren*, pages 37–51. 1997.

[78] U. Krause. A discrete nonlinear and non-autonomous model of consensus formation. *Communications in Difference Equations*, pages 227–236, 2000.

[79] L Krick. Application of graph rigidity in formation control of multi-robot networks. Master's thesis, Graduate Department of Electrical and Computer Engineering University of Toronto, 2007.

[80] S.O. Krumke and H. Noltemeier. Graphentheorische Konzepte und Algorithmen. B.G. Teubner, 2005.

[81] Y. Kuramoto. Cooperative dynamics of oscillator community. *Supplement of the Progress of Theoretical Physics*, (79):224–240, 1984.

[82] H.J. Kushner. *Stochastic Stability and Control*. Academic Press, New York (NY), USA, 1967.

[83] W. Lai and I. Ch. Paschalidis. Optimally balancing energy consumption versus latency in sensor network routing. *to appear in ACM Transactions on Sensor Networks*, 2007.

[84] G. Laman. On graphs and rigidity of plane skeletal structures. *Journal of Engineering Mathematics*, 4:331–340, 1970.

[85] M.A. Lewis and K.H. Tan. High precision formation control of mobile robots using virtual structures. *Autonomous Robots*, 4(4):387–403, 1997.

[86] S. Li and H. Wang. Multi-agent coordination using nearest neighbor rules: a revisit to Vicsek model. *arXiv:cs/0407021v2 [cs.MA]*, 2004.

[87] J. Lin, A.S. Morse, and B.D.O. Anderson. The multi-agent rendezvous problem part i, the synchronous case. *in press for SIAM Journal on Control and Optimization*.

[88] J. Lin, A.S. Morse, and B.D.O. Anderson. The multi-agent rendezvous problem. In *Proceedings of the 42th IEEE Conference on Decision and Control (CDC'2003)*, pages 1508– 1513, Hawaii (HA), USA, December 2003.

[89] J. Lin, A.S. Morse, and B.D.O. Anderson. The multi-agent rendezvous problem: An extended summary. In *Lecture Notes in Control and Information Sciences*, number 309, pages 257–289. 2005.

[90] J. Lorenz. A stabilization theorem for continuous opinion dynamics. *Physica A*, 355(1):217–223, 2005.

[91] J. Lorenz. Consensus strikes back in the Hegselmann-Krause model of continuous opinion dynamics under bounded confidence. *Journal of Artificial Societies and Social Simulation*, 9(1):http://jasss.soc.surrey.ac.uk/9/1/8.html, 2006.

[92] J. Lorenz. Continuous opinion dynamics under bounded confidence: A survey. *Preprint, arXiv:0707.1762v1 [physics.soc-ph]*, July 2007.

[93] J. Lorenz. *Repeated Averaging and Bounded Confidence - Modeling, Analysis and Simulation of Continuous Opinion Dynamics*. PhD thesis, Universität Bremen, http://nbn-resolving.de/urn:nbn:de:gbv:46-diss000106688, 2007.

[94] J. Lorenz and "undergraduate students of psychology in the third semester". Meinungsdynamik-Experiment, ss 2006. *available on request from post@janlo.de*, 2006.

[95] L. Lovasz and Y. Yemini. On generic rigidity in the plane. *SIAM Journal on Algebraic and Discrete Methods*, 3.

[96] A. Mantler and J. Snoeyink. Banana spiders: A study of connectivity in 3d combinatorial rigidity. In *Proceedings of the 16th Canadian Conference on Computational Geometry*, pages 44–47, August 2004.

[97] S. Martinez and F. Bullo. Optimal sensor placement and motion coordination for target tracking. *Automatica*, 42(3):661–668, 2006.

[98] S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, London, UK, 1993.

[99] L. Moreau. Stability of continuous-time distributed consensus algorithms. In *Proceedings of the 43st IEEE Conference on Decision and Control (CDC'2004)*, volume 4, pages 3998–4003, Paradise Island, Bahamas, December 2004.

[100] L. Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005.

[101] C. Moukarzel. An efficient algorithm for testing the generic rigidity of graphs in the plane. *Journal of Physics A*, 29(24):8079–8098, 1996.

[102] K. Nagaich. Report on the work pursued during the intership at ucl.

[103] R. Olfati-Saber, J.A. Fax, and R.M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, January 2007.

[104] R. Olfati-Saber and R.M Murray. Distributed structural stabilization and tracking for formations of dynamic multi-agents. In *Proceedings of the 41st Conference on Decision and Control (CDC'2002)*, volume 1, pages 209–215, Las Vegas (NV), USA, December 2002.

[105] R. Olfati-Saber and R.M Murray. Graph rigidity and distributed formation stabilization of multi-vehicle systems. In *Proceedings of the 41st Conference on Decision and Control (CDC'2002)*, volume 3, pages 2965–2971, Las Vegas (NV), USA, December 2002.

[106] M. Pachter and J. Hebert. *Cooperative aircraft control for minimum radar exposure*, pages 199–211. Kluwer Academic, 2002.

[107] I. Ch. Paschalidis, W. Lai, and D. Starobinski. Asymptotically optimal transmission policies for large-scale low-power wireless sensor networks. *IEEE/ACM Transactions on Networking*, 15(1):105–118, 2007.

[108] H. Peyton Young. *Individual Strategy and Social Structure*. Princeton Paperbacks, Princeton, Providence (NJ), USA, 1998.

[109] A. Recksi. A network theory approach to the rigidity of skeletal structures. i: Modelling and interconnection. *Discrete applied mathematics*, 7(3):313–324, 1984.

[110] A. Recksi. A network theory approach to the rigidity of skeletal structures. ii: Laman's theorem and topological formulae. *Discrete applied mathematics*, 8(1):63–68, 1984.

[111] W. Ren and R.W. Beard. Consensus seeking in multiagent systems under dynamically changing interaction. *IEEE Transactions on Automatic Control*, 50:655–661, May 2005.

[112] W. Ren, R.W. Beard, and E.M. Atkins. *IEEE Control and Systems Magazine*, 27(2):71–82, April 2007.

[113] J.A.A. Rogge. *Dynamic Behavior of Oscillator Networks*. PhD thesis, Ghent University, http://erente.googlepages.com/ErenPhDThesis.pdf, 2006.

[114] B. Roth and W. Whiteley. Tensegrity frameworks. *Transactions of the American Mathematical Society*, 265(2):419–446, june 1981.

[115] H. Sakaguchi and Y. Kuramoto. A soluble active rotator model showing phase transitions via mutual entrainment. *Progress of Theoretical Physics*, 3(76):576–581, September 1986.

[116] A. Sarlette, L. Scardovi, and R. Sepulchre. Synchronization and balancing on the n-torus. *Systems and Control Letters*, 56(5):335–341, 2007.

[117] A. Sarlette and R. Sepulchre. Consensus optimization on manifolds. *Preprint*, 2007.

[118] A. V. Savkin. Flocking motion of groups of autonomous mobile robots. In *Proceedings of the European control conference 2007 (ECC'2007)*, pages 896–899, Kos, Greece, July 2007.

[119] B. Smith, M. Egerstedt, and A. Howard. Persistent formations for multi-agent networks under range constraints. In *Proceedings of the first International Conference on Robot Communication and Coordination*, Athens, Greece, October 2007.

[120] I. Suzuki and M. Yamashita. Distributed anonymous mobile robots: Formation of geometric patterns. *Siam Journal of Computing*, 28(4):1347–1363, 1999.

[121] H.G. Tanner, G.J. Pappas, and V. Kumar. Leader-to-formation stability. *IEEE Transactions on Robotics and Automation*, 4(3):433–455, 2004.

[122] A. Tarski. *A Decision method for elementary algebra and geometry*. University of California Press, 1951.

[123] T.S. Tay. Rigidity of multigraphs. I. Linking rigid bodies in n-space. *J. Combin. Theory Ser. B*, 36(1):95–112, 1984.

[124] T.S. Tay and W. Whiteley. Generating isostatic frameworks. *Structural Topology*, 11:21–69, 1985.

[125] J.N. Tsitsiklis. *Problems in decentralized decision making and computation*. PhD thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, http://web.mit.edu/jnt/www/PhD-84-jnt.pdf, 1984.

[126] D. Urbig. Attitude Dynamics With Limited Verbalisation Capabilities. *Journal of Artificial Societies and Social Simulation*, 6(1):http://jasss.soc.surrey.ac.uk/6/1/2.html, 2003.

[127] T. Vicsek, A. Czirok, I. Ben Jacob, I. Cohen, and O. Schochet. Novel type of phase transitions in a system of self-driven particles. *Physical Review Letters*, 75:1226–1229, 1995.

[128] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal. Interacting agents and continuous opinions dynamics. Technical report, arXiv:cond-mat/0111494v1 [cond-mat.dis-nn], 2001.

[129] W. Whiteley. Infinitesimally rigid polyhedra I. Statics of frameworks. *Transactions of the American Mathematical Society*, 285(2):431–465, 1984.

[130] W. Witheley. Rigidity and scene analysis. In J.E. Goodman and J. J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 60, pages 1327–1354. Hall/CRC Press, Boca Raton (FL), USA, 2004.

[131] J. Wolfowitz. Products of indecomposable, aperiodic, stochastic matrices. *Proceedings of the American Mathematical Society*, 14:733–737, 1963.

[132] C. Yu, B.D.O. Anderson, S. Dasgupta, and B. Fidan. Control of minimally persistent formations in the place. *Preprint*.

[133] C. Yu, J.M. Hendrickx, B. Fidan, and B.D.O. Anderson. Structural persistence of three dimensional autonomous formations. In *Proceedings of the First International Workshop on Multi-Agent Robotic Systems (MARS'2005)*, pages 47–55, Barcelona, Spain, September 2005.

[134] C. Yu, J.M. Hendrickx, B. Fidan, B.D.O. Anderson, and V.D. Blondel. Three and higher dimensional autonomous formations: Rigidity, persistence and structural persistence. *Automatica*, 43:387–402, March 2007.

List of Figures

Appendix A

Definitions related to Part I

Rigidity and Persistence of graph representations

A **graph representation** is a function $p : V \rightarrow \Re^D : i \rightarrow p_i$. The **position** of a vertex i is its value by p. The distance between two representations is defined by $d(p, q) = \max_{i \in V} ||p_i - q_i||$.

A **distance set** associated to a graph is a set of desired distances $d_{ij} = d_{ji}$ for each edge (i, j) of the graph.

A distance set is **realization** of a distance associated to a graph is a representation p of this graph such that $||p_i - p_j|| = d_{ij}$ holds for each edge (i, j). A distance set is **realizable** if it admits a realization. Moreover, each graph representation induces a distance set of which it is a realization.

A **Euclidean transformation** is a bijection $E : \Re^D \rightarrow \Re^D : x \rightarrow E(x)$ such that $||x - y|| = ||E(x) - E(y)||$ holds for any $x, y \in \Re^D$. It is a combination of rotations, translations and reflections.

Two representations are **congruent** if one is the image of the other by a Euclidean transformation.

A representation is **an equilibrium representation** for a distance set $\{d_{ij}\}$ associated to a graph if there is no vertex i and point $p^* \in \Re^D$ such that the following strict inclusion holds

$$\{(i, j) \in E : ||p_i - p_j|| = d_{ij}\} \subset \{(i, j) \in E : ||p^* - p_j|| = d_{ij}\}.$$

In other word, it is impossible to increase the set of satisfied constraints corresponding to edges leaving an agent by modifying the position of this vertex, all other positions remaining fixed.

A **representation** p **of a graph** G **is constraint consistent** if there is a neighborhood of p in which every representation p' at equilibrium for the distance set induced by p and G is a realization of the same distance set.

A **representation** p **of a graph** G **is rigid** if there is a neighborhood of p in which all realizations p' of the distance set induced by p are congruent to p.

A **representation** p **of a graph** G **is persistent** if there is a neighborhood of p in which every representation p' at equilibrium for the distance set induced by p and G is congruent to p. A representation is persistent if and only if it is rigid and constraint consistent.

Infinitesimal notions

An **infinitesimal displacement** δp_i of a vertex is a vector or $\Re^D$ for which the approximation $\delta p_i^T \delta p_i = 0$ is made. An infinitesimal displacement of a graph or of a representation is a vector $\delta p \in \Re^{Dn}$ obtained by juxtaposing the infinitesimal displacements of all vertices: $\delta p^T = \left(\delta p_1^T, \ldots, \delta p_n^T\right)$.

A **partial infinitesimal displacement** is a subset of vertices $V_c \subseteq V$ together with an infinitesimal displacement δp_i for each $i \in V_c$. We denote it by δp_{V_c}. For a partial displacement δp_{V_c}, any displacement δp whose restriction to V_c is δp_{V_c} is called a **completion** of δp_{V_c}.

The **rigidity matrix** $R_{G,p} \in \Re^{|E| \times Dn}$ associated to G and p is obtained by associating one line to each edge, and D columns to each edge. . Vertices of the graphs correspond to D columns of the matrix, and each line of the latter corresponds to an edge (i, j) in G and is equal to

$$\begin{pmatrix} \ldots & 0 & (p_i - p_j)^T & 0 & \ldots & 0 & (p_j - p_i)^T & 0 & \ldots \end{pmatrix},$$

the non-zero column being the $(i-1)D+1^{st}$ to the iD^{th} and the $(j-1)D+1^{st}$ to the jD^{th}. Multiplying δp by the line corresponding to the edge (i, j) gives $(p_i - p_j)^T (\delta p_i - \delta p_j)$.

An **infinitesimal displacement** δp **is admissible** for a representation p of a graph G if $\delta p \in \mathrm{Ker} R_{G,p}$, that is, if $R_{G,p}\delta p = 0$.

An **infinitesimal displacement** δp **is Euclidean** if it is a combination of infinitesimal translations and rotations. More formally, an infinitesimal displacement δp is Euclidean if there exists a time-continuous length-preserving transformation $E : \Re^D \times \Re^+ \to \Re^D : (x,t) \to E(x,t)$ such that $\delta p_i = K \frac{dE(x,t)}{dt}|_{p_i,0}$ holds with the same K for all $i \in V$. We denote by Eu_p the set of Euclidean infinitesimal displacements for a representation p. If the number of vertices is larger than D, Eu_p has generically a dimension $f_D = \frac{1}{2}D(D+1)$. For every representation p of any graph G, there holds $Eu_p \subseteq R_{G,p}$.

An **equilibrium infinitesimal displacement** δp is an infinitesimal displacement such that for each i, δp_i satisfies a maximal set of equations of

$$\begin{pmatrix} (p_i - p_{j_1})^T \\ (p_i - p_{j_2})^T \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \end{pmatrix} \delta p_i = \begin{pmatrix} (p_i - p_{j_1})^T \delta p_{j_1} \\ (p_i - p_{j_2})^T \delta p_{j_2} \\ \vdots \\ (p_i - p_{j_{d_i^+}})^T \delta p_{j_{d_i^+}} \end{pmatrix},$$

considering p and all δp_{j_k} as fixed, where $j_1, \dots, j_{d_i^+}$ are the neighbors of i. Note that each equation of this system is equivalent to $(p_i - p_{j_k})^T (\delta p_i - \delta p_{j_k}) = 0$. We denote by $\mathrm{Equil}_{G,p}$ the set of all equilibrium infinitesimal displacements associated to a representation p of a graph G. For every representation p of any graph G, there holds $\mathrm{Ker} R_{G,p} \subseteq \mathrm{Equil}_{G,p}$.

A **representation p of a graph G is infinitesimally constraint consistent** if all its equilibrium infinitesimal displacements are admissible, i.e., if $\mathrm{Equil}_{G,p} \subseteq \mathrm{Ker} R_{G,p}$.

A **representation p of a graph G is infinitesimally rigid** if all its admissible infinitesimal displacements are Euclidean, i.e., if $\mathrm{Ker} R_{G,p} \subseteq Eu_p$.

A **representation p of a graph G is infinitesimally persistent** if all its equilibrium infinitesimal displacements are Euclidean, i.e., if $\mathrm{Equil}_{G,p} \subseteq Eu_p$. A representation is infinitesimally persistent if and only if it is infinitesimally rigid and infinitesimally constraint consistent.

A **representation p of a graph G is structurally persistent (respectively constraint consistent)** if it is persistent (respectively constraint consistent) and if every partial equilibrium can be completed to obtain an equilibrium displacement, that is an displacement for which every vertex is at equilibrium.

Generic notions

Let P be a property defined for graph representations. A graph is **generically P** if the set of its representations not having the property P has zero measure. A graph is **generically not P** if the set of its representations having the property P has zero measure. The property is **a generic property** if every graph

is either generically P or generically not P.
For the sake of conciseness, we omit the word "generically" in the sequel, unless when the context could allow ambiguities.

For a representation p of a graph G, we say that a set of edges is **independent** if the corresponding lines in $R_{G,p}$ are linearly independent. It can be proved [6,7] that the independence of edges is a generic notion. We call **generic representations** those representations for which every generically independent set of edges is independent.

A representation p of a directed graph G is **non-degenerate** if for any i and any subset $\{j_1, j_2, \dots, j_{n'_i}\}$ of at most D of the vertices to which it is connected by directed edges, the collection of vectors $\left\{(p_{j_1} - p_i), (p_{j_2} - p_i), \dots, (p_{j_{n'_i}} - p_i)\right\}$ spans a n'_i-dimensional space.

Rigidity and persistence of graphs

A graph is (generically) **rigid** if almost all its representations are rigid, or equivalently if almost all its representations are infinitesimally rigid.

A graph is (generically) **constraint consistent** if almost all its representations are constraint consistent, or equivalently if almost all its representations are infinitesimally constraint consistent.

A graph is (generically) **persistent** if almost all its representations are persistent, or equivalently if almost all its representations are infinitesimally persistent.

Minimal rigidity and persistence

A graph is **minimally rigid** if it is rigid and if the removal of any one or several of its edges leads to a loss of rigidity. Containing a minimally rigid subgraph on all its vertices is necessary and sufficient for a graph to be rigid.

A graph is **minimally persistent** if it is persistent and if the removal of any one or several of its edges leads to a loss of persistence. Every persistent graph contains a minimally persistent graph on all its vertices, but containing such a subgraph is not sufficient for persistence.

When a graph has more than D vertices, it is minimally rigid (respectively persistent) if and only if it is rigid (respectively persistent) and contains exactly $nD - \frac{1}{2}D(D+1)$ edges.

Subgraphs

A **closed subgraph** $G'(V', E')$ of a graph $G(V, E)$ is a subgraph for which $d^+_{i,G'} = d^+_{i,G}$ holds for all i. Every edge leaving a vertex of G' arrives thus in G' and belongs to E'.

A $\min D$**-subgraph** $G'(V', E')$ of a graph $G(V, E)$ is a subgraph for which $d^+_{i,G'} \geq \min\left(D, d^+_{i,G}\right)$ holds for all i. A vertex of a $\min D$-subgraph may thus be left by an edge of $E \setminus E'$ if it is left by at least D edges of E'.

A **strict** $\min D$**-subgraph** $G'(V', E')$ of a graph $G(V, E)$ is a subgraph for which $d^+_{i,G'} = \min\left(D, d^+_{i,G}\right)$ holds for all i. A vertex of a strict $\min D$-subgraph is thus left by exactly D edges of E' if its out-degree is larger than or equal to D. If its out-degree is smaller than D, all its outgoing edges belong to E'.

A **subgraph of** $G(V, E)$ **on all its vertices** is a subgraph $G'(V, E')$ of G sharing the same vertex set.

Miscellaneous

The **number of degrees of freedom** (in $\Re^D$) of a vertex in a directed graph is the generic dimension of the set of its possible equilibrium displacement (in $\Re^D$), considering the displacements of the other agents as fixed. It equals $D - \min\left(D, d^+_i\right)$. The sum over all vertices in a persistent graph of their number of degrees of freedom is at most $f_D = \frac{1}{2}D(D+1)$.

A **double edge** in a directed graph is a cycle of length 2. There is a double edge between i and j if $(i, j) \in E$ and $(j, i) \in E$.