



HAL
open science

Distributed and scalable methods for the analysis and control of networks of dynamical agents

Paolo Frasca

► **To cite this version:**

Paolo Frasca. Distributed and scalable methods for the analysis and control of networks of dynamical agents. Systems and Control [cs.SY]. Université de Grenoble, 2024. tel-04674170

HAL Id: tel-04674170

<https://theses.hal.science/tel-04674170v1>

Submitted on 21 Aug 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HABILITATION À DIRIGER DES RECHERCHES

Pour obtenir le diplôme de

L'UNIVERSITÉ GRENOBLE ALPES

Spécialité : **Automatique**

Arrêté ministériel :

Présentée par

Paolo Frasca

préparée au sein du **Laboratoire Grenoble Images Parole Signal Automatique (GIPSA-Lab)**

et de l'**École doctorale Électronique, Électrotechnique, Automatique, Traitement du Signal (EEATS)**

Distributed and Scalable Methods for the Analysis and Control of Networks of Dynamical Agents

Thèse soutenue publiquement le **12/07/2024**,
devant le jury composé de :

M. Carlos Canudas de Wit

Directeur de Recherche, CNRS, GIPSA-Lab, Grenoble, Président

M. Claudio Altafini

Professor, University of Linköping, Sweden, Rapporteur

M. Ming Cao

Professor, University of Groningen, Netherlands, Rapporteur

M. Constantin Morescu

Professeur, Université de Lorraine, France, Rapporteur

M. Mario di Bernardo

Professor, University of Naples, Examineur

Mme. Elena Panteley

Directrice de Recherche, CNRS, L2S, Paris-Saclay, Examinatrice



Acknowledgements

Writing this Habilitation memoir has naturally provided an opportunity not only to reflect on my past and future research but also to acknowledge the many people whose influences have shaped my work, scientific interests, and career—students, mentors, and collaborators. I feel privileged to have collaborated with and learned from such remarkable individuals. While I will refrain from attempting to make a full list here, I hope their positive influence resonates with the readers of this document, for whose flaws I bear sole responsibility.

On this occasion, I would especially like to acknowledge the roles of Francesca Ceragioli, Roberto Tempo, Federica Garin, and Carlos Canudas de Wit in my career following my PhD and in the steps that have led to this Habilitation. Francesca graciously supported me at a critical juncture in my career. Roberto was both a role model and a friend until his untimely passing; I wish to dedicate this document to his inspiring memory. Federica strongly encouraged me to apply for a CNRS position in Grenoble and has provided invaluable support since then. Carlos's continued support since 2016 has been crucial in my successful application for this Habilitation from the University of Grenoble.

I would also like to extend my gratitude to the esteemed colleagues who have kindly agreed to serve on my Habilitation committee, particularly the three referees.

Finally, my special thanks go to Christian Commault, who has kindly and effectively encouraged me over several years to pursue this Habilitation.

Contents

Acknowledgements	iii
1 Introduction	1
1.1 Motivations, themes, and methods of my research	2
1.2 Collaborations, community spirit, and interdisciplinarity	6
1.3 Mentoring students	7
1.4 The rest of this document	9
I Sampling from my research past	11
2 Discrete-time consensus algorithms	13
2.1 The mean square error of randomized consensus algorithms	14
3 Nonsmooth consensus-seeking systems	17
3.1 Basic consensus dynamics	17
3.2 Differential equations with discontinuous right-hand side	19
3.3 Bounded confidence	21
3.4 k -neighbor interactions	22
3.5 Quantization in social interactions	23
3.6 Quantization for the consensus problem	26
4 Distributed multi-agent and network systems	29
4.1 Coverage control by gossiping robots	30
4.2 Distributed estimation from relative measurements	31
4.3 Ergodic randomized network dynamics	32
5 Approximate methods for large-scale networks	37
5.1 Node aggregation methods	38
5.2 The continuation method	39
5.3 Graphons	40
5.4 Illustrative example: Epidemics on networks	41
6 Applications in engineering, biology, social sciences	45
6.1 Optimal deployment of sensor and robotic networks	45
6.2 Animal groups: global shape from local interactions	46
6.3 The origins of disagreement in opinion dynamics	47
6.4 Vehicle groups: communication and heterogeneity	49

II Scientific Perspective	53
7 Dynamics of large-scale socio-technical networks	55
7.1 The challenges of large uncertain networks	56
7.2 Open networks	57
7.3 The “cyber” loop	60
7.3.1 A toy model of news recommendations	62
8 Ongoing work: theory and applications	65
8.1 Graph limits: dynamics on graphs and on graphons	65
8.2 Online social media and their disorders	67
8.3 Information loops in mobility: the case of navigation apps	68
References	70
Appendix	86
A Curriculum Vitae	87
A.1 Positions and employment	87
A.2 Education	88
A.3 Awards	88
A.4 Invited faculty positions & academic visits	88
A.5 Supervised students and postdocs	89
A.6 Editorial and referee activities	92
A.7 Leadership roles, service to the community and organization activities	93
A.8 Funding and research grants	97
A.9 Outreach and non-academic collaborations	99
A.10 Teaching	99
A.11 Invited presentations	100
B Publications	105
B.1 Overview	105
B.2 Selected papers on key topics	106
B.3 Full list of publications, per type	107
B.3.1 Books	107
B.3.2 Edited books	108
B.3.3 Chapters in books	108
B.3.4 Journal papers	108
B.3.5 Conference proceedings papers	113
B.3.6 Technical reports	117
B.3.7 Preprints	117

Chapter 1

Introduction

This document has been written as a requirement to obtain the French Habilitation degree *Habilitation à Diriger des Recherches*. Therefore, it is meant to allow the habilitation committee to appreciate the quality and originality of my research, my scientific maturity, and my ability to independently supervise junior researchers. To this purpose, the document summarizes my career, my research after my PhD thesis and my activities as supervisor, in a critical perspective and with due reflection about their future.

Sketching this critical summary will be the main objective of this first chapter. Section 1.1 contains an overview of what I consider to be the main themes, original features, and overarching objectives of my research. Section 1.2 comments on my inclination to collaborate with colleagues, even across disciplinary boundary lines. Section 1.3 zooms in onto my fruitful experiences as advisors. This chapter will thus be rich in references to my own work and pointers to later sections of this document.

In the 15 years since defending my PhD, I have worked on a broad range of research questions and I have collaborated with more than a hundred researchers based in more than a dozen countries, including a fair number of students whose work I have advised. This work and these collaborations have led to many significant results, which have been published in journals and conference proceedings and have been presented at scientific audiences at conferences, workshops and schools. My publications include more than sixty journal papers published in more than twenty different international journals, which cater to a broad range of scientific communities, from my core domain of automatic control to network science, transportation research, robotics, signal processing, mathematical biology, sociology and media studies.

In view of the variety of this material, I will refrain from attempting to provide a comprehensive account. Instead, I will try to highlight the *fil rouge* that runs through it: the mathematical modeling of dynamics and interactions through notions from control systems and network theory. In doing so, I will have to make difficult choices about what to devote space to, and I will spare the readers almost all technical details. For some topics, but not for all, I will use mathematical expressions and language, but I have not striven to provide all the necessary background and I have not made notation consistent across different chapters: I hope the readers will be forgiving and, most of all, I am confident they will manage to make their way through the material.

In the rest of this chapter, my objective will be to provide a concise account of the




Figure 1.1: Representation of a network system as a graph with nodes (agents) and arcs (interactions). Each agent i has a state x_i that evolves dynamically per $\dot{x}_i = f_i(x)$, where x is the stack vector of all agents states and f_i only depends on the components of x that correspond to agents j that are neighbors of i .

main overarching objectives, driving questions, techniques, and recurring themes of my research. I will also aim to emphasize how advising junior researchers has been integral to my work.

Later, in the chapters that constitute Part I of this manuscript, I will provide a more detailed description of some selected lines of work: not all the research that is mentioned in this introduction will be detailed in these chapters. After this perspective about the past, I will present a research perspective in Part II. The document is completed by relevant annexes, such as a detailed curriculum vitae and a full list of publications.

1.1 Motivations, themes, and methods of my research

Since my doctoral studies, my research interests have been at the intersection of network science and control systems theory. I mostly worked on mathematical problems that feature multiple *dynamical systems that interact through a network*. The network is mathematically described by a graph, whose nodes represent the dynamical systems and whose edges represent the interactions between the systems. I often refer to the interacting dynamical systems as *agents* and to their whole interconnection as either a *multi-agent system* or a *network system*. A graphical representation with some recurring notation is given in Figure 1.1.

Depending on the context and the motivating application, these agents can represent various entities, such as mobile robots, sensors, autonomous vehicles, flocking animals, or individuals in a society. The states of the agents can represent their positions in space or other relevant variables, such as the opinions of individuals in social applications. The agents' states evolve according to differential equations or discrete-time processes. Interactions can be physical in nature or, more often, exchanges of information. The topology of the interaction graph can either be fixed or depend on time and on the agents' states: for instance, interactions may be constrained to take place between agents that are close to each other.

Within this framework, the main underlying question in my research has been to understand how multiple dynamical agents interact and produce a collective behavior. This

general question can be seen either as an analysis question, where one aims to understand the origins of an observed collective behavior or to deduce the consequences of an interaction rule, or as a design question, where one tries to engineer the right interaction rules to achieve a desired collective outcome.

Along the years, this general question has brought me to explore several **application domains**, some of which are more extensively described in Chapter 6:

1. **Sensor networks**, for which I have studied distributed algorithms for the estimation of global quantities from local measurements [234, 248].
2. **Robotic teams**, for which I have designed and implemented algorithms that ensure their optimal deployment and their cooperation with minimal, asynchronous, and unreliable communication [100, 88].
3. **Connected and autonomous vehicles**, for which I have designed control methods to ensure stable platooning behavior notwithstanding unreliable communication [5, 4], heterogeneities [170] and nonlinearities [140].
4. **Animal groups**, in which I have highlighted (through theoretical and empirical work) the role of the communication patterns in defining the shapes of their flocks [82, 14].
5. **Epidemics on networks**, for which I have studied stability and long-time behavior on large graphs [280, 93], designed strategies for effective cure allocation [184], and evaluated non-pharmaceutical interventions in a realistic simulation setting for schools facing COVID-19 outbreaks [189].
6. **Social networks**, regarding which I have studied original models of opinion formation [124, 62, 241] and social influence [274]. More recently, I have also experimentally explored the attention dynamics in social media and, more specifically, on YouTube [49].

In exploring these applications, I always strove to reach out and communicate my results to the relevant scientific communities. This effort can be seen in my publication record, which comprises a “core” of methodological publications in control systems journals, together with publications in the fields that are associated to the above applications: signal processing, robotics, transportation systems, biology, and social sciences.

These applications have motivated my interest in several **methodological questions**: I describe here **five major themes** that are weaved into most of my research.

Distributed control of networks. In a network system, a (possibly time-varying) graph defines which agents interact or communicate with each other. Therefore, the network graph encodes the most fundamental communication constraints. In an engineering perspective, we may take the set of agents and the graph as given, and design control algorithms to achieve a collective task. These algorithms require coordination, and therefore communication, between the agents. I say that an algorithm is *distributed* when it satisfies the communication constraint given by the network. A prototypical example of coordination problem is the *average consensus problem*, on which I have worked extensively and which will be formally presented in Chapter 2: in the average consensus problem, the agents have to coordinate to compute the global average of some locally available quantities. For dynamical networks with distributed interactions, the topology of the network determines the ability to achieve a coordination task and any measure of collective performance. For instance, the average consensus problem can be

solved so long as the network is connected, and the computation of the average will be faster if the graph is “more connected”, in a sense that can be made precise [105].

Communication constraints. In addition to being constrained by the graph, communication between agents in network systems can suffer from additional constraints. For instance, in engineering applications, the implementation of control procedures is affected by the need to communicate measurements and controls through digital channels. Communication channels are subject to limited precision (because of using finite numbers of bits), delays, noise and message losses. The impact of these issues on coordination algorithms has been one of my main research topics. One frequent issue is that communication can be unreliable because subject to unpredictable *losses* on the communication channel: I extensively considered this issue both in consensus problems [45, 104, 120, 3] (see Section 2.1) and in other coordination problems, such as string stability in lines of vehicles [4] (see Section 6.4). In other cases, the exchange of information can be affected by additive noise [281]. Another frequent issue is the presence of *quantization* effects. In particular, I devoted a series of papers [46, 47, 117, 44, 59, 116, 61] to explore the impact of limited precision in average consensus dynamics, both for discrete-time and for continuous-time systems. Quantization issues are not limited to engineering systems, but also appear in social systems, individuals may only exchange coarse signals during social interactions [61]. Whether in social or engineering systems, quantization effects that can degrade the ability of the multi-agent system to coordinate [117, 58, 116, 61]. Systems subject to quantization are described in Sections 3.5 and 3.6.

Minimal and gossiping interactions. Given the restrictions that affect the communication, and the costs (e.g. energy consumption) that communication entails, it makes sense in an engineering perspective to push the design so to ensure that the collective task is achieved with minimal coordination. In this engineering perspective, I have particularly explored the framework of *gossiping* communications, which I discovered through the seminal work [34]: the agents communicate in pairs and the pairwise interactions take place according to a random process [36, 99, 236]. In my work, partly described in Chapter 4, I have applied this framework to several different problems: consensus, optimal deployment, PageRank computation, estimation from pairwise measurements.

Heterogeneity. Heterogeneity is ubiquitous in multi-agent and network systems, but is often disregarded in the theoretical analysis for the sake of simplicity. However, an assumption of homogeneity may hide important features and prevent the understanding of relevant phenomena. A remarkable example is the problem of *synchronization*. Synchronization is an ubiquitous phenomenon whose mathematical representation in the multi-agent system framework boils down to the agents’ states converging to each other. Synchronization can be either a naturally emerging property or a design objective. When the agents are identical to each other (homogeneous case), synchronization is quite natural to achieve through distributed interactions: it is enough for each agent to use a static controller that combines the outputs of the neighbors into an input for the agent. Heterogeneity, however, is a huge obstacle to synchronization and makes its analysis rather delicate [214]: in fact, a landmark paper [291] has proved that achieving asymptotically exact synchronization through continuous static couplings is impossible between heterogeneous agents. Instead, the agents should either use non-smooth couplings [78, 77] or use dynamic controllers that contain identical copies of a common dynamical system (the so-called internal model). Since some heterogeneity is unavoidable in reality, its apparent ability to disrupt synchronization has struck my intellectual curiosity. I have therefore investigated in detail the consequences of not using an inter-

nal model and thus suffering some synchronization error [250] and I proposed a learning approach to circumvent the need for the precise *a priori* knowledge of the internal model [17].

Scalability to large networks. In networks, size matters and large networks are the most challenging ground for network science and engineering. For this reason, I devoted a special attention to performance analysis on large networks and to the scalability of my methods (by scalability, I mean the feasibility of the approach on networks with any number of nodes N). Distributed control methods, in which we design the interactions between the agents, are naturally scalable in their implementation. However, their performance may depend on N . Most performance measures are prone to degrade for large N : for instance, the convergence rate of the linear consensus algorithm can grow as N^2 in poorly connected graphs [105]. Nevertheless, a large size can be beneficial when it allows to “average out” undesired random effects [122], see Section 2.1. In many engineering and natural systems, however, we are unable to design the network or the interactions. In such cases, one has to ensure the scalability of the methods, whether they are deployed for analysis, estimation, or control. Multiple design strategies are possible, which resort to suitable approximations: one strategy is to approximate the large graph by a smaller one, by clustering nodes together, and then work on the clustered graph [183, 185]; another strategy is to approximate the large graph by a continuum and then work in the continuous domain [280, 285, 208]. These approaches are described in Chapter 5.

As per the **mathematical instruments**, I have resorted to many different tools, chiefly from the theory of systems and control, but also from optimization, learning, and more generally applied mathematics. In comparison with the typical control system toolbox, my signature tools have been graph-theoretical methods to account for the network structure, non-smooth dynamical systems and hybrid systems, and tools from probability theory to deal with various sources of randomness.

Graph theory is obviously crucial for an effective modeling of networks and particularly to achieve a precise understanding of the role of network topology. Therefore, I have been extensively using graph theory in most of my research: for instance, I have related the eigenvalues of the graph Laplacian to various measures of performance of distributed algorithms [105]. In some occasions, I have provided contributions that may be of independent interest to graph theory, for instance regarding the properties of the Laplacian matrix and of its spectrum: on the relationship between graph dimension and average effective resistance of the graph (a.k.a. Kirchoff index) [247] and on the approximation of large dense graphs by graphons [285].

Nonsmooth systems, that is, systems of differential equations with a discontinuous right-hand side, arise in my research because of two reasons. The first reason is *quantization*, that is the presence of continuous variables (representing the states of the agents) and of discrete variables (representing discrete values for communication or control inputs) [59, 86]. The second reason is that, when interactions are encoded by a graph, topology changes can abruptly change the interaction laws. In mathematical terms, such abrupt changes result in a discontinuity of the right-hand side of the differential equation that describes the dynamics. For instance, in many agent-based models it is reasonable to assume that communication can take part only between agents that are close to each other [60]. In some cases, quantization and topology changes are better represented by discontinuities (i.e., jumps) of the state itself (as opposed to mere discontinuities of the right-hand side): in this case the dynamics takes the form of a

hybrid system, that is, a dynamics that involves both continuous flows and discontinuous jumps. I have specifically dealt with hybrid systems on several occasions: in [59] I have proposed a novel and effective quantizer that exploits hysteresis to ensure the stabilization of consensus without chattering; in [87] I have used hybrid systems to define a self-triggered consensus dynamics with ternary controls; and in [127] I have exploited the flexibility of hybrid models in opinion dynamics. In this context of non-smooth and hybrid systems, I have devoted special care to (technical, but important) questions about the existence of solutions and the occurrence certain pathologies, such as having multiple (non-unique) solutions and having attractive points that are not equilibria [62]. Some results about nonsmooth consensus-seeking systems are described in Chapter 3.

Randomness. In many relevant contexts, such as mobile robotic teams and social dynamics, interactions between agents happen asynchronously, either because of the absence of a common clock or because they only happen sporadically. The times of asynchronous interactions and the consequent changes in the agent states are often well described by a **random processes**. I have considered randomized interactions in several papers [122, 237]. Furthermore, I have extensively used random models to describe failures in communication [45, 4] or in sensing [118]: in fact, contrarily to much of the control system community, I tend to prefer stochastic failure models and probabilistic guarantees over deterministic worst-case analysis. In this context of modeling failures, I often resort to traditional assumptions of independence in time and space: however, I have also studied models with some degrees of correlation [104, 122]. Finally, I have often used random graph models to describe or simulate networks [27, 93].

1.2 Collaborations, community spirit, and interdisciplinarity

My research has been marked by a strong drive towards collaborating with fellow scientists. I have always had the sense that research is the collective outcome of the scientific community's labor and I have tried to behave accordingly. During my international career that has developed across Italy, the Netherlands and France, my research has been shaped by many fruitful collaborations with colleagues based not only in my host countries but also all over Europe, in the USA, in China and in Japan.

My inclination towards interactions with colleagues is reflected in my involvement for the scientific community in multiple aspects, including several national and international committees (see Appendix A). One instance are several Technical Committees of the International Federation of Automatic Control (IFAC) and of the IEEE Control Systems Society: see the reports [233, 232, 72, 70, 71] of the IEEE CSS Technical Committee on Networks and Communications. Another form of interaction that I prize is the peer review process and, therefore, the editorial service. Besides an extensive service as anonymous referee since 2009, I have been serving since 2013 as Associate Editor for several journals, including (in chronological order) the International Journal of Robust and Nonlinear Control, the Asian Journal of Control, the IEEE Control Systems Letters, and Automatica.

Even though the above roles have been within the control systems community, in my research I have actually always been keen on interdisciplinary collaborations. I have explored collaborations both with disciplines that are traditionally closer to control

systems, such as robotics, communications engineering, or data science, as well as with disciplines that may seem further away, such as zoology, epidemiology, and sociology and media science. Some of these collaborations are reflected in Sections 6.1, 6.2 and 8.2. In fact, I believe that the theoretical perspective of dynamics, control, and networks offers a powerful lens through which multiple real-world problems can be fruitfully approached.

1.3 Mentoring students

As I believe that research is done by people, I have always devoted a special attention to mentoring students and younger researchers. I have been (co)advisor to seven doctoral students (two ongoing) and about fifteen master students. I have published peer-reviewed research papers with all of my PhD students and postdocs and with ten of my master students. Several of my master students have pursued PhD studies and three of my former students (Luca Vassio, Wilbert Samuel Rossi, Nelson Chan) have obtained Assistant Professor positions.

A full list of the students whom I have formally supervised is given in Section A.5. In this section, I prefer to give an account of my mentoring experience by describing the collaborative research work that I have done with ten selected students, including some informal supervisions. For each of these students (in reverse chronological order), I will briefly present their research topic and results and I will refer to our joint publications and to the content of later chapters for more details.

Raoul Prisant (EEATS doctoral school, Grenoble, October 2023–ongoing)

Jointly advised with Federica Garin. Supported by ANR grant COCOON of which I am Principal Investigator.

Raoul’s ongoing PhD deals with using graphons and other graph limits for the analysis and control of large networks of interconnected dynamical systems (see Section 8.1). Some preliminary results on opinion dynamics on large graphs are available in the report [119].

Tommaso Toso (EEATS doctoral school, Grenoble, October 2021–ongoing)

Jointly advised with Alain Kibangou. Supported by IRGA (*Initiative de Recherche Grenoble Alpes*) grant ON-ROUTE.

Tommaso’s ongoing PhD deals with online navigation systems and their impact on traffic dynamics and mobility networks, through the route recommendations that they provide to users. Crucial to his thesis is the definition of a mathematical model of traffic flows on a network, which accounts for the users’ knowledge. Together with Tommaso, we have identified a suitable class of dynamical models and obtained results about their stability and steady-state behavior. These problems and results are described in Section 8.3 and in a series of publications [267, 266].

Maria Castaldo (EEATS doctoral school, Grenoble, November 2019–November 2022).

Jointly advised with Tommaso Venturini and Floriana Gargiulo. Supported by CNRS MITI 80 PRIME project “DOOM” (Disorders of Online Media) of which I was Principal Investigator.

Title: “Attention dynamics on YouTube: conceptual models, temporal analysis of engagement metrics, fake views”.

Maria’s PhD thesis [49] has focused on attention dynamics in social media [50], both by the study of original mathematical models [52] and the analysis of empirical data about engagement metrics in social media and more particularly on YouTube [53]. This line of work, which is still ongoing, is presented in Section 8.2. Maria’s interdisciplinary

work has been co-advised by two researchers in social science and more specifically in media studies (Venturini) and in computational social sciences (Gargiulo). Our collaboration has also included a study of the collaboration dynamics within the Polymath project [137]. After her graduation, Maria has joined Boston Consulting Group in Milan, Italy, as a senior data scientist.

Renato Sebastian Vizuete Haro (University of Saclay, Greater Paris, September 2019–September 2022).

Jointly advised with Elena Panteley and supported by ANR project “HANDY– Hybrid and Networked Dynamical Systems”.

Title: “Contributions to open multi-agent systems: consensus, optimization and epidemics”.

My collaboration with Renato has in fact taken place in two phases. In 2019, I co-advised his Master thesis with Federica Garin: together we worked on graphon approximations to large network systems [280, 285]. Renato’s internship effectively launched my research on graphons, presented in Section 5.3.

Afterwards, I co-advised his PhD [144] on Open Multi-Agent Systems (OMAS), where he explored several approaches to OMAS [281, 197, 198] and applications in large-scale distributed optimization [279, 282, 198]. This ongoing line of work is presented in Section 7.2. Renato Vizuete is now a FNRS researcher at the Université catholique de Louvain, Belgium.

Denis Nikitin (EEATS doctoral school, Grenoble, September 2018–September 2021) Jointly advised (50%) with Carlos Canudas De Wit and supported by ERC project “Scale-freeBack”

Title: “Scalable large-scale control of network aggregates”.

During his thesis [205], Denis has developed several original methods for scalable network control, working mainly towards two directions. The first approach was based on controlling local averages in the network [207], along the lines of the thesis of N. Martin (Section 5.1). The second approach involves constructing and studying a suitable PDE approximation of the large network-based ODE system [208, 206, 209], Section 5.2. After graduation, Denis has joined Wayze, an autonomous vehicles company based in London, UK.

Nicolas Martin (EEATS doctoral school, Grenoble, December 2016–February 2020) Jointly advised (50%) with Carlos Canudas De Wit and supported by ERC project “Scale-freeBack”

Title: “Network partitioning algorithms with scale-free objective”.

Nicolas, my first student in France, developed graph-theoretical methods towards a scale-free approach to large networks [183, 185]. In this approach, suitable regions of the graph are lumped together into “super-nodes”, thus producing benefits for estimation and control (Section 5.1). We have also explored applications to transportation networks and to epidemics [184]. After his graduation [182], Nicolas has been working in the popularization of science as a free-lancer.

Francesco Acciani (PhD student, University of Twente, September 2014–June 2022) Jointly advised (50%) with Geert Heijenk and Anton Stoorvogel

Title: “Control over unreliable networks: Consensus and platooning”.

Francesco was my first formal PhD student: his work, coadvised with an expert in networking and wireless communications (Heijenk), has introduced me to the problems of vehicle coordination (Section 6.4), which I have later developed with several other colleagues [139, 140]. The leitmotiv of his thesis has been accounting for packet losses in the control loop, either for the sake of analysis or to design mitigation strategies.

We published work on consensus algorithms [3] and string stable platooning [5, 4], in which packet losses are modeled as a random process and control design is adapted to their presence. Francesco completed his PhD [2] in 2022 while employed by Rabobank in Utrecht.

Wilbert Samuel Rossi (Turin & Twente, 2010–2022)

Wilbert and I had a multifold and long collaboration. I first co-advised his master thesis at Politecnico di Torino, then collaborated with him during his PhD, and finally had him work with me as a postdoc at the University of Twente. We published work on several topics: the effective resistance of graphs [247], the problem of estimation from relative measurements [248], opinion dynamics with k -neighbor interactions [241, 244] (Section 3.4), the convergence of message-passing algorithms for influence maximization in social networks [243, 240, 242], and the closed-loop evolution of opinions and personalised recommendations in online social networks [249] (Section 7.3). After his postdoc in Twente, Wilbert joined the University College, University of Groningen, Netherlands as Assistant Professor.

Joey Durham & Rushab Patel (UC Santa Barbara, 2008–2015)

Joey and Rush were students at the University of California Santa Barbara under the supervision of Francesco Bullo. During a series of extended visits there, I got deeply involved in their research about coordination of robotic teams for optimal deployment and coverage control. We eventually published three journal papers [220, 219, 99] and several conference papers together: this material is described in Sections 4.1 and 6.1. After graduation, Joey has become senior staff at Amazon Research, while Rush has been with several US companies including Raytheon and Skyryse.

1.4 The rest of this document

The rest of this document is organized into two parts and two appendices. Part I collects chapters that describe a few selected lines of research. I have chosen these lines of research based on criteria of originality, impact, coherence with the rest of my work, and whether they were developed through series of papers. These elements, as well as the contributions of my students, are discussed in each relevant Chapter. The topics are presented roughly in their chronological (and indeed also logical) order of development.

Chapter 2 contains results on randomized discrete-time consensus algorithms: this work features a prominent role of randomness in discrete-time consensus algorithms, which had previously been the main focus of my PhD thesis. Chapter 3 contains results on continuous-time consensus dynamics with discontinuous right-hand side: this is a substantial body of work that I started right after my PhD and that I have been developing until now. Chapter 4 contains results on other problems of distributed control and computation, including coverage optimization, PageRank computation, and distributed estimation: also in this chapter randomization plays a crucial role. Chapter 5 contains results on scalable methods for the control of large networks: all of these methods involve approximating the actual, large, network by another mathematical object, which can be a smaller network or a continuum. Chapter 6 contains results that pertain to specific application domains in engineering, biology, and social sciences. All chapters of Part I discuss the contributions in light of the literature and, where I found it suitable, present relevant open problems.

Part II is about future and ongoing work. Chapter 7 contains a general perspective, in which I set to work on the control of large-scale socio-technical systems. Chapter 8

contains some concrete research ideas, which I intend to develop in the near future or which I am already pursuing to implement the general perspective. Both chapters contain accounts about recent work of mine that is immediately relevant to the research perspective.

The main body of the document is followed by two appendices. Appendix A contains a detailed Curriculum Vitae, which includes the complete lists of my students (Section A.5) and of the research grants that have supported my research (Section A.8). Appendix B describes my publication output and in particular contains the complete list of my publications.

Part I

Sampling from my research past

Chapter 2

Discrete-time consensus algorithms

In the average consensus problem, each node of a network is given a number and has the goal of computing the average of these numbers, with the constraint of communicating with its neighbors only. That is, the averaging algorithm has to be distributed and communication is only allowed following the links of the network¹. This problem can be solved by simple linear methods, which can take the form of discrete-time dynamics or continuous-time dynamics. The former case is addressed in this chapter, the latter will be addressed in the next one. In both cases, given a set of nodes \mathcal{I} of cardinality N , the state x shall take values in $\mathbb{R}^{\mathcal{I}}$.

In discrete time the dynamics (also referred to as algorithm) is defined as

$$x_i(t+1) = \sum_{j \in \mathcal{I}} a_{ij} x_j(t) \quad \text{for all } i \in \mathcal{I}, t \in \mathbb{Z}_{\geq 0}, \quad (2.1)$$

or more compactly the form

$$x(t+1) = Ax(t),$$

where for every $i, j \in \mathcal{I}$, a_{ij} is nonnegative and can be different from zero only if i can receive information from j (there is a link from j to i) and where $\sum_{\ell \in \mathcal{I}} a_{i\ell} = 1$. Matrix A thus has non-negative entries that sum to one along each row: matrices with these properties are said to be *stochastic* matrices. This iterative dynamics converges to consensus as long as the interaction parameters a_{ij} s encode a connected and aperiodic network (in the sense that there is a globally reachable aperiodic² node). If additionally³ $\mathbf{1}^T P = \mathbf{1}^T$ (for instance because the matrix is symmetric), then the convergence value is the average of the initial values and thus the average consensus problem is solved.

¹Until we explicitly say otherwise, the network is assumed to be static (independent of time). I will not be recalling the basic graph-theoretic jargon and background, for which I refer to my exposition in [105, Chapter 1].

²Given a graph $G = (\mathcal{I}, E)$ and $i \in \mathcal{I}$, let the period of i be the greatest common divisor of the lengths of the circuits (i.e. closed paths) in G to which i belongs. The node is said to be aperiodic if its period is one. Notice that if a self-loop (i, i) is present, then i is certainly aperiodic.

³We denote by $\mathbf{1}$ a column vector of appropriate dimension whose entries are all equal to 1. For matrix M , M^T denotes its transpose.

Extended background about this dynamics can be found in several textbooks including the most recent revision of [35] and my own monograph [105]. In this manuscript I will present some of my work on important variations on this basic dynamics. In this chapter, Section 2.1, I will present a discrete-time variation in which the updates are randomized (other randomized dynamics are presented in Section 4.3 and 6.1). Next, in Chapter 3 I will present several continuous-time variations in which the right-hand side is discontinuous.

2.1 The mean square error of randomized consensus algorithms

In the average consensus problem, sample values are available at the nodes of a communication network, and their average is approximated by running an iterative algorithm that has the sample data as the initial condition. Clearly, one should ensure that along the iterations of the consensus algorithm, no (or little) deviation from the correct average is introduced. This guarantee is not always simple to achieve: one issue is that many algorithms and dynamics over networks are not synchronous. In an average consensus algorithm, lack of synchrony may prevent the preservation of the average: for instance, if only one node changes its value at each update time, it is clear that preservation of the global average is impossible [108]. In such a case, it is crucial to understand whether these errors accumulate into a potentially large deviation between the computed average and the initial one.

In a joint work [122] with Julien Hendrickx, I consider linear randomized asynchronous averaging algorithms, and I analyze the mean square deviation of the consensus value from the initial average. We want to ensure that this error is small, so that averages are computed accurately. In particular, we provide conditions under which the mean square error tends to zero when the number of samples, i.e. the number of nodes, grows. More precisely, we consider discrete-time consensus systems with random updates that preserve the average in expectation, and we provide new bounds on the mean square deviation of the current average from the initial average. We show that under certain conditions the expected increase of the deviation is bounded proportionally to the expected decrease of the disagreement.

Problem statement and results. Given a set of nodes \mathcal{I} of finite cardinality N , we consider the discrete-time random process $x(\cdot)$ taking values in $\mathbb{R}^{\mathcal{I}}$ and defined as

$$x_i(t+1) = \sum_{j \in \mathcal{I}} a_{ij}(t)x_j(t) \quad \text{for all } i \in \mathcal{I}, t \in \mathbb{Z}_{\geq 0}, \quad (2.2)$$

where for every $i, j \in \mathcal{I}$, we assume $\{a_{ij}(t)\}_{t \in \mathbb{Z}_{\geq 0}}$ to be a sequence of independent and identically distributed random variables such that $a_{ij}(t) \geq 0$ and $\sum_{\ell \in \mathcal{I}} a_{i\ell}(t) = 1$ for all $t \geq 0$. System (2.2) can be conveniently rewritten as

$$x_i(t+1) = x_i(t) + \sum_{j \in \mathcal{I}} a_{ij}(t)(x_j(t) - x_i(t))$$

for all $i \in \mathcal{I}$ and $t \in \mathbb{Z}_{\geq 0}$, or in matrix form as

$$x(t+1) = x(t) - L(t)x(t) \quad t \in \mathbb{Z}_{\geq 0}, \quad (2.3)$$

where the matrix $L(t)$ is defined so that $L_{ij}(t) = -a_{ij}(t)$ if $i \neq j$ and $L_{ii}(t) = \sum_{j:j \neq i} a_{ij}(t)$. Namely, $L(t)$ is the Laplacian matrix of a (time-dependent) weighted graph $(I, E(t), A(t))$ where the entries of matrix $A(t)$ are defined as $[A(t)]_{ij} = a_{ij}(t)$ and the set of edges $E(t)$ contains all pairs (i, j) such that $a_{ij}(t) > 0$.

Let us remind that dynamics (2.2) is run in order to bring each node to a good estimate of the initial average $\frac{1}{N} \sum_{i \in \mathcal{I}} x_i(0)$. It is then convenient to define the notation

$$x_{\text{ave}}(t) = \frac{1}{N} \sum_{i \in \mathcal{I}} x_i(t)$$

and observe that the average evolves according to $x_{\text{ave}}(t+1) = x_{\text{ave}}(t) - \mathbf{1}^T L(t)x(t)$. Since under our assumptions $L(t)$ is independent from $x(t)$, we immediately deduce that the expected evolution of $x_{\text{ave}}(t)$, conditional on the previous state, is such that $\mathbb{E}[x_{\text{ave}}(t+1)|x(t)] = x_{\text{ave}}(t)$ if and only if $\mathbf{1}^T \mathbb{E}[L(t)] = 0$. In view of this fact, we assume the average to be preserved in expectation, that is, we will assume $\mathbf{1}^T \mathbb{E}[L(t)] = 0$, implying that

$$\mathbb{E}[x_{\text{ave}}(t)] = x_{\text{ave}}(0) \quad \text{for all } t \geq 0.$$

Consequently, we are left with the problem of studying the mean square error, or variance of $x_{\text{ave}}(t)$, that is $\mathbb{E}[(x_{\text{ave}}(t) - x_{\text{ave}}(0))^2]$. This study is the object of the following key result, which was proved in [122, Thm 1]. Define $\text{Var}(t) = \frac{1}{N} \sum_{i \in \mathcal{I}} (x_i(t) - x_{\text{ave}}(t))^2$: notice how this quantity naturally evaluates the dispersion of a vector around its average.

Theorem 2.1 (Mean square error estimate). *Let x be a solution of system (2.3).*

If $\mathbf{1}^T \mathbb{E}[L(t)] = 0$ and there exists $\gamma > 0$ such that

$$\mathbb{E}[L(s)^T \mathbf{1} \mathbf{1}^T L(s)] \leq \gamma \mathbb{E}[L(s) + L(s)^T - L(s)^T L(s)], \quad (2.4)$$

then for every $t \geq 0$, there holds

$$\mathbb{E}[(x_{\text{ave}}(t) - x_{\text{ave}}(0))^2] \leq \frac{\gamma}{N + \gamma} \text{Var}(0).$$

Notice that this result holds true for every t and is unrelated to the convergence properties of (2.3), which have been addressed elsewhere in the literature and which depend on the ‘‘average connectivity’’ of the network (see [107, 262] for convergence proofs and [95] for a survey of related results). Our focus, instead, is on the quality of the convergence value, in terms of its distance from the initial average, irrespective of convergence. If additionally the system converges to consensus (that is, $x(t) \rightarrow x_\infty \mathbf{1}$, for $x_\infty \in \mathbb{R}$), then Theorem 2.1 immediately implies that

$$\mathbb{E}[(x_\infty - x_{\text{ave}}(0))^2] \leq \frac{\gamma}{N + \gamma} \text{Var}(x(0)).$$

The interest of this general result becomes apparent when we are able to find a constant γ that is independent of N (or at least, such that $\gamma = o(N)$). In such a case, the mean square error goes to zero as N grows to infinity: that is, *the mean square error becomes negligible for large networks*.

Remarkably, this turns out to be possible in many relevant cases. Indeed, a constant γ can be easily found for systems in which only one agent updates its state at each time

step, or in which at most a bounded number of agents update their states simultaneously, as well as for systems in which the updates have small statistical dependence across the network: in all these case the mean square error tends to zero when the number of nodes grows. One relevant example is the *Broadcast Gossip Algorithm* proposed in [108, 107, 106]. As another application of this bound, I have shown that *random packet losses* have a negligible impact on consensus performance of large networks [121]. The impact of random independent packet losses is negligible in the sense that the consensus algorithm still converges with probability 1 (albeit more slowly) and the mean square deviation between the computed average and the initial average is $O(1/N)$. If also this deviation is undesired, it is possible to adapt the algorithm in order to compensate for the losses and preserve the average, at the price of further slowing its convergence: see my later work [3].

Discussion and impact. I am specially fond of on this result because it has been one of the first ones that I obtained without my advisors or other senior collaborators. I felt that the result was very innovative with respect to prior work, which focused on preserving the average on expectation and had no effective way to estimate the mean square error. Indeed, previous estimates [107, 262] are by-products of convergence properties and therefore depend on global properties of the communication network, like connectivity or graph spectrum. Instead, only local network properties, like the degrees of the nodes, play a role in the application of our result. Moreover, Theorem 2.1 is also very general as it makes very mild assumptions on the dynamics: I had previously obtained similar results under restrictive assumptions on the graph topology [104].

In summary, our result offers effective and easy-to-implement guidelines to the designer who needs to choose a network and an algorithm to solve an estimation problem, thanks to their generality and to their dependence on local network properties only. Due to their attractive features, our work has been noted by several researchers: for instance, they have been commended by the inaugural EUCA European Control Award laureate Luca Schenato during his keynote address at the 2014 European Control Conference.

Chapter 3

Nonsmooth consensus-seeking systems

In this chapter, I present several continuous-time consensus-seeking dynamics, in which the right-hand side is discontinuous. The backbone of the presentation is based on a tutorial exposition of mine [63], even though the scope of this chapter is broader as it includes dynamics that feature bounded confidence interactions (Section 3.3), k -neighbor interactions (Section 3.4), and quantized interactions (Sections 3.5 and 3.6). Before discussing the role of discontinuities and introducing the discontinuous dynamics that I have studied, let us begin by some background on the basic continuous-time consensus dynamics, against which the other dynamics should be compared.

3.1 Basic consensus dynamics

The continuous-time counterpart of the discrete-time consensus dynamics (2.2) takes the following form. Let a scalar x_i be associated to an individual $i \in \mathcal{I} = \{1, \dots, N\}$ and evolve in time according to the ordinary differential equation

$$\dot{x}_i = \sum_{j=1}^N a_{ij}(x_j - x_i) \quad i \in \mathcal{I}. \quad (3.1)$$

Since we assume that the interaction weights a_{ij} are nonnegative, this linear dynamics postulates that each individual is attracted by the other individuals with whom it interacts. Under very mild assumptions on the interaction pattern, this dynamics converges to a state of *agreement*, or *consensus*, where all components x_i are equal. See an illustration in Figure 3.1.

This dynamics can be written in vector form as

$$\dot{x} = -Lx,$$

where x is the vector of the state variables (one for each node) and L is a Laplacian matrix that encodes the network structure, has nonnegative entries on the diagonal, has nonpositive entries $-a_{ij}$ outside the diagonal and its rows sum to zero. This Laplacian




Figure 3.1: Evolution of a solution of (3.1) from random initial conditions on a cycle graph with 25 nodes.

dynamics converges to consensus as long as the interaction parameters a_{ij} s encode a connected network (in the sense that the network contains a globally reachable node). If additionally $\mathbf{1}^T L = 0$ (for instance because the Laplacian is symmetric), then the convergence value is the average of the initial values and thus the average consensus problem is solved¹.

Theorem 3.1 (Basic consensus). *If the graph underlying (3.1) is connected and the adjacency matrix A is symmetric, then for any solution $x(t)$ of (3.1) the following properties hold true:*

1. (contractivity and boundedness) $\overline{\text{co}}\{x_i(t), i \in \mathcal{I}\} \subseteq \overline{\text{co}}\{x_i(0), i \in \mathcal{I}\}$;
2. (average preservation) $x_{\text{ave}}(0) = x_{\text{ave}}(t)$;
3. (equilibria) x^* is an equilibrium point of (3.1) if and only if x^* is a consensus point, that is, $x_i^* = x_j^*$ for all $i, j \in \mathcal{I}$;
4. (average consensus) $\lim_{t \rightarrow +\infty} x(t) = x_{\text{ave}}(0)\mathbf{1}$.

More general versions of this result are available in the literature [35]: for instance, we can relax the connectivity assumption by not requiring symmetry and just assume the network to have a globally reachable node, at the price of the consensus value not being the average of the initial conditions.

In this chapter I will introduce and study four dynamics that are variations of this basic consensus dynamics. These variations, which are motivated either by coordination problems with communication constraints in engineering systems [37] or by opinion dynamics in social networks [130], will feature equations whose right-hand side is discontinuous (in fact, piecewise continuous). More generally, if we seek to model the dynamics of multi-agent systems, there are multiple situations where discontinuities arise, often in connection with discrete variables. I see two main categories of causes:

¹Here and later in this document, given a subset S of the Euclidean space, we denote by \overline{S} its topological closure, by ∂S its border, and by $\overline{\text{co}}S$ its closed convex hull.

- The presence of state-dependent limitations on the allowed interactions. This includes the case of physical, technological, or cognitive constraints that limit interactions to agents that are spatially or socially close to each other [146, 18, 37, 147, 13]).
- The presence of quantities taking values in discrete sets, when a finite number of choices is given (e.g. whether or which product to buy) or when communication takes place by means of a finite set of symbols ([67, 187, 62]).

When models are defined in discrete time, discontinuities of the right-hand side do not necessarily cause mathematical difficulties. Instead, *in continuous time discontinuities give rise to technical difficulties*, as the qualitative theory of ODEs is deeply based on the notions of continuity and differentiability.

Given the difficulties they entail, one could wonder whether discontinuities of the right-hand side are needed at all. Indeed, the discontinuities of some models can be avoided by defining suitable smoothed counterparts, which contain continuous approximations of the discontinuous functions. However, the connections between continuous and discontinuous variants are not trivial, as I have argued in [60]. Most importantly, discontinuities cannot always be avoided. These unavoidable discontinuities include the cases when the agents are allowed to interact with a fixed number of neighbors, but the set of neighbors depends on the state (Section 3.4), or when the interactions depend on “actions” taken from a finite set (Sections 3.5 and 3.6).

Since the rigorous analysis of differential equations with discontinuous right-hand side requires specific tools, I will first briefly introduce them in the next section. In the remaining sections of this chapter, I will define the dynamics of interest and state some of the most relevant results about them². We shall see that these variations of the consensus dynamics can (and in most cases, will) actually fail to converge to consensus, despite the fact that they describe phenomena of attraction between individuals. For this reason, I refer to these dynamics more broadly as *consensus-seeking* systems.

3.2 Differential equations with discontinuous right-hand side

In this section we summarize some notions which are essential in order to deal with systems whose right-hand side is discontinuous with respect to the state variable. I did produce a more extended introduction, tailored to the study of consensus-seeking systems, in [63]. The readers are in any case advised to consult the rich literature on the topic, including the tutorial [79] and the books [11, 112], as well as the original works about stability [16] or generalized solutions [57].

Let us consider the Cauchy problem

$$\dot{x} = f(x) \quad x(0) = x_0 \tag{3.2}$$

where $x_0 \in \mathbb{R}^N$ and $f : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is measurable and locally bounded. We will denote by Δ_f the subset of \mathbb{R}^N where f is discontinuous. When facing system (3.2), one should first of all choose which type of generalized solution is the most suitable for the system of interest. We shall consider Carathéodory solutions and Krasovskiy solutions.

²Most of these results are work of mine, often with students and almost always with my colleague Francesca Ceragioli who has introduced me to the analysis of non-smooth systems

The notion of solution nearest to the classical one is that of Carathéodory solution.

Definition 3.1 (Carathéodory solution). *Let $I \subset \mathbb{R}$ be an interval with $0 \in I$ and let $x_0 \in \mathbb{R}^N$. An absolutely continuous function $\varphi : I \rightarrow \mathbb{R}^N$ is a Carathéodory solution of equation (3.2) on I with initial condition x_0 if $\varphi(0) = x_0$ and if it satisfies (3.2) for almost all $t \in I$ or, equivalently, if it is a solution of the integral equation*

$$\varphi(t) = x_0 + \int_0^t f(\varphi(s))ds.$$

We say that a local Carathéodory solution corresponding to the initial condition $x_0 \in \mathbb{R}^N$ exists if there exist a neighbourhood $I(x_0)$ of x_0 , an interval of the form $[0, T)$ and an absolutely continuous function $\varphi : [0, T) \rightarrow I(x_0)$ such that $\varphi(0) = x_0$ and $\varphi(t)$ is a Carathéodory solution of (3.2) on $[0, T)$.

As we shall see for some of the dynamics in this chapter, Carathéodory solutions may easily fail to exist. In order to ensure existence of solutions, other generalized solutions have been introduced in the literature. In the context described here, Krasovskiy solutions can be easily and successfully used.

Definition 3.2 (Krasovskiy solutions). *Let $I \subset \mathbb{R}$ be an interval with $0 \in I$ and let $x_0 \in \mathbb{R}^N$. An absolutely continuous function $\varphi : I \rightarrow \mathbb{R}^N$ is a Krasovskiy solution of (3.2) with initial condition x_0 if $\varphi(0) = x_0$ and if for almost all $t \in I$ it satisfies the differential inclusion*

$$\dot{\varphi}(t) \in \mathcal{K}f(\varphi(t)), \tag{3.3}$$

where

$$\mathcal{K}f(x) = \bigcap_{\delta > 0} \overline{\text{co}}\{f(y) : y \text{ such that } \|x - y\| < \delta\}.$$

We say that a local Krasovskiy solution corresponding to the initial condition $x_0 \in \mathbb{R}^N$ exists if there exist a neighbourhood $I(x_0)$ of x_0 , an interval of the form $[0, T)$ and an absolutely continuous function $\varphi : [0, T) \rightarrow I(x_0)$ such that $\varphi(0) = x_0$ and $\varphi(t)$ is a Krasovskiy solution of (3.2) on $[0, T)$.

The following existence theorem is an immediate consequence of [11, Theorem 3, page 98], as the vector field $f(x)$ is measurable and locally bounded.

Theorem 3.2. *For any initial condition $x_0 \in \mathbb{R}^N$ there exists a local Krasovskiy solution of (3.2).*

We underline that *any Carathéodory solution is also a Krasovskiy solution*, even though the set of Krasovskiy solutions is much larger.

Remark 3.1 (Unicity of solutions & strong properties). *For non-smooth systems, generalized solution typically fail to be unique. In my work (and therefore in this Chapter), I have not much sought conditions to ensure uniqueness: instead, I have striven to deal with multiple solutions by proving results that are valid for all solutions. I refer to such results as strong, as opposed to weak results, which hold for some solutions only. For instance, a set is strongly invariant if all solutions that originate inside it remain in it.*

Remark 3.2 (Completeness of solutions). *For all the dynamics considered in this chapter, all solutions can easily be proved to be bounded. Therefore, following the argument of [63, Proposition 2], they can always be extended to the time interval $[0, +\infty)$ (that is, they are complete), as long as local solutions are guaranteed to exist for all initial conditions.*

Since the equations that we consider here have complete solutions, the analysis of their asymptotic behavior, for t going to infinity, is justified. In this perspective, the notion of equilibrium is useful. Equilibria are points where a solution can remain indefinitely³. In the context of generalized solutions, this general definition leads to distinguish between Carathéodory equilibria and Krasovskii equilibria.

Definition 3.3 (Equilibria). *A point x^* is a Carathéodory (Krasovskii) equilibrium of (3.2) if the function $\varphi(t) \equiv x^*$ is a Carathéodory (Krasovskii) solution of (3.2) for $t \geq 0$.*

Carathéodory equilibria are characterized by the equation $f(x) = 0$ while Krasovskii equilibria are characterized by the inclusion $0 \in \mathcal{K}f(x)$. Thanks to the multiplicity of solutions, there are examples of non-constant solutions issuing from an equilibrium point.

3.3 Bounded confidence

An interesting case of state-dependent interactions is the following, which is termed *bounded confidence* in the literature: two individuals are assumed to influence each other if their states are closer than a certain threshold (that we choose to be equal to 1 for simplicity). This interaction rule translates into the following differential equation:

$$\dot{x}_i(t) = \sum_{j=1}^N a(x_i(t), x_j(t))(x_j(t) - x_i(t)) \quad i \in \mathcal{I} \quad (3.4)$$

where $a(y, z) = \begin{cases} 1 & \text{if } |y - z| < 1 \\ 0 & \text{if } |y - z| \geq 1. \end{cases}$

This model⁴, which is a continuous-time counterpart of the opinion dynamics studied by Hegselmann & Krause [146], has been proposed by [28] and further considered in [60]. Very similar models have been considered in [200, 157, 293, 66, 288], while other models that involve assumptions of bounded confidence include [91, 193, 41, 295, 126, 127]. A specific survey on this kind of models has recently appeared [24].

The Bounded Confidence dynamics (3.4) allows for the existence of complete Carathéodory solutions and all Krasovskii solutions are proved to converge to an equilibrium. The structure of these equilibria is a set of separated clusters of individuals sharing the same opinion: see Figure 3.2 for an illustration.

Theorem 3.3 (Properties of BC [60]). *The following properties hold true for dynamics (3.4):*

- *for every initial condition, there exist a complete Carathéodory solution;*
- *the set of Carathéodory equilibria coincides with the set of Krasovskii equilibria of (3.4) and is equal to*

$$\mathcal{F} = \{x \in \mathbb{R}^N : \text{for every } (i, j) \in \mathcal{I} \times \mathcal{I}, \text{ either } x_i = x_j \text{ or } |x_i - x_j| \geq 1\};$$

- *all Krasovskii solutions converge to a point in \mathcal{F} .*

³Note that this is a “weak” notion of equilibrium: in case of multiple solutions, we do not require that all solutions remain at the equilibrium.

⁴This dynamics is sometimes referred to as *metric* bounded confidence, in order to distinguish it from dynamics (3.5), which is referred to as *topological* [64].




Figure 3.2: Evolution of a solution of (3.4) from a random initial condition on 25 nodes.

Dynamics (3.4) does not produce agreement, but instead clustering of individuals into groups characterised by agreement within each group and disagreement between groups. The equilibria are precisely those states in which individuals either agree or are enough apart not to influence each other. In [28, 60] it is proved that, due to robustness issues, one can expect the opinion values of different clusters to be approximately twice the threshold apart. This conjecture has motivated quite some work [288] but hasn't been settled yet: actually, a fine understanding of how the final opinions depend on the initial ones is still missing. However, some very interesting contribution has been recently obtained by a mean-field approach on the closely related Deffuant dynamics in [97].

3.4 k -neighbor interactions

The k -neighbor interaction model is obtained when agent i interacts only with a fixed number k of neighbors, where $1 \leq k \leq N$. More precisely, for every agent $i \in V$, his neighborhood $N_i(x)$ is defined in the following way: the elements of $V \setminus \{i\}$ are ordered by increasing values of $|x_j - x_i|$; then, the first k elements of the list (i.e. those with smallest distance from i) form the set $N_i(x)$ of current neighbors of i . Should a tie between two or more agents arise, priority is given to agents with lower index:

$$\dot{x}_i = \sum_{j \in N_i(x)} (x_j - x_i), \quad (3.5)$$

This continuous-time topological⁵ interaction model was first pointed out in [13], while several other models have considered topological interactions in different forms: see [82, 244] and references therein. Topological interactions can be motivated by the notion of

⁵The word “topological” is sometimes used to contrast this dynamics against others, like (3.4), which feature purely “metric” interactions. These topological interaction are much harder to study. Notice indeed that in (3.4), whether agents i and j interact only depends on x_i and x_j . Instead, in (3.5) the presence or absence of interaction between i and j depends on the stats of all agents.




Figure 3.3: A solution of (3.5) with $k = 3$ from random initial conditions. The non-smooth nature of the trajectories is also very visible.

Dunbar number [98] that indicates a cognitive limit in the number of significant relationships among individuals. This concept is particularly meaningful in the contemporary world, where potential contacts and available information seem to be unlimited but the attention span of the individuals is not⁶.

During the last several years I have devoted specific attention to dynamics (3.5) through a series of papers [65, 64, 125]: the following result is a summary of what is currently known about this dynamics.

Theorem 3.4 (Properties of k -neighbors). *The following properties hold true for dynamics (3.5):*

- For almost any initial condition there exists a Carathéodory solution and for any initial condition there exists a Krasovskiy solution.
- If $k = 1, 2$, then for any initial condition there exists a Carathéodory solution.
- If $k = 1$, then any Carathéodory solution converges to some x^* . Moreover, the limit vector is such that $x_i^* = x_j^*$ if and only if there is a path from i to j or from j to i in \bar{G} , where $\bar{G} = (V, \bar{E})$ is the associated graph for $t > 0$.

In fact, in contrast with a good understanding of the case of with only one neighbor, the convergence analysis is still completely open for $k > 1$ and for general Krasovskiy solutions.

3.5 Quantization in social interactions

Another relevant phenomenon is *quantization*, which occurs both in engineering and in social systems. In social systems, quantization may originate because the state variable, which is understood as a continuous opinion or belief of the agent, is “communicated” by the display of an action or behavior, which can take on discrete values only: for

⁶See Section 8.2 for an account of my interests in the dynamics of online platforms.

instance, the purchase of certain products. The discretization of the opinions in social systems has been observed by social scientists [132, Chapter 10] and addressed in several models including [271, 187, 67]. Quantization is also of relevance in opinion dynamics that take place on social media, where individuals are unable to share the full shades of their opinions, but instead typically react by stereotyped signals such as clicking *like* buttons.

In this context, I have proposed [60, 61] to investigate the following “quantized behaviors” model:

$$\dot{x}_i(t) = \sum_{j=1}^N a_{ij} [q(x_j(t)) - x_i(t)] \quad i \in \mathcal{I} = \{1, \dots, N\} \quad (3.6)$$

where we define the quantization of a real number simply by rounding it to the closest integer: $q(s) = \lfloor s + \frac{1}{2} \rfloor$.

This dynamics has good properties in terms of existence of solutions.

Proposition 3.1 (Existence for QB). *For any initial condition there exists a local Carathéodory solution of (3.6).*

However, more pathological behaviors arise when we consider the asymptotic properties of Carathéodory solutions, which can converge to points that are not (Carathéodory) equilibria⁷.

Definition 3.4 (Extended equilibrium). *Let $k \in \mathbb{Z}^n$ and f_k be*

$$(f_k)_i(x) = \sum_{j=1}^n a_{ij} (k_j - x_i).$$

An extended equilibrium is a point $x^ \in \mathbb{R}^n$ such that there exists $k^* \in \mathbb{Z}^n$ such that $f_{k^*}(x^*) = 0$ and $x^* \in S_{k^*}$.*

Notice that all (Carathéodory) equilibria are extended equilibria, but extended equilibria need not be (Carathéodory) equilibria. Indeed, $f_{k^*}(x^*) = 0$ does not imply that $f(x^*) = 0$ (where we use f to denote the right-hand side of (3.6) in vector form).

In [62] it was found that on complete and complete bipartite graphs, convergence to consensus is achieved for all initial conditions. However, simulations suggest that, in other type of graphs, solutions usually converge to non-consensus extended equilibria, hence the importance of such points.

The presence of these pathological attractors motivates us to consider also Krasovsky solutions, which have two important advantages. Firstly, Krasovsky’s definition is more general than Carathéodory’s, meaning that Carathéodory solutions are particular Krasovsky solutions. Hence, results that are established for Krasovsky also apply to Carathéodory. Secondly, Krasovsky solutions can not converge to points that are not (Krasovsky) equilibria, thus solving the mentioned pathological convergence.

In case of the quantized behaviours equation (3.6) we do not have a characterization of the set of equilibria, but we already have some significant information. First, we observe that consensus points of the form $h\mathbf{1}$ with $h \in \mathbb{Z}$ are Carathéodory equilibria.

⁷For smooth systems, if a (classical) solution converges to a point, then such point is a (classical) equilibrium. This property is not necessarily true for Carathéodory solutions of systems with discontinuous right-hand side. We refer the reader to [61] for the construction of these peculiar solutions.

Second, there exist equilibria that are far from consensus and are attractive for some Carathéodory solutions. An example (on cycle graphs) can be observed in Figure 3.4; another (on path graphs) is provided in the next result.

Proposition 3.2 (Far-from-consensus equilibrium of QB). *Consider (3.6) with an N -node path as underlying graph and all non-zero entries of the adjacency matrix A equal to 1. Then, there exists a Krasovsky equilibrium x^* such that*

$$x_N^* - x_1^* = \begin{cases} \frac{(N-2)^2}{4} & \text{if } N \text{ is even} \\ \frac{(N-1)(N-3)}{4} & \text{if } N \text{ is odd.} \end{cases}$$

Theorem 3.5 (Distance from consensus). *If $\varphi(t)$ is any Krasovsky solution of (3.6) and*

$$M = \left\{ x \in \mathbb{R}^N : \inf_{\alpha \in \mathbb{R}} \|x - \alpha \mathbf{1}\| \leq \frac{\|A\| \sqrt{N}}{\lambda_*} \right\},$$

then $\text{dist}(\varphi(t), M) \rightarrow 0$ as $t \rightarrow +\infty$.

We remark that this result is tight in the following sense: on some graphs, the estimate on the limit set is asymptotically tight for large networks in the sense of the Euclidean distance from the consensus. More precisely, if the graph is a path with N nodes and weights are uniform, for all points in the attractor M it holds true that $\frac{1}{\sqrt{N}} \|x - x_{\text{ave}}\| = O(N^2)$ as $N \rightarrow \infty$. At the same time, the equilibrium x^* that was constructed in the proof of Proposition 3.2 is such that (for odd N)

$$\frac{1}{\sqrt{N}} \|x^* - x_{\text{ave}}^*\| = \frac{1}{\sqrt{120}} N^2 + o(N^2) \quad \text{as } N \rightarrow \infty.$$

Hence, the estimate of M can not be improved in general in terms of distance from consensus. Details of these computations can be found in [62].

This model allows for the existence of complete Carathéodory solutions for every initial conditions, but Krasovsky solutions are preferred to avoid the pathology of solutions converging to non-equilibrium point. In general, a result of convergence to equilibria is missing, but a tight result of convergence to a set is available. Remarkably, there can be equilibria very far from consensus, in which the difference among different opinions of individuals is proportional to N^2 .

The sociological interest of this dynamics, also recalled later in Section 6.3, is highlighted by the following example, which is due to my student Raoul Prisant and is presented in [227].

Example 3.1 (Hidden consensus). *On an undirected ring graph, every node i interacts with $i-1$ and $i+1$ (modulo n). For x^* to be an extended equilibrium, it must hold that $x_i^* = \frac{q(x_{i-1}) + q(x_{i+1})}{2}$. If we consider $S_{k^*} = S_{(0,0,1,1,0,0,1,1)}$, we have that $x^* = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \dots) \in \overline{S_{k^*}}$ is a non-integer consensus extended equilibrium. The peculiarity of this point is that nodes agree on the opinion $\frac{1}{2}$, yet they actions behave differently. In the solution shown in the left plot of Figure 3.5, four opinions converge to $\frac{1}{2}$ from below, and as such are quantized to 0, and four from above, and consequently take action 1. From a social interpretation point of view, since only actions can be seen, not opinions, this situation would appear as a disagreement scenario. Thus, this type of consensus is undetectable.*




Figure 3.4: Evolution of a solution of (3.6), assuming the same initial conditions and graph as Figure 3.1.

The dynamics presents a number of additional pathologies, which are just now starting to be studied. My most recent contributions [227], which are joint work with my students Luca Cataldo and Raoul Prisant, include

1. proving that, on any graph, each extended equilibrium has a basin of attraction of positive measure;
2. proving convergence from any initial condition on line graphs: the limit point is an extended equilibrium but needs not be a consensus ;
3. showing the existence of limit cycles and of Zeno solutions on directed ring graphs. Interestingly, the constructions of these cyclic and Zeno solutions rely on the properties of the golden ratio.

These results highlight the richness of the qualitative behavior of the dynamics and the complexity of its study.

3.6 Quantization for the consensus problem

In engineering, quantization can represent communication constraints, where the state variable is communicated between individuals via a digital channel with finite data rate, and thus constrained to take on discrete values. In my paper [59], limited data rate is described as a quantization of the communicated values, so that the values $q(x)$ are communicated across the links, instead of the actual values x .

In this context, an effective consensus-seeking system is the following “quantized states” system studied in [59]:

$$\dot{x}_i(t) = \sum_{j=1}^N a_{ij} (q(x_j(t)) - q(x_i(t))) \quad i \in \mathcal{I}. \quad (3.7)$$

A simulation is provided in Figure 3.6. Note that the right-hand side features the quantized values of both states x_j and x_i : the presence of the quantized state $q(x_i)$ is




Figure 3.5: Quantized Dynamics on the undirected ring. Time evolution of the “hidden consensus” solution in Example 3.1.

crucial to ensure the “good” properties of this dynamics, which will be discussed below. Indeed, this dynamics preserves the average of the initial states. On the contrary, the right-hand side of equation (3.6) featured the quantized value of x_j , but not of x_i , which leads to more complex dynamics that may end up far from consensus.

Consensus dynamics with quantization have first been studied with engineering motivations, while seeking controlled dynamics that could lead to (approximate) consensus despite the constraint of quantization [48, 202]. Proposed in this context by me and my coauthors Francesca Ceragioli and Claudio De Persis [59], dynamics (3.7) better fits engineering applications than social dynamics, since in many engineering applications we have freedom to design the right-hand side, so long as we satisfy the communication constraints.

As per its mathematical properties, the quantized states dynamics (3.7) does enjoy global existence of Carathéodory solutions and thus requires to consider Krasovskiy solutions. The best property is that all Krasovskiy solutions converge to equilibria such that the quantized opinions are equal. This is not exactly consensus, as individuals’ opinions may slightly differ, but they agree on their quantized values.

The most powerful theorem about the convergence properties of (3.7) has been proved in [290], significantly refining my initial results [59].

Theorem 3.6 (Convergence to equilibria). *Let*

$$\mathcal{D} = \{x \in \mathbb{R}^N : \exists h \in \mathbb{Z} \text{ such that } h - \frac{1}{2} \leq x_i < h + \frac{1}{2}, \forall i \in \mathcal{I}\}$$

and assume that the graph has a globally reachable node. Then, the set of Krasovskiy equilibria of (3.7) is $\overline{\mathcal{D}}$ and any Krasovskiy solution to (3.7) converges to \mathcal{D} .




Figure 3.6: Evolution of a solution of (3.7) with the same initial conditions and graph as Figure 3.1.

Chapter 4

Distributed multi-agent and network systems

My research about distributed algorithm has not been limited to the important special case of consensus and closely-related consensus-seeking systems. Instead, I have considered several other problems and a variety of solutions and methods.

A natural extension of the consensus problem is the **synchronization** problem¹, in which the dynamics of the agents converge to a common trajectory. In essence, the consensus problem is a special case in which synchronization is sought towards a constant value (not evolving in time). Even though in this manuscript I will not discuss synchronization at length, this important and well-established problem has been a reference in my research. It has given to me the opportunity to address the issue of heterogeneity in multi agent systems in [250, 17], as briefly described in Section 1.1, and a benchmark problem to showcase the effectiveness of the continuation method that I described in Section 5.2: the latter method has allowed me to study large arrays of non-isochronous oscillators [209].

Further away from consensus and synchronization, there are distributed algorithms that solve other problems of estimation, control, and optimization. An example is the problem of **controlling the average opinion of consensus-seeking systems** in large-scale networks in an optimal way. When stubborn agents, which do not change their opinion, are mixed with agents that follow the usual consensus dynamics, the opinions of the latter converge to a non-trivial profile [69, 278]. With my students Luca Vassio [274], we studied the problem of optimally selecting stubborn nodes so to maximally move the average opinion of the system. Interestingly, we came up with a *message-passing* algorithm that allowed for a distributed solution. It was easy to show that this algorithm was exact and convergent on trees, but its behavior on general graphs with cycles was much harder to study. In subsequent papers with my postdoc Wilbert Samuel Rossi [239, 243, 240], we provided a general convergence proof and an analysis of the approximation error that is committed by the algorithm due to the presence of

¹Notice how in an engineering perspective, synchronization is a control objective that requires steering the agents states or outputs towards each other by exploiting their interactions. Instead, in a natural science perspective, synchronization is a phenomenon whose explanation is found in the interactions between the agents.

cycles. More recently, with my student Massimo Bini [27], we have also considered the related problem of controlling the average opinion by optimally adding links.

In the rest of this chapter I will concentrate on three sets of contributions. In Section 4.1 I will focus on a class of distributed algorithms for the optimal deployment of groups of robots (robotic networks considered again in Section 6.1 as an application domain). Next, in Section 4.2 I will quickly describe my work on distributed estimation with relative measurements: this problem has important similarities but also differences with the consensus problem. Finally, in Section 4.3 I will present a unifying result about the randomization of linear network dynamics.

4.1 Coverage control by gossiping robots

When I was working on robotic networks, we had a vision that in the not-too-distant future, networks of coordinated autonomous robots and drones would perform a broad range of environmental monitoring and logistic tasks. For example, groups of robots would enable novel logistic capacities in the transportation of goods and the delivery of services and resources to customers. For these future applications to see the light, load balancing algorithms will need to dictate how the workload is shared and assigned to the individual robots. In other words, robotic resources should be assigned and deployed to competing requests in such a way as to optimize some performance metric. Remarkably, load balancing problems in robotic networks are often equivalent to robotic deployment and environment partitioning problems. Indeed, in the transportation of goods or delivery of services, minimizing the customer waiting time is equivalent to a multivehicle routing problem and, in turn, to computing optimal depot positions and regions of responsibility.

Motivated by these scenarios, my award-winning paper [36] considers the two following interrelated problems. The *deployment* problem for a robotic network amounts to the design of coordination algorithms that lead the robots to be optimally placed in an environment of interest. Deployment performance is characterized by an appropriate network utility function that measures the deployment quality of a given configuration. The *partitioning* problem is the design of coordination algorithms that lead the robots to optimally partition the environment into subregions of interest. Each robot would then be tasked with servicing events in its own subregion.

In [36], Francesco Bullo, Ruggero Carli and I tackle partitioning and coverage control algorithms in innovative ways. First, we design algorithms that require only “gossip” communication, i.e., asynchronous, pairwise, and possibly unreliable communication. Gossip communication is a simple, robust, and effective protocol for noisy and uncertain wireless environments and can be implemented in wandering robots with short-range unreliable communication. Second, at a methodological level, we propose a change of perspective in coverage control and multicenter optimization. In earlier works [80], the state space for the coverage algorithms are the agents’ positions, i.e., as a function of the agents’ positions the environment is divided into regions and regions are assigned to each agent. Note that in this classical approach, every movement of an agent is reflected in a change of both its own assigned region and its neighboring regions. Clearly, this rigidity conflicts with allowing unreliable and asynchronous communication. Instead, the agents’ positions are no longer a concern in our approach: the state space is the space of partitions of the environment and the algorithm dictates how to update the regions. As the space of partitions is much richer than the space of the agents’ positions, we gain

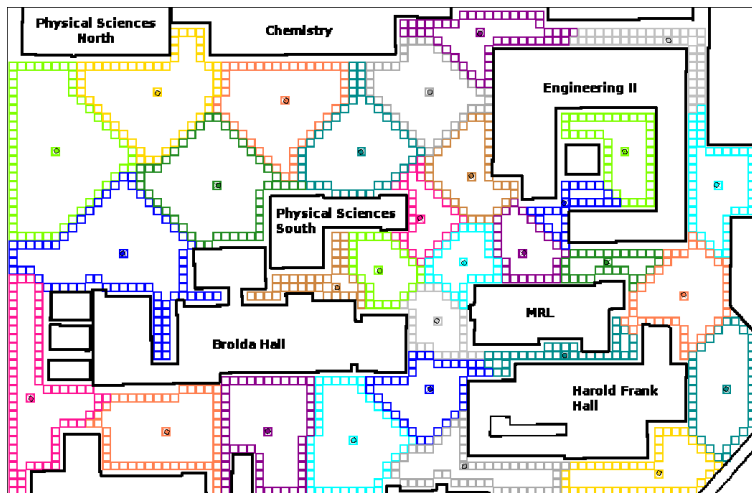


Figure 4.1: Simulated configuration of a self-deployed network of sensors in a complex environment (the campus of the University of California at Santa Barbara). For each sensor, its own responsibility region is shown using the same colour. This is joint work with UCSB student Joey Durham [100].

more freedom in the design of partition optimization algorithms and in particular the possibility to use gossip communication.

Our partitions-based framework was originally developed for compact subsets of the Euclidean space in [36], but can be naturally extended to partitioning of other metric spaces. Indeed, the paper [100] studies partitioning of graphs and proposes its application to partitioning non-convex environments with complex geometry. Crucial to this endeavour is producing a discretization of the continuous state space into a network of relevant points: our algorithm then operates on such network. In the example of Figure 4.1 the network is obtained by using a regular mesh and considering adjacent cells as neighbors. Paper [220] further extends this approach to cloud-based interactions between the robots, by allowing for the additional flexibility of covering the environment of interest with regions that can possibly overlap. The possibility to have overlapping regions also allows to assign multiple robots to service an event, as I studied in the context of Unmanned Aerial Vehicles [88].

4.2 Distributed estimation from relative measurements

In the problem of distributed estimation from relative measurements, we assume that each agent is attributed an unknown value and that relative measurements are pairwise (noisy) differences taken across the edges of a given graph. Possibly, also (noisy) absolute measurements of the unknown values can be available to the nodes. The objective of this problem is reconstructing the full vector of the agent values (each agent estimates her own value); see [22] for an introduction to this problem. Similar problems are studied in machine learning, where they are referred to as statistical ranking problems [160, 10, 212, 252, 148], and in mobile robotics, where they appear as cooperative localization problems in autonomous navigation [154, 42, 83]. These problems in machine learning and robotics are in general multidimensional and nonlinear, whereas the problem I

worked on with my students Wilbert Samuel Rossi and Nelson Chan is linear.

The linear nature of the problem allowed me and my students to provide explicit results that very clearly highlight the role of the network topology. Indeed, the network that describes the measurements determines the fundamental limitations of the achievable performance [245, 246, 248]. Using an intuitive interpretation of the problem that employs an electrical network of resistors, we have shown that the error of the least squares optimal estimator depends on the topology of the graph that encodes the measurements. The topology determines whether the error decreases to zero as the number of unknown variables grows to infinity: for instance, this happens on fully connected networks, but not on cycle networks.

Besides these fundamental limitations, I also worked on distributed algorithms for its solution². The simplest approach is a synchronous gradient algorithm [248], which bears a lot of similarities with a consensus algorithm, albeit for the presence of constant input term. Our study of this gradient descent algorithm has revealed that the number of steps required to achieve a certain accuracy (relative to the optimal estimator) does not depend on the measurement graph's topology or even on the number of nodes. Running the algorithm for k steps means that the estimate of each node is based on information from nodes at distance at most k on the graph. Therefore, the above finding indicates that only a limited portion of the graph, independent of its total size, needs to be explored by each node through the algorithm. In other words, some cooperation among the nodes is necessary for this distributed estimation problem, but too much cooperation is superfluous and even counter-productive. In contrast, other networked estimation problems, such as estimating a common parameter through an average consensus algorithm, always benefit from more cooperation (that is, a longer running time). Indeed, when all nodes measure the same quantity and use consensus to estimate it, the time to achieve a given accuracy can grow proportionally to the square of the number of nodes, necessitating full cooperation across the network. In this sense, the advantages of cooperation are limited in this problem, because its objective is the estimation of a “*diffused*” parameter. This property is not shared by other estimation problems where a *common* parameter is to be estimated and additional cooperation is always beneficial to improve the estimate.

4.3 Ergodic randomized network dynamics

Several passages of this document have already elaborated that randomization is specially natural in network dynamics, either by the uncertain nature of the network at hand, or by a design aimed at improving performance and robustness.

Despite the variety of relevant cases, there used to be no general tools for the design or the analysis of randomized algorithms over networks. In the paper [237], which is joint work with Chiara Ravazzi, Roberto Tempo and Hideaki Ishii, I challenged this gap in the literature. Our work focuses on a large class of randomized linear dynamics that do not have deterministic equilibria but are stable on average. This stability property ensures that the dynamics, although featuring persistent random oscillations, has an ergodic behavior. Our main contribution is precisely to prove an ergodicity result that can

²Besides the synchronous gradient algorithm, I also developed randomized algorithms with gossiping communication [237], which leverage the mathematical machinery presented in the Section 4.3, and algorithms that are able to automatically distinguish whether the measurements have small or large noise and adapt their estimator accordingly [234].

be readily applied to several network-based dynamics, where randomization apparently prevents convergence. As a consequence, the desired convergence property—holding in expectation—can be recovered by each node through a process of time-averaging. Remarkably, time-averages can be computed locally by each node and even without access to a common clock.

Consider the affine dynamics representing a time-invariant discrete-time dynamical system over a network, described by a directed graph $G = (V, E)$ with N nodes, with state $x(t) \in \mathbb{R}^V$, $t \in \mathbb{Z}_{\geq 0}$

$$x(t+1) = Px(t) + u, \quad (4.1)$$

where the matrix $P \in \mathbb{R}^{V \times V}$ is adapted to the graph \mathcal{G} and $u \in \mathbb{R}^V$ is a constant input. We have the following simple fact.

Proposition 4.1. *If P is Schur stable, then the dynamics in (4.1) converges to*

$$x^* = (I - P)^{-1}u$$

for any initial conditions $x(0) = x_0$.

We are interested in randomized versions of the dynamics in (4.1). More precisely, let $\{\theta(t)\}_{t \in \mathbb{Z}_{\geq 0}}$ be a sequence of independent identically distributed (i.i.d.) random variables taking values in a finite set Θ . Given a realization $\theta(t)$, $k \in \mathbb{Z}_{\geq 0}$, we associate to it a matrix $P(t) = P(\theta(t)) \in \mathbb{R}^{V \times V}$ and an input vector $u(t) = u(\theta(t)) \in \mathbb{R}^V$, obtaining a time-varying discrete-time dynamical system of the form

$$x(t+1) = P(t)x(t) + u(t), \quad (4.2)$$

with initial condition $x(0) \in \mathbb{R}^V$. We observe that the state $\{x(t)\}_{t \in \mathbb{Z}_{\geq 0}}$ is a Markov process because, given the current position of the chain, the conditional distribution of the future values does not depend on the past values.

It may happen that the dynamics (4.2) oscillates persistently and fails to converge in a deterministic sense: this behavior is apparent in the example of Figure 4.2. In view of this fact, we aimed to provide simple conditions which guarantee probabilistic convergence to the vector x^* . We say that the process $\{x(t)\}_{t \in \mathbb{Z}_{\geq 0}}$ is *ergodic* if there exists a random variable $x_\infty \in \mathbb{R}^V$ such that almost surely

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\ell=0}^{t-1} x(\ell) = \mathbb{E}[x_\infty]. \quad (4.3)$$

We now establish our main result for randomized dynamics over networks. Its proof can be found in [237].

Theorem 4.1 (Ergodicity of affine dynamics). *Consider the random process $\{x(t)\}_{t \in \mathbb{Z}_{\geq 0}}$ defined in (4.2), where $\{P(t)\}_{t \in \mathbb{Z}_{\geq 0}}$ and $\{u(t)\}_{t \in \mathbb{Z}_{\geq 0}}$ are i.i.d. and have finite first moments. If there exist $\alpha \in (0, 1]$, P and u such that*

$$\mathbb{E}[P(t)] = (1 - \alpha)I + \alpha P, \quad \mathbb{E}[u(t)] = \alpha u, \quad (4.4)$$

with P and u satisfying Proposition 4.1, then

1. $x(t)$ converges in distribution to a random variable x_∞ , and the distribution of x_∞ is the unique invariant distribution for (4.2);
2. the process is ergodic;




Figure 4.2: The randomized dynamics (left panel) fails to converge: however, its ergodicity can be highlighted by plotting its temporal averages (right panel).

3. the limit random variable satisfies $\mathbb{E}[x_\infty] = x^*$.

An example of the relation between randomized dynamics and time average is illustrated in Figure 4.2: this relation hinges on the underlying deterministic dynamics (4.1). Indeed, under the assumptions of the theorem, P is Schur stable in Proposition 4.1 and so is $\mathbb{E}[P(t)]$. Consequently,

$$\begin{aligned}\mathbb{E}[x(t+1)] &= \mathbb{E}[P(t)]\mathbb{E}[x(t)] + \mathbb{E}[u(t)] \\ &= ((1-\alpha)I + \alpha P)\mathbb{E}[x(t)] + \alpha u,\end{aligned}$$

and $\lim_{t \rightarrow \infty} \mathbb{E}[x(t)] = x^*$. The expected dynamics of the process (4.2) can thus be interpreted as a “lazy” (slowed down) version of the synchronous dynamics (4.1) associated to the matrix P .

This theoretical contribution can be applied to a wide set of linear dynamics, ranging from distributed PageRank computation to estimation in power systems, as we illustrated [123]. I am particularly proud of two applications: one about opinion dynamics in social networks (also discussed in Section 6.3) and one about estimation from relative measurements (defined in Section 4.2).

Regarding **opinion dynamics in social networks**, in [124] we addressed an open problem about opinion dynamics in presence of prejudices, emerging from the seminal work of Noah Friedkin [131] in mathematical sociology. Friedkin’s work on social influence exhibited a mismatch between the experimental setup, in which individuals communicated by phone (thus, in pairs), and the mathematical model, which featured synchronous linear iterations akin to (4.1). Theorem 4.1 offered a way to justify the equivalence between the dynamics postulated by Friedkin’s model and a dynamics based on pairwise interactions.

Regarding the problem of **distributed estimation from relative measurements**, we proposed in [235] an original ergodic algorithm to solve it. The classical algorithm can be interpreted as a gradient descent with a constant input: since we wanted to propose a randomized, asynchronous, algorithm, our approach was to replace the gradient descent with a *random coordinate descent* algorithm. Unfortunately, the presence of an input term representing the measurement was creating oscillations and preventing convergence of the algorithm. Theorem 4.1 has allowed us to characterize and understand these

oscillations: we were thus able to overcome the lack of convergence and propose an effective algorithm, by combining randomization and time-averaging.

In a broader perspective of algorithm design, the interest of this work of mine is that time-averaging can be a tool for smoothing stochastic oscillations that may appear in other randomized network systems³. Another example are randomized algorithm for PageRank computation. Indeed, oscillations are commonplace in randomized network systems: other authors have proposed different solutions, which essentially damp the system inputs in the long run: this goal is achieved through “under-relaxations”, that is, by using gains that decrease along time. The analysis of the resulting dynamics is often based on tools from stochastic approximation [33] or semi-martingale theory [226, Ch. 2].

³As introduced in Section 1.1, much research of mine has dealt with randomized asynchronous algorithms and with mitigating their potential drawbacks: another type of issue, that is, the failure to preserve the average in distributed averaging algorithms, was addressed in Section 2.1.

Chapter 5

Approximate methods for large-scale networks

The task of controlling large-scale networks is very difficult in the first place because of their large dimensionality, making the computation of traditional control algorithms too expensive. In systems of large dimensions, the number of sensors is often much lower than the number of states, which makes hard to identify the mathematical model of the system and to estimate its state. Similar issues arise regarding the number of actuators, which limits the control capabilities. Another difficulty is that the energy needed to control all nodes of the network can grow exponentially with the number of nodes, at least for some network structures [218, 169].

Therefore, in some cases, it can be preferable to define control or estimation problems for some output of the network rather than all individual states. These questions can in principle be addressed by leveraging classical control tools, such as the notion of structural controllability, which determines whether control objectives are achievable by conditions that only depend on the network structure [168, 96]. This approach has actually been quite popular around a decade ago [173, 102, 210, 222, 141]. In this context of structural linear systems theory, I have first explored quantitative notions of controllability/observability radii for networks in [25] and later the notion of functional controllability [68]: conditions for functional controllability of networks have intuitive interpretation in terms of paths on the network graph and allow to design effective algorithms with good complexity.

However, the very large size of many networks that are encountered in complex socio-technical systems¹, such as social networks, is bound to make classical control approaches not viable and demands radically different approaches. In the last decade, I have thus devoted much of my work to address this need and I have obtained encouraging results. The rest of this chapter is indeed devoted to present general control methods for large-scale linear and nonlinear dynamical systems, which I have developed with my doctoral students Nicolas Martin, Denis Nikitin, and Renato Vizuete. All of these methods, beyond their differences, originate from the recognition that a full description of the network system dynamics, down to each individuals agents, is impossible

¹I will give in Chapter 7 a broad presentation of the complex socio-technical systems that I am interested in, together with their challenges, including the challenge of size.




Figure 5.1: Scale-Free network with four identified regions, which we call “hubs” (shaded in yellow), which only interact with the rest of the network through a smaller number of boundary nodes (double circles). Hub regions can thus be controlled (or observed) by acting on (or measuring) only the boundary nodes.

to achieve and unnecessary for meaningful control objectives. Instead, one should focus on aggregate quantities that describe the system at large, such as the distribution of the agent states, or the average state of the whole network, its variance, or the average values in different regions of the network.

5.1 Node aggregation methods

A general approach to scalable control methods, which is referred to as “scale-free”, is based on the aggregation of the variables that belong to neighboring nodes. The aggregation is done in such a way to construct a scale-free reduced network, where the goal is to control the averaged state and the variance of the hubs, corresponding to regions or groups of nodes, and the control is applied to the boundaries of the hubs. A graphical illustration of the approach is provided in Figure 5.1. The same approach, *mutatis mutandis*, can be applied to estimation problems [204]. I contributed to this multifold approach by participating in the ERC Advanced Scale-freeBack project, which was lead by Carlos Canudas de Wit in Grenoble from 2016 to 2022. My contributions span several aspects of this scale-free approach in collaboration with our students Nicolas Martin and Denis Nikitin.

In the scale-free modelling and analysis of dynamic networks, the purpose is to reduce the network complexity by finding the appropriate level of scale aggregation, while imposing the control and observation model properties that are desired and at same time preserving the natural properties of the original, large, network system (such as, preservation of mass). My papers [183, 185] focus on a prototypical problem: reducing the network size while induce a power-law distribution in the aggregated network and at the same time preserving the invariant measure of its associated random walk. Additional properties of the partition of the nodes into hub regions can be imposed in this framework: for instance, in our paper [186] (in collaboration with Tokyo Institute of Technology) we focus on the constraint of the regions being connected.

This understanding of the scale-free network reduction, which was achieved by the thesis of N. Martin, has allowed us to move to a next phase of work in which we developed control design methods. In the thesis of D. Nikitin, we were able to develop a rather




Figure 5.2: Control design can be performed on the PDE level and then discretized back to be implemented on each node.

general control framework [207] in which the objective of controlling the network state is spelled out into a twofold objective: (i) control the average value within each hub area; (ii) reduce the dispersion around that average by leveraging extremum-seeking techniques.

This scale-free approach has been successfully applied to monitoring and controlling traffic networks [238]. Other applications are possible, however, such as thermal monitoring of buildings [203] and epidemics control. For instance, in [184], we have presented an application to the design of optimal vaccination schemes that target the hubs of the reduced network.

5.2 The continuation method

When considering approximate models for large networks, we naturally fall into continuous models. These limit models can take different forms. One way to define continuous limits is to regard, instead of the agent states, their *distribution*. The evolution of the distribution would then be naturally described by a partial differential (PDE) or integro-differential equation. A good approximation implies that control actions can be designed on the continuous system and have guaranteed performance on the original (graph-based) one; see illustration in Figure 5.2.

By the thesis work of D. Nikitin (also funded by the Scale-freeBack project) and a series of papers, we have approached this problem for sparse spatially-distributed systems by a very original approach that we have referred to as *continuation*. This work has reached a twofold research objective:

1. We have developed a rigorous and complete theoretical framework for the PDE approximation of high-dimensional systems of coupled ODEs [208]. This framework draws on a varied mix of mathematical tools including control and numerical methods and Kolmogorov–Arnold representation theorem. For linear systems, our results ensure that the PDE approximation can be arbitrarily accurate if the order of the PDE is sufficiently high, independently of the number of coupled ODEs (that is, of nodes in the network).
2. We have applied this framework to multiple applications in linear and nonlinear control, including swarms of autonomous robots [208], laser arrays [206], and spin-torque oscillators [209]. The latter application has been an absolute novelty for

the Automatic Control community.

Nikitin’s continuation method has also been applied to traffic networks [270] and, despite being quite recent, has already aroused interest from the control community [180, 181, 101]. Instead of reporting here the general theory, for which I refer to [208], I prefer to illustrate its main ideas and features by a concrete example, the SIS dynamics. This will be done in Section 5.4, which will also illustrate the approach described in the next section and compare them.

5.3 Graphons

Another effective way to define continuous limits is by the concept of graph function, or *graphon*. A visual intuition of what a graphon is can be gained with the help of a pixel picture. Starting from the adjacency matrix of a graph, if we visualize a 0 as a small white square and a 1 as a small black square, we can construct a pixel diagram, as we can appreciate in Figure 5.3. Letting the number of nodes go to infinity (and under technical conditions), we obtain a continuous object, a “heat map” of connectivity that characterizes the whole sequence: a pictorial representation of this convergence is given in Figure 5.4. Conversely, finite graphs can be generated by sampling from the continuous graphon: in this case, the properties of the finite networks can be inferred from the properties of the graphon.

The theoretical foundations of graphon theory were developed two decades ago [176, 31, 32] and are presented in Lovasz’ book [175]. Their usage is supported by a growing body of literature and specifically by a well-developed approximation theory that dictates how large the number of nodes N must be for the approximation to be accurate [12]. Graphons have thus become a popular tool to describe large networks in machine learning, and recent applications include centrality measures [12], link prediction [296], and large population games [215, 40]. Dynamics over large networks have also been approached by using graphons at least since the work of Medvedev [190, 191] and recent developments and applications include [166, 30, 7, 15, 201]. However, graphons have not been so much used in control theory, despite the ground-breaking work by Gao and Caines [136, 135]. The latter has laid down a theory of the approximate control of linear network systems by combining the theories of graphons and of infinite dimensional systems. By this theory, graphon dynamical systems are formulated in an appropriate infinite dimensional space, in order to represent arbitrary-size networks of linear dynamical systems. Consistency between the models is defined as the convergence of sequences of network systems to the graphon system.

Inspired by previous results on centrality measures [12], with my student Renato Vizuete, we have been able to use graphons to define performance *metrics* that quantify system-theoretic properties like stability, controllability, or sensitivity to noise. So long as a graphon approximation of the network is available, these metrics can be computed from the graphon at low computational cost and approximate well the system-theoretic properties of the corresponding dynamics on graphs of large-but-finite size [280, 285]. An intuitive example are the stability properties of the SIS dynamics, which I will discuss in Section 5.4. Other examples are given by the spectral properties of the graph Laplacian: indeed graphons approximate well the spectral gap of the Laplacian and the Kirchoff index (or average effective resistance of the graph). I have presented these results about the Laplacian spectrum in [285].




Figure 5.3: Pixel diagram for a simple graph.




Figure 5.4: A sequence of graphs (and their pixel diagrams) that converge to a graphon.

Another way to use graphons is to define dynamics that are associated to a graphon, just like the usual network dynamics are associated to a graph. The idea behind such definitions is that the dynamics on the graph should be, under suitable but hopefully mild conditions, a good approximation of the original dynamics on a (large) graph. Later in this manuscript I will give two concrete examples of this endeavor. First, in the next section I will introduce the graphon SIS (5.3) as an approximation of the SIS dynamics on networks that I have recently studied with my postdoc Aurélien Velleret. Second, in Section 8.1 I will describe some ongoing work on the graphon Laplacian dynamics. Other examples can be found in the literature of the last decade, including random walks [224], consensus dynamics [30], and Kuramoto models [191, 190].

5.4 Illustrative example: Epidemics on networks

For the illustration and concrete comparison of the two methods, continuation and graphons, that I introduced in the last two sections, let us take the example of a simple epidemic model on a network.

Let us consider the susceptible-infectious-susceptible (SIS) epidemic model, which describes the propagation of a single communicable disease in a susceptible population. The transmission of the pathogen occurs when infectious hosts transmit the disease pathogen to healthy susceptible individuals. The infectious period extends throughout the whole course of the disease until the recovery of the patient, warranting a two-stage model: either infected or susceptible. In its classical version, the model is governed by the following equation

$$\frac{dI}{dt} = -\delta I + \beta(1 - I)I, \quad (5.1)$$

where I is the fraction of infected individuals in a homogeneous population and δ, β

positive parameters. This model is of *mean-field* type, that is, it assumes that the whole population can be described by a scalar variable. For such a description to be sound, the population has to be homogeneous and well-mixed.

However, we very well know that populations may happen not to be homogeneous and their interactions not fully mixing. In an inhomogeneous population (where individuals are clustered in communities such as cities or other social groups), a better model would be a *networked SIS*

$$\frac{dI_i}{dt} = -\delta I_i + \beta \sum_{j=1}^N a_{ij}(1 - I_i)I_j, \quad (5.2)$$

where the index $i \in \{1, \dots, N\}$ corresponds to a given community and the nonnegative coupling coefficients a_{ij} s are the entries of the adjacency matrix that encodes the network of inter-community interactions. In such a model, each of the populations is assumed to be well-mixed and the inter-community network takes care of the heterogeneity.

The *graphon SIS* [92] extends on this idea of interaction network and assumes the infected fraction to be $\mathcal{I}(t, x)$, where t is time and the continuous variable $x \in [0, 1]$ generalizes the index i . The resulting model is

$$\frac{\partial \mathcal{I}}{\partial t} = -\delta \mathcal{I} + \beta(1 - \mathcal{I}) \int_0^1 W(x, y) \mathcal{I}(t, y) dy, \quad (5.3)$$

where $W : [0, 1]^2 \rightarrow [0, 1]$ is the graphon that replaces the graph and the sum that was present in (5.2) has been duly replaced by an integral. The continuous graphon is suitable to replace the graph so long as the graph is large, because the graphon is the limit, in a suitable sense, of a sequence of graphs with increasing size, as was informally discussed earlier. A rigorous justification of Equation (5.3) can be obtained as the limit of an agent-based model and is presented in the forthcoming publication [93] with my postdoc Aurélien Velleret.

The properties of (5.3) are indeed consistent with those of (5.2). Let us consider the question of the stability of the origin (that is, the disease-free equilibrium). The stability condition for the finite-dimensional model for a given network reads

$$\lambda_1(A) \frac{\beta}{\delta} < 1,$$

where $\lambda_1(A)$ is the largest eigenvalue of the interaction matrix A . In the continuous SIS model, the stability condition [92] is $\|T_W\| \frac{\beta}{\delta} < 1$, where $\|T_W\|$ is the largest eigenvalue of the integral operator T_W associated to W . Besides the analogy, the graphon condition is insightful for large networks sampled from the graphon. It has been shown in [280] that the disease-free state is stable with high probability for networks of large dimension N if

$$\frac{\beta_N}{\delta_N} N (\|T_W\| + \phi(N)) < 1, \quad (5.4)$$

where $\|T_W\|$ is the spectral norm of the graphon and $\phi(N) = O((\log N/N)^{1/2})$. Note that the scaling factor N originates from the definition of graphon.

A *continued SIS* can be derived by applying the general method from [208]. For the sake of this discussion, we assume for simplicity that the network is space-invariant and one dimensional, so that agent i (located at position x_i) interacts with agents $i - 1$ and $i + 1$




Figure 5.5: Space-invariant system of nodes aligned in 1D line with dynamics determined by couplings a_{ij} that are nonzero for $j = i + 1$ and $j = i - 1$.

(located at distance Δx), as illustrated in Figure 5.5. We then perform second-order Taylor expansions of function $\mathcal{I}(t, x)$, centered in x_{i-1} and x_{i+1} , to obtain

$$\mathcal{I}(t, x_{i\pm 1}) = \mathcal{I}(t, x_i) \pm \frac{\partial \mathcal{I}}{\partial x}(t, x_i) \Delta x + \frac{\partial^2 \mathcal{I}}{\partial x^2}(t, x_i) \frac{\Delta x^2}{2}$$

and we substitute these into the networked SIS equation (5.2). This calculation leads to obtain the partial differential equation

$$\frac{\partial \mathcal{I}}{\partial t} = -\delta \mathcal{I} + 2\beta \mathcal{I}(1 - \mathcal{I}) + \beta \epsilon (1 - \mathcal{I}) \frac{\partial^2 \mathcal{I}}{\partial x^2}, \quad (5.5)$$

where $\epsilon = \Delta x^2$ is a positive weight that depends on the spatial distribution of the population. For simplicity we wrote the continued SIS in a one-dimensional space and using derivatives up to the second order: approximations in higher dimensions (featuring a Laplace operator) and with higher-order derivatives can be easily derived. In fact, our results in [208] ensure that the higher the order of the derivatives, the more accurate the approximation of the equation.

We can see from this example that approximations by graphon and continuation give rise to infinite-dimensional systems that are rather different from each other: the former produces an integral equation that is a direct counterpart of the graph-based equations, whereas the latter produces a PDE that is visually more distant from the networked dynamics but lends itself to an intuitive physical interpretation as a diffusion process.

Joint application of both graphons and continuation can be envisaged for systems that feature both short-range, spatially localised, interactions and long-range interactions. For instance, a joint SIS approximation

$$\frac{\partial \mathcal{I}}{\partial t} = -\delta \mathcal{I} + 2\beta \mathcal{I}(1 - \mathcal{I}) + \beta \epsilon (1 - \mathcal{I}) \frac{\partial^2 \mathcal{I}}{\partial x^2} + \beta' (1 - \mathcal{I}) \int_0^1 W(x, y) \mathcal{I}(t, y) dy \quad (5.6)$$

could be appropriate for an epidemic model on a large geographical scale, featuring both short-range contacts within families or neighbourhoods (the diffusion term) and long-range contacts by air travel (the integral term).

Chapter 6

Applications in engineering, biology, social sciences

As I briefly explained in the Introduction, Section 1.1, my research has been directly inspired by challenging problems from various application fields, ranging from engineering to biology and social sciences. In this chapter, I am going to provide an account of some of these interests and their relations with my theoretical results.

6.1 Optimal deployment of sensor and robotic networks

Distributed algorithms for information fusion in **sensor networks** have been a motivating application for much of my work around the consensus algorithm, including the gossiping and asynchronous versions that I reported about in Section 2.1. The need for control algorithms in sensor networks, however, is not limited to sensor fusion and cooperative estimation. Instead, other kinds of (distributed) algorithms are also necessary in order to ensure the proper functioning of sensor networks. For instance, the sensors may need to be suitably deployed in a target environment in order to provide good coverage of it. This kind of problems has lead me to consider distributed algorithms for the *optimal deployment of robotic teams*, where the robot is a self-propelled unit that moves in the environment in such a way to optimize its ability to sense the environment or to service possible events that happen therein [37]. Deployment algorithms for optimal coverage that require minimal, gossiping or asynchronous communication have been described in Section 4.1. Besides the design of distributed algorithms, my interests have also included the robustness of the optimal deployment to sensor failures: this concern has lead to an elegant analysis of optimal 1-dimensional coverage under random failures [118].

My algorithms have also been tested in simulations and hardware experiments, showcased in Figure 6.1 and where the contribution of students J. Durham and M. De Roo has been fundamental. In [100] we tested our algorithm for partitioning non-convex environments with complex geometry on a small group of physical rovers, augmented by a larger set of simulated rovers. In [88], we used small drones to test a coverage algorithm with overlapping regions, which allows multiple robots to be tasked with jointly

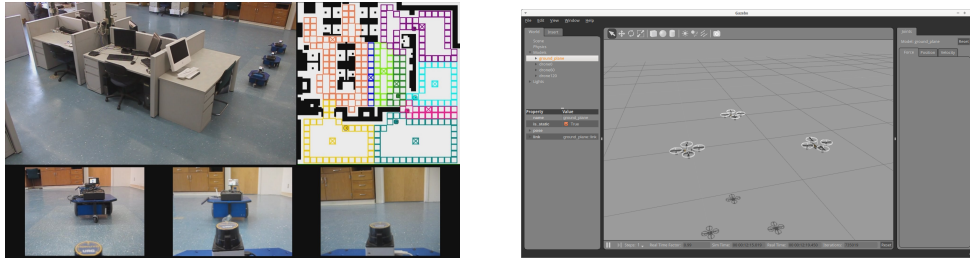


Figure 6.1: Robotic platforms in my coverage control experiments. Left: a group of rovers (some real, some simulated) [100]. Right: gazebo simulation prior to deployment of a small drone fleet [88].

servicing an event.

6.2 Animal groups: global shape from local interactions

The collective motion of animal groups has been a fascination of mine for many years, which was fostered by the inspiring readings of [164, 81, 18, 19]. Unlike consensus or synchronization problems, where individuals typically experience attraction forces alone, the behavior of animal groups arises from a delicate interplay of attraction and repulsion forces.

My approach to studying these phenomena hinges on the premise that *the shape and movement of animal groups emerge directly from the interactions among individuals*. These interactions are shaped by cognitive limitations and sensing constraints. Cognitive limitations dictate that individuals can only monitor a small number¹ of closest neighbors [18], while sensing constraints, often visual in moving vertebrate groups, restrict the field of view to the front or sides (see Figure 6.2).

Combining these few elements (attraction, repulsion, limited neighborhood awareness, and anisotropic sensing) has proved to be extremely fruitful. In fact, I have shown in [82] that these ingredients can replicate² a diverse array of group shapes and patterns, from rounded swarms common among insects and certain fish schools to linear formations seen in elephant or wolf packs, and even V-shapes and echelons observed in flying birds like geese (see Figure 6.3). I have also proposed to apply these ideas to explain the pat-




Figure 6.2: The shape of the sensitivity zones around each animal (moving from left to right). Solid red lines delimit the repulsion zone; dashed blue lines delimit the attraction zone.

¹Section 3.4 delves into the consequences of limiting the size of neighborhoods in opinion dynamics.

²Similar connections between collective motion patterns and visual capabilities have recently been highlighted, independently of our work, by [55].




Figure 6.3: V-like formations obtained with $\alpha_r = 60^\circ$, $N = 30$. The group is moving horizontally from left to right. The plot on the right is a close-up on one of the Ve'es.

terns of birds roosting on wires: the resulting model has received *empirical confirmation* in a joint work with a team of biologists from Rutgers University [14].

6.3 The origins of disagreement in opinion dynamics

In the last twenty years the control community has developed a strong interest in studying differential models that describe, or are at least inspired by, the evolution of opinions and beliefs in human groups and social networks [29]. These models have been proposed by social scientists since the fifties [128] and more recently pushed forward by applied mathematicians and control theorists: most of these models are variations of consensus-seeking systems, akin to those studied in Chapter 3.

In this context, the nodes of the graph are individuals, an edge between two nodes means that they socially interact and the i -th component of the state represents the value of the i -th individual's opinion. This graph-based modeling approach follows a solid tradition in mathematical sociology [129, 229, 230] as well as in economics [158] and in the physics of complex systems [54, 134, 13]. A broad variety of models have been studied, as illustrated by several surveys [229, 230, 264, 24] and multiple results have been obtained, which connect the topology of the social network with the asymptotic behavior of the dynamics. Problems of identification, estimation, optimization and control have also been considered by myself and by others [94, 294, 27].

The fundamental assumption in these models of opinion dynamics is that when an individual communicates with another, their opinion is influenced or attracted by the other's. If one translates this assumption into a set of differential equations, then one gets a dynamics similar to (3.1): the latter was already proposed as an opinion dynamics model more than 50 years ago by [1]. This dynamics asymptotically leads to consensus, i.e. agreement of the individuals on the same opinion, except when there are different groups of individuals that do not communicate with each other, i.e. when the communication graph has separated connected components. However, it has been noted that agreement is rare in societies [130], even if individuals do communicate: for this reason, more complex models have been elaborated with the aim of explaining both agreement and disagreement.

In my reflection and research, I have identified and studied three key mechanisms that prevent consensus, which I am going to rapidly survey with an emphasis on my contri-

butions and on their sociological interpretations.

Limited confidence. A first explanation postulates that individuals, despite interaction with others, may disregard their opinion if the latter is too different from their own. Depending on the precise definition of “too different”, we end up with dynamical models like those in Chapter 3 (confidence radius, k neighbors), where different kinds of threshold phenomena reduce connectivity and prevent consensus. In Chapter 3 we focused on hard thresholds that lead to discontinuities, but smooth dynamics with similar interpretations can be envisaged: in my paper [60] I have shown that the qualitative behavior of their solutions is similar to those with hard thresholds.

Obstinacy and prejudices. In the seminal work [131], disagreement is explained as the effect of obstinacy, that is translated into the dependence of any individual’s opinion on its initial value. Stubbornness as the source of disagreement is also considered in other models, such as [194, 217], also in presence of randomized asynchronous interactions [6]. A contribution of mine in this context, which leverages the mathematical results of Section 4.3, has been to show that a mathematical model with synchronized interactions can be justified as the average dynamics of a randomized, asynchronous, model [124].

Discrete actions. Another possible source of disagreement is the phenomenon of “quantization” or the coexistence of discrete and continuous variables in the opinion dynamics. Dynamical models with continuous opinions and discrete actions have attracted significant attention in the last decade. These models are often referred to as CODA models (Continuous Opinions Discrete Actions) and have been proposed by [187] and later studied by several authors [188, 67, 272, 62]. The simplest model of this kind has been presented in Chapter 3 as equation (3.6). This way of including quantization in opinion dynamics can represent situations in which social influence is mediated through binary (or, more generally, discrete) choices by the individuals. In that continuous-time opinion evolution the individuals hold scalar *opinions* and take discrete *actions*. The aim of model (3.6) is not to have a realistic description of social interactions, but to emphasize the effect of quantization on the asymptotic behaviour, which can actually be quite disrupting. Whereas the consensus dynamics asymptotically reaches consensus as long as the nodes are connected, the dynamics with discrete actions yields convergence to non-consensus equilibria, and can even exhibit limit cycles and Zeno trajectories for specific graph topologies³. Owing to this complexity, obtaining a complete picture of its convergence properties has proved elusive, as we have described in Chapter 3. The sociological interpretation of these non-consensus solutions is a strong warning to mind the difference between inner opinions and displayed actions in social dynamics.

These three explanations above are my favourite: I find them compelling and sound representations of real-life phenomena, even though the specific mathematical models that I have contributed to elucidate should be seen as conceptual models rather than detailed descriptions of reality that fit experimental data.

However, I should point out that the literature also includes other explanations. Popular ones are the presence of contrarians [133] or of negative interactions, i.e. negative

³The discovery of these special solutions was made possible by the work of my students L. Cataldo and R. Prisant.




Figure 6.4: Representation of an heterogeneous line of vehicles.

weights in the adjacency matrix. An example of the latter is the celebrated Altafini model [8] that has been later extended in multiple directions including time-varying networks [228]. I have recently studied its generalization to graphons in [119]. Similar dynamics on “signed graphs” may also feature randomized interactions [254] or bounded confidence [9]: a comprehensive survey is provided in [253].

6.4 Vehicle groups: communication and heterogeneity

The management of vehicular traffic is a major issue in modern society as it has an impact on air and noise pollution, carbon gas emissions, travel time, and their social and economic consequences. As connected and automated vehicles (CAVs) are introduced in the traffic stream, there is need for control technologies to manage their coordinate behavior and their interaction with traditional ones. The engineering community is hoping to use CAVs both as sensors to detect congestion and as actuators to fluidify the traffic flow. Until recently, CAV research has focused on methods to reach and maintain the desired behavior of a group of CAVs. More recently, instead, researchers have collectively realized the importance of understanding the interaction between traditional and connected autonomous vehicles, their system-theoretical properties, and the need to design novel control laws to steer them according to performance measures such as *string stability* [260] and *coherence* [20]. The purpose is not only to ensure good behavior of “mixed” groups (comprising both CA and traditional vehicles) in realistic situations, but also to assess the potential of CAVs to complement existing approaches to traffic control such as variable speed limits and intelligent traffic lights.

My approach to these questions has been to focus on models of *lines of vehicles*⁴. Vehicle lines are useful models that serve to define analysis, estimation, and control design problems, relevant to both engineering autonomous vehicles and studying traffic in urban and highway roads. In studying a line of vehicles, we focus on longitudinal motion and neglect lateral motion, although we sometimes account for lane changes. A line of vehicles model is thus determined by three ingredients:

- The *vehicle model*, with a more or less detailed internal dynamics;
- The *interaction topology* that describes which vehicles interact with each other;
- the *interaction model* that specifies which information is shared between vehicles, via sensing or via communication.

These line of vehicle models are suitable to formulate and study several control problems:

⁴My work on vehicle string started thanks to the PhD thesis of Francesco Acciani [2] and later continued through several master students including V. Giammarino [139] and C. Magnetti Gisolo [140]. On this research I was also much influenced by my collaborations with Maria Laura Delle Monache and Simone Baldi.




Figure 6.5: Representation of a string-stable system through its relevant variables. These plots are based on simulations for a line of 14 vehicles with a 0.3s headway from my work [4].

- *Platooning*, by which I mean the stability of steady, equi-spaced and equal-velocity uniform motion of the vehicles.
- *String stability*, which is an enhanced form of stability that accounts for transient behaviors. String stability requires a uniform bound on the relevant variables (inter-vehicle distances, speeds, accelerations, control inputs): see Figure 6.5. Crucially, a string stable interconnected system dampens disturbances as they travel downstream along the platoon. Disturbance amplification is not only a risk for a safe operation of the platoon, but it also compromises traffic flow stability and road throughput.
- *Maneuvering*, which I intend broadly to include several complex tasks that are needed to operate an automated platoon: performing lane changes, reconfiguring the platoons by merging or splitting them, and allowing vehicles to join and to leave the platoon.

My research in the broad field of multi-vehicle coordination has dealt with all three problems above and has focused on two key issues, which are actually general to all multi-agent and network problems (see Section 1.1), but which are well exemplified by vehicles groups: the role of communication constraints and the presence of heterogeneities.

Communication constraints are an essential component of all coordination problem. In vehicle groups, sensing is possible only between immediate neighbors and is limited to position and velocity. Any information about further away vehicles, as well as internal variables to the vehicle and its engine, must be communicated through a radio channel.

Such information is crucial to the performance of coordination algorithms. However, the radio channel may have limited communication radius and can be prone to issues like delays and packet losses. Contrarily to much research that emphasizes the effects of delays [298, 211], I have focused my attention on message losses, which I have described through a random loss model. In collaboration with my PhD student Francesco Acciani, I have designed cooperative adaptive cruise control algorithms that ensure a graceful degradation of string-stability performance in presence of stochastic losses [4]. Crucially, this contribution has required to adapt the deterministic notion of string stability to a stochastic setting⁵.

Heterogeneities are a thorny issue for multi-agent systems. Since homogeneous models are easier to study, most research assumes homogeneity. However, it is pretty clear that only in very specific situations one can assume all vehicles to obey the same dynamics and behavioral roles: think of a platoon that comprises lorries with different loads, or otherwise vehicles with different drivelines (Figure 6.4). In collaboration with experts in adaptive control, I have developed adaptive strategies for the coordination of systems with heterogenous and uncertain (that is, unknown) dynamics. My results, published in a series of paper, span from abstract synchronization problems [17] to applications to merging of platoons [170, 171].

Another important issue, though more commonplace in control science, is the presence of non-negligible nonlinearities. The latter can appear in the vehicle dynamics or in the interactions. Most of the multi-vehicle control literature concentrates in linear models, on which the analysis can go very far but which can fail to represent real systems. Even if I have mainly used linear models, I do have included two important sorts of nonlinearities in my studies.

- In [171] I have included engine and braking constraints, which are modeled as saturations.
- In [140] I have considered nonlinear inter-vehicle interactions according to the well-known Optimal Velocity Model of vehicle dynamics.

In presenting my working tools, I have been referring to lines of vehicles. In fact, an important tool for me has been studying **rings of vehicles**. Studying rings of vehicles, instead of straight lines, brings several advantages for the analysis. First of all, it removes border effects and renders the system *space-invariant*⁶. The space-invariance of rings allows to focus on the crucial point of the interaction between vehicles. Indeed, arranging vehicles on a ring is not only convenient for the analysis, but also for the empirical experiments: a good example are the experiments that reproduce stop-and-go waves in [258, 257].

For these reasons, I have made rings of vehicles prominent in my recent research [139, 140]. Indeed, rings can be seen as a proxy model for infinite lines, but with the crucial advantage of a finite-dimensional state space that is both more treatable from the

⁵An independent analysis of stochastic losses has been provided by [231], where losses are seen as the cause of stochastic delays in the communication.

⁶Elsewhere in my research I have taken advantage of on space-invariance properties. An example is the continuation methods in [208], which is more intuitively developed for space-invariant systems, even though it is not restricted to them. Moreover, since my PhD thesis I have extensively worked on Abelian Cayley graphs [105, Section 1.5], a family of graphs that well represents networks that have a space-invariant interconnections in a d -dimensional space [178]. For instance, I have used them as case studies in my analysis of the Kirchoff index of graphs [247] and in the study of randomized consensus [104].

mathematical point of view and more intuitive for the transition to practice. Indeed, infinite strings of vehicles have been occasionally considered in the literature [109, 113], but caution is advised in extrapolating conclusions from the infinite to the finite case or vice-versa [84].

Part II

Scientific Perspective

Chapter 7

Dynamics of large-scale socio-technical networks

My research objective for the near future is contributing to the theoretical foundations of analysis and control tools for large-scale socio-technical networks. I will devote this chapter to describe what are such networks, their features and the challenges that they bring, together with the tools that I envision to develop.

Socio-technical systems are dynamical systems that describe large and evolving populations of heterogeneous agents, whose physical or social interactions can be mediated by digital and algorithmic systems. I refer to this mediation as the “cyber” component, which complements the “natural” dynamics at the social or physical level.

This broad definition can cover multiple examples of complex systems, including social media, the global markets of commodities and goods, the electric power grid, and mobility and transportation systems. Personally, I identify mobility systems and social media as the application domains that I want to focus on.

- In mobility systems, individuals make choices (for instance, regarding their mode of transportation and their route) that lead to physical flows of vehicles and travellers on the transportation networks. At the same time, individuals elaborate their choices depending on their needs and perceived utility, but also due to imitation, peer-pressure, and other social interactions. Nowadays, the individuals’ choices can be based on real-time information about the state of the network, which is provided by navigation apps.
- In social media, individuals interact, consume contents, and eventually form their opinions within digital platforms that act as algorithmic gate-keepers of the flow of information. Information is generated and transmitted by the individuals themselves as users or by content creators. The platforms adapt to their users by learning the preferences of the latter and thus personalizing their offer of content, via their recommendation systems.

Similar definitions of socio-technical (or techno-social) networks, with a similar scope of application domains, have been around for some time: early examples can be found in the physicists’ science of complex systems, for instance [277]. This early perspective on complex techno-social networks already emphasizes the interplay between the social

and technological components and identifies the opportunity brought by the explosion of sensing and computing power. However, it fails to describe the role of digital systems in mediating and enhancing the interaction: simply, such role was not prominent yet¹.

In what follows, I will first concentrate on the key features of large-scale networks, which make them challenging for control, and next focus on the importance of the cyber component. Section 7.1 is concerned with the size of the networks, its consequences, and its challenges for control. Section 7.2 is concerned with dealing with populations of agents whose composition can evolve with time, i.e. Open Multi-Agent Systems. Then, Section 7.3 is concerned with understanding and describing the cyber loops and contains a case study of modeling a recommendation system.

7.1 The challenges of large uncertain networks

The systems I am interested in are first and foremost networks, and networks with specific features that make their study challenging. This first section is devoted to describe these features with respect to their network structure and their size.

I believe that networks constitute an essential paradigm to describe a huge variety of large complex systems, including social networks, transportation networks, supply chains, power networks, and multi-agent systems. These networks are the environment of complex dynamics that (depending on each case) may take the form of dynamics of opinions, epidemics, traffic flows, or collective motion and navigation of individuals. The theory of Automatic Control needs substantial advancements to manage such complex network dynamics, not only because of the need for accommodating the network structure, but also because achieving control and estimation objectives using standard methods is made computationally and practically intractable by the ever-growing network size. Instead, these large networks and the dynamics therein require adapted tools for modeling, learning, monitoring, and control: these tools must be scalable to networks of large size.

Compelled by the need for scalability, which has already been the topic of Chapter 5, I have become convinced of the interest of *continuous* network models. Traditionally, networks are described by graphs, that is by discrete mathematical objects composed of a set of nodes and a set of edges (that is, pairs of nodes) that encode their relations. This classical approach is inherently prone to scalability issues and should be transcended to deliver truly scalable mathematical methods. Instead, being continuous models independent of the network size, they will lead to methods that are inherently scalable. However, it is clear that graphs cannot be forgotten altogether, since real networks are in practice observed as nodes and edges. Hence, the crucial part of my research will be on the interface between continuous and discrete domains, where we aim to take advantage of continuous tools to understand and manage real discrete networks. Figure 7.1 illustrates this general idea of going back and forth between discrete and continuous domains.

¹In order to emphasize the importance of this “cyber” component, I have often referred to the complex socio-technical networks that I am interested in as *cyber-social* networks, or sometimes *cyber-physical-social* networks when the physical nature of the dynamics cannot be neglected. Nevertheless, I prefer to put aside this language here, mainly because of its similarity with the keyword “cyber-physical systems” that was popular in automatic control some years ago. The latter connection might convey the idea that the cyber component is open to be designed. In fact, in the modern socio-technical networks, the cyber component is already in place and any design idea should contend with its current technological and economical structures.

Continuous methods bring two different types of challenges with their specific mathematical difficulties. Firstly, it is clear that we need to deal with continuous systems, which have infinite-dimensional state spaces and may be delicate to treat, especially in presence of nonlinear dynamics, thus requiring advanced mathematical methods. Most importantly, however, we need to carefully deal with the relationship between the continuous and discrete models. This observation implies the need to derive conditions for consistency between the continuous and discrete models and to derive quantitative approximation theorems, in order to ensure that results about the former bear significance for the latter. Relevant descriptions of models and methods can be found in Chapter 5: some ongoing work in this direction will be presented in Section 8.1.

Large size, moreover, is not only a difficulty in itself, due to the need for scalable methods. Large size, instead, naturally brings additional challenges. One such challenge is uncertainty. The larger the network, the less likely it is that any external manager can know it accurately. Instead, knowledge about the network is bound to be imprecise, especially if the network structure evolves with time². In such a case one may have to manage the network based on outdated information. Uncertainty asks for methods that are robust, in the sense that they do not require precise information about the structure of the network. Continuous methods are again promising in this perspective, because they abstract from the detailed discrete structure of the network.




Figure 7.1: Common framework of reducing (discrete) networks to continua, via graphons or continuation techniques.

7.2 Open networks

In many applications, including online social networks, teams of autonomous vehicles, smart power grids and the Internet of Things, the structure of the network is subject to change. First of all, the network nodes can drop some of their connections or can create new ones, for instance because they are moving in some environment and are bound to interact with their physical neighbors. Furthermore, also the set of the interacting node can change with time, as new nodes can join the network and others can leave it. Despite this empirical evidence, control-theoretic methods for networks often assume the network structure to be static. While on a short time scale this assumption can be justified, on a longer time scale the network is bound to change due to the addition or removal of both nodes and arcs. Actually, the larger the network, the more likely it is that it changes during the dynamics. Therefore, the open character of large-scale networks cannot be ignored in their study. I will refer to networks whose set of nodes evolves with time as to *Open Networks* or as to *Open Multi-Agent Systems*.

Despite their ubiquity, open multi-agent systems have received surprisingly little attention in control or contiguous fields. Instead, definitions of open multi-agent systems can

²The fact that network structures, and specifically their set of nodes, can evolve with time will be the topic of Section 7.2.

be found in the computer science literature, for instance in [156, 225], where they refer to software agents and to the problem of evaluating reputation in open environments: however, these are not dynamical systems. Dynamically evolving populations of agents have also been considered in other streams of literature. For instance, game theorists have looked at dynamic sets of players in population games [251, 179], while researchers in complex systems have explored the role of evolving populations in opinion dynamics models [43, 142, 265].

Returning to the control systems literature, it is apparent that despite the abundance of works in multi-agent systems from the systems and control community, openness is rarely explicitly included in a rigorous analysis, but rather just presented as a possibility to be left to the practitioner or to be explored through simulations: see for instance [297]. In multi-robot systems, where adaptivity to addition/removal of robots is crucial, some architectures accommodate for dynamic teams but the analysis is not developed enough to derive rigorous performance guarantees: an example can be found in my own work [220]. I blame the relative scarcity of results on the fact that openness implies some conceptual difficulties in adapting control-theoretic notions such as state or stability. In fact, the key intrinsic difficulty of open multi-agent systems is that persistent arrivals and departures of agents keep the system away from equilibrium.

Against the literature background that I was sketching in the previous paragraph, in the last few years there has been a growing awareness and interest in OMAS in control. Since the pioneering work [149] in 2016, a fair number of papers has appeared and some PhD theses have been devoted to the issue, including those of Zohreh Sanai, Charles Monnoyer de Galland, and Renato Vizuete (the latter under my supervision). Some of these progresses are described in a review paper that we have recently produced in [286]. Nevertheless, a comprehensive theory remains elusive.

In the rest of this section, I want to survey several possible approaches to OMAS in order to present a critical summary of the literature, including my own work, and to identify some remaining open problems. I will thus describe three scenarios, which correspond to open multi-agent systems of increasing generality: OMAS with constant number of agents, OMAS with bounded number of agents, and OMAS with unbounded number of agents. All three scenarios have seen contributions of mine and my students, but remain active and promising areas of research.

Systems with constant number of agents (replacements only). The simplest scenario is when the number of nodes is constant, but the nodes themselves can be replaced by new ones. One then analyzes such dynamic networks by modeling the replacement of a node as a jump in the state of the overall system and a change in the connectivity matrix of the network. This scenario clearly avoids some of the features of more general OMAS, but does incorporate the persistent disturbance due to the new agents. This disturbance requires to look for suitable notions akin to practical stability or input-to-state stability. More generally, this scenario allows to formulate questions related to the information that is available in the system. Indeed, in some engineering applications one would like the system to forget the information that is no longer relevant because it originated from units that have left the system [195]. Multiple types of dynamics have been studied in this context. For instance, through a series of papers [197, 198, 282, 279] with my student Renato Vizuete, I have studied the performance of coordinate descent algorithms for optimal resource allocation when nodes can be replaced. More generally, distributed optimization methods are being

extended to open networks by several researchers [151, 155, 145].

Finally, this scenario can be seen as an approximation of more general cases in which the number of nodes varies, subject to independent arrivals and departures that take place with similar rates. I have taken this perspective in studying the SIS dynamics in OMAS [284]. At the level of large populations of individuals (e.g., countries, cities), it is common to have similar rates of arrivals and departures such that variations of the size of the populations are negligible. However, even if the total number of individuals in a specific place can remain approximately constant, it is unrealistic to assume that exactly the same individuals stay in a fixed location during all the evolution of the epidemic. If we consider the same rate of departures and arrivals, which preserves the number of individuals, epidemics under these conditions can be approximated as a dynamical system subject to replacements of individuals: a similar approach has been taken for opinion dynamics by [43, 265].

Systems with bounded number of agents. The second scenario is when one can define a uniform upper bound on the number of nodes or one can assume that, at all times, the nodes in the network belong to a finite (though possibly large) set. In such a case, one can perform the analysis by taking this large set to be the set of nodes and distinguishing between *active* nodes, which interact and participate to the collective dynamics, and *inactive* nodes that wait idle. This approach allows to leverage the substantial body of methods that have been developed for (stochastically) time-varying networks. This type of approach has been taken by [273] for consensus and later by several others [261, 145], including myself in [285]. In that paper, we consider nodes that are randomly activated at each time step and run a consensus algorithm, which is disturbed by additive noise.

Systems with unbounded number of agents. The third, and more challenging, scenario is when *the dimension of the system is not known* in advance or limited [114, 115]. This generality entails some mathematical difficulty in defining an appropriate state space. In order to overcome these issues, two approaches have been proposed in the literature.

The first approach is to define more general notions of trajectories that can take place across different spaces. Such switching between spaces can in principle be achieved by the formalism of multi-mode multi-dimensional (M3D) systems introduced in [276], where the switching is allowed between a fixed number of spaces of possibly different dimensions. Despite some recent interest [292], the M3D approach to OMAS remains unsatisfactory because it requires to know all the possible vector spaces to define the discrete vector bundle over which the system evolves. A more interesting development in this spirit is the framework of *live systems* that has been very recently proposed in [192]: this modeling framework seems quite flexible and allows to define and study Input-to-State Stability notions.

The second approach involves taking a *mean-field* perspective³ and give up a precise description of all agents in the system. Instead, one would use some low-dimensional (or even scalar) *descriptor*, that is, scale-independent quantities. The problem then becomes how to measure the changes on the descriptors that are produced by the arrivals and departures. For relatively simple systems, the choice of this descriptor is evident, as

³Being an application of the mean-field perspective, it is clear that this approach can benefit from the continuum methods for networks that are described in Sections 5.2 and 5.3.

in the case of gossiping consensus where the variance is a good measure of disagreement [150, 196]. Nevertheless, the choice is not straightforward for more complex dynamics, for two reasons. Not only one has to select a descriptor that's actually informative about the system, but the descriptor should ideally have a proper, closed, dynamics. This cannot be expected on systems that are not homogeneous and “well-mixed” in their interactions. Even though one can make do with some approximations, this “closure” issue, which is in essence another instance of the ubiquitous moment-closure problem [165], strongly constraints the potential reach of of this approach. In [283], I provide some details on this issue, and a partial discussion, through the interesting case of Open Hegselmann Krause dynamics.

These two main approaches do not fully cover the literature. For instance, in [115] I have studied a problem of dynamic consensus (that is, consensus with inputs [162]) in OMAS and I proposed an intermediate approach in which asymptotic stability is studied in normalized norms, given by the usual norms of finite-dimensional vectors divided by the square root of the dimension of the vector. This approach turned out to be effective, at least for contractive systems, but is currently lacking a generalization to a broader scope of systems.

7.3 The “cyber” loop

The systems I am interested in are networks of social or physical interactions. Besides the network aspects, the systems I am interested in share a prominent role of digital elements that mediate the interactions. I generically refer to these components, which feature artificial intelligence or algorithmic decisions, as to the “cyber” component. This cyber component typically produces some sort of feedback loop into the physical or social dynamics: these feedback may be unintended and always deserve close scrutiny.

Indeed, more and more complex algorithms are nowadays deployed in complex socio-technical systems with a variety of objectives: assisting or even replacing human decision making, searching for relevant information, and personalising content and advertising. Examples of application domains include: recommendation systems for e-commerce and social media; micro-targeted and personalized advertising; navigation systems and route recommendations for drivers and pedestrians; AI-assisted decision making in the judiciary and credit system.

In online services, loop closures through recommendations are ubiquitous. The fundamental role of recommendations is apparent in e-commerce, where customers are happy to receive suggestions based on their previous actions, such as viewing or purchasing some other product. Recommendations are perhaps less apparent in social media [103]: nevertheless, everything that is shown in a user's Facebook NewsFeed (or a similar home page on other platforms) is the output of a recommendation algorithm that selects the most relevant content [143]. Relevance depends on popularity, immediacy, and the user's previously displayed interests. In all of these cases, the loop closure is apparent, its effects on the engagement are actively sought, but its long-term effects on the users are only partially understood [159, 138]. Even though recommendation systems have been recognized as feedback loops since [21], *this intuition has not been fully exploited to date*, possibly because the details of the algorithms are shrouded by industrial secrecy or because control theorists have not perceived and taken up this challenge.

In these domains, the analysis cannot be limited to the properties of the algorithms

in isolation but should embrace the complexity of *the interaction between the algorithmic component and the collective social/human component, with its dynamics and its inherent biases*. The economics literature, indeed, has long realised that human behaviours are influenced by multiple forms of cognitive biases, including confirmation bias, conformity bias, and authority bias. How do AI algorithms react to human biases? Unfortunately, it appears that their relationship is potentially problematic, for multiple reasons. First, algorithms that are trained on biased data may incorporate and even amplify such biases. Second, algorithms may potentially exploit biases to deploy unfair “persuasion techniques” against the humans, in order to improve engagement or sale metrics. To solve this problem, a substantial research effort in algorithmic fairness has developed solutions that can mitigate these biases at different stages of the AI pipeline by enforcing some metrics of individual or group fairness [56]. Although these attempts can be successful in the short term, they often do not perform equally well in the long term, i.e., after multiple rounds of the decision-making process [259, 172].

I am convinced that the underlying reason for the persistence and emergence of biases lies in the fact that the methods to mitigate them are designed for stationary systems, while the system itself dynamically evolves over time. Indeed, the system changes over time because its output (the decision) feeds back as input to the system itself. The result is that biases are perpetuated (or even reinforced) due to the existence of the feedback loop, despite enforcing static mitigation techniques.

Therefore, I deem crucial to understand the role of the feedback loops, and how they relate to the amplification of different types of bias that may be present at the human and at the algorithmic level. The understanding of the user-algorithm feedback loops will lay the necessary foundation for analyzing the dynamics of automated decision-making systems and pave the way for the design of long-term bias mitigation techniques in the future.

Inputs from multiple disciplines, including behavioral economics, computer science, machine learning and optimisation are needed to deliver an integrated approach to understanding the interplay of human biases and AI algorithms. Studying the interconnections, and specifically the feedback interconnections, of different systems is a central theme in the theory of Automatic Control, which therefore will be a crucial instrument. Within this integrated approach, I can define a few objectives:

- The definition of a mathematical framework, inspired by systems & control theory, for a rigorous analysis of the human/algorithm interaction. This framework shall be geared towards qualitatively identifying and quantitatively assessing how human biases, algorithmic biases, and feedback loops interplay.
- The development of detailed feedback loop models for concrete case studies in recommendation systems, their rigorous analysis and their application to algorithm design.
- The collection and analysis of data, both from online platforms “in the wild” and from controlled laboratory experiments, to measure user bias and how AI algorithms react to them.

The ultimate goal of this research will be the design of algorithms that are fair, in the sense that they avoid incorporating biases or amplifying biases and refrain from exploiting users’ weaknesses.

Central to my efforts to achieve this goal is defining and studying tractable models in which a user is affected by cognitive bias and interacts with an online recommendation




Figure 7.2: The closed loop between the user and the news aggregator. The diagram includes the variables exchanged by the two interacting dynamical systems, as well as their internal state variables.

system that provides her with personalized content recommendations. The models that we will define shall make explicit the dynamical feedback loop between the evolution of the user's internal state (representing for instance her preferences) and the personalized recommendation of items. The purpose of these models is twofold:

1. Rigorously studying properties of the interconnection in the long run and in the transient: in control theoretic terms, these may include boundedness and stability of trajectories or controllability, to complement performance measures written in terms of loss or regret.
2. Design recommendation systems that keep users' states stationary in a suitable sense [89]: in practice, designing recommendation systems which do not influence the preferences of individuals too much.

I will now present a relevant case study in this perspective.

7.3.1 A toy model of news recommendations

In online platforms, recommender systems are responsible for directing users to relevant content. In order to enhance the users' engagement, recommender systems adapt their output to the reactions of the users, who are in turn affected by the recommended content. I have identified the case of a user that interacts with an online news aggregator as a prototypical example of unintended feedback, as it exemplifies all aspects that were discussed above.

In joint work with my postdoc W.S. Rossi, we have studied the empirical and domain literature to reach a good understanding of the phenomenon and to distil its key elements. Based on this analysis, we have defined a tractable analytical model of a user that interacts with an online news aggregator, with the purpose of making explicit the feedback loop between the evolution of the user's opinion and the personalised recommendation of content. The model is summarized in Figure 7.2. More specifically, we assume that the user is endowed with a scalar opinion about a certain issue and receives news about it from a news aggregator: her opinion is influenced by all received pieces of news, which are characterized by a binary position on the issue at hand. The user is affected by a confirmation bias, that is, a preference for news that confirm her current opinion. The news aggregator recommends news items with the goal of maximizing the number of user's clicks (as a measure of her engagement): in order to fulfil its goal, the recommender has to compromise between exploring the user's preferences and exploiting what it has learned so far.

After defining suitable metrics for the effectiveness of the recommender systems (in our case, the click-through rate) and for its impact on the opinion, we perform in [249] both extensive numerical simulations and a mathematical analysis of the model. Our findings indicate that personalised recommendations markedly affect the evolution of opinions and favor the emergence of more extreme ones: the intensity of these effects is inherently related to the effectiveness of the recommender. We also show that by tuning the amount of randomness in the recommendation algorithm, one can seek a balance between the effectiveness of the recommendation system and its impact on the opinions.

Our work is starting to attract attention from different communities (control, complex systems, computer science) and its ideas are influencing opinion dynamics models [223, 167] and design strategies for recommendation systems [213, 90, 256].

Chapter 8

Ongoing work: theory and applications

I devote this final chapter to the description of research lines that are currently active and that implement the perspective from the previous chapter.

The first research line is a theoretical line of work that, along the lines of Sections 5.3 and Section 7.1, aims to develop methods for the analysis and control of dynamical systems on large graphs by using graph limits (essentially graphons, even though I am also interested in other graph limits).

The second research line is a more empirical one that aims at studying dynamics in social media through an interdisciplinary approach that combines modeling, control-theoretic analysis, and the analysis of data from social media platforms (with a focus on YouTube).

The third research line regards the effects of the widespread usage of navigation apps in mobility. In my broad perspective, I see it as another key instance of the “cyber loop” that I have emphasized in Section 7.3, in which information flows between digital systems (and ultimately, between people) have significant consequences for physical flows in traffic networks.

8.1 Graph limits: dynamics on graphs and on graphons

As I already explained in Sections 5.3 and 7.1, I believe that graph limits are very promising tools to address the challenges of large size and uncertainty that socio-technical networks present. I am thus actively working on them, with two main research directions. This work is ongoing through the PhD thesis of Raoul Prisant, co-advised with my long-time collaborator Federica Garin.

The first research direction deals with the properties of network Laplacians and consensus-seeking dynamics associated to the Laplacian (see Chapter 3). The study of graphon Laplacians is quite delicate because the associated operator is not, in general, compact. Building upon earlier work in [287, 12], I have provided with my student Renato Vizuete a detailed study of its spectrum in [285]. I currently intend to apply this knowledge to study dynamics of consensus, synchronization, and opinion dynamics. For concreteness,

I will provide an example drawn from ongoing work on the Laplacian dynamics (3.1). Let us remind the associated initial value problem with a slightly different notation, in which n is the number of nodes, g_i the initial opinion of node i , and $u_i^{(n)}$ the opinion of each node i that satisfies

$$\begin{cases} \dot{u}_i^{(n)}(t) = \frac{1}{n} \sum_{i=1}^n A_{ij}^{(n)} (u_j^{(n)}(t) - u_i^{(n)}(t)), \\ u_i^{(n)}(0) = g_i. \end{cases} \quad (8.1)$$

This should be compared with the dynamics on the graphon

$$\begin{cases} \frac{\partial u}{\partial t}(x, t) = \int_I W(x, y)(u(y, t) - u(x, t))dy \\ u(x, 0) = g(x), \end{cases} \quad (8.2)$$

where $u : [0, 1] \times \mathbb{R}^+ \rightarrow \mathbb{R}$ and $g : [0, 1] \rightarrow \mathbb{R}$ is the initial opinion distribution.

In order to appropriately compare solutions, define intervals $I_i = (\frac{i-1}{n}, \frac{i}{n}]$ for $i = 1, \dots, n$ and the piece-wise constant functions

$$\begin{aligned} W_n(x, y) &= \sum_{i,j} A_{ij}^{(n)} \mathbb{1}_{I_i}(x) \mathbb{1}_{I_j}(y), & g_n(x) &= \sum_{i=1}^n g_i \mathbb{1}_{I_i}(x), \\ \text{and } u_n(x, t) &= \sum_{i=1}^n u_i^{(n)}(t) \mathbb{1}_{I_i}(x), \end{aligned}$$

where $\mathbb{1}_{I_i}(x) = 1$ if $x \in I_i$ and 0 otherwise. Notice that W_n is the piece-wise constant graphon that represents the original finite graph. We thus have the following result¹.

Theorem 8.1 (Bound on approximation error). *Consider $W : [0, 1] \rightarrow [0, 1]$ measurable and $g \in L^\infty(I)$. If $u^{(n)}(t)$ and $u(x, t)$ are solutions of (8.1) and (8.2), then for all $t \in [0, T]$*

$$\|(u - u_n)(\cdot, t)\|_2^2 \leq (\|g - g_n\|_2 + C_u \|W - W_n\|_2) \exp(2T),$$

where $C_u = \sup_{t \in [0, T], x \in I} |u(x, t)|$.

The bound shows that the solutions of the two dynamics are close and converge to each other when the number of nodes grows to infinity.

The second research direction regards the relationship between the steady-state behaviors of dynamics on the graphons and those of dynamics on large graphs (as those in the previous paragraph and those I presented in Section 5.4 for the SIS). As a matter of fact, a closer look at the bound in Theorem 8.1 shows that although the result guarantees that solutions converge to each other as n diverges on any bounded interval, the presence of the exponential-in-time factor prevents drawing immediate conclusions about the steady-state. Instead, in control systems we are deeply interested in steady states, such as equilibria or limit cycles, and in their stability properties.

¹This statement is based on [190]: a variation that applies to signed graphon, that is, graphons that take values in $[-1, 1]$, is presented in my preprint [119]. I have preferred to include the statement from [190] because it has the same spirit and requires to define some less notation.

8.2 Online social media and their disorders

Online social networks, such as online blogging platforms and social media, are chief examples of complex systems where social and technological components interact. I refer to these systems as cyber-social networks: social components are human individuals, whereas technological (or cyber) components are devices or platforms endowed with sensing, computation, and communication capabilities. In these systems, the interactions between the individuals are mediated and determined by the ubiquitous presence of digital technology and artificial intelligence algorithms. Furthermore, artificial entities like social bots interact with humans [110] and are becoming more and more sophisticated and effective with the recent advances of generative artificial intelligence.

On online networking platforms such as Twitter, Facebook and the like, it is plain that users do not interact in person, but via the platform: *interactions are mediated by the technological interface*. The platform enables interactions, but also limits and orients them, in at least three fundamental ways. Firstly, many interactions are *stereotyped* because users are allowed a limited set of possible reactions to express their thoughts and emotions (e.g. re-tweeting or not, liking or not). These limitations naturally require dynamical models where continuous variables (defining the individual states) and discrete variables (defining the communication) interplay [271, 62]. Secondly, the platform provides the user with *live updates* about her/his impact on her/his friends and followers: feeding back this information alters the user's behavior. Thirdly, platforms use *recommendation systems* to filter and select which contents should be brought to the users' attention, although users themselves are often unaware of this filtering [103]. These live updates and recommendations introduce nonlinearities and feedback loops in the dynamics, such as I studied in [249] and Section 7.3.

Online social services easily record behaviors and interactions, producing large data streams that enhance our understanding of social dynamics [221]. Beyond the analysis power, these tools also offer new opportunities to influence (sometimes inadvertently) the behaviors of the individuals. This influence can be obtained in various ways, including advertising, diffusing sensitive information, or altering the way individuals interact. A remarkable example of this potential was the controversial Facebook experiment [163], which induced an emotional contagion in thousands of unaware users by tweaking their News Feeds so to emphasize displays of sadness.

Understanding these dynamics and how to manage them is thus not only a scientific challenge, but also an urgent need for the society. Exemplary and well-known issues in this context are the formation of *filter bubbles* [216] that hinder the transmission of relevant information and create self-confirmatory echo-chambers; the viral diffusion of *fake news* [153] and *junk news* [275] that distort and clutter public discourse; and the ever-accelerating *pace of collective attention* [174], which jeopardizes a meaningful public debate. These interrelated phenomena, which I sometimes refer to as the *disorders of online media*, are a byproduct of the very nature of social media, which make interactions highly dynamical and introduce unprecedented effects of feedback and scale [255].

In the last few years I have set up a research effort to contribute to understanding and mitigating these phenomena. This research builds upon my interest in opinion dynamics (which I discussed in Section 6.3) but has broader scope that includes the dynamics of content production, collective attention and popularity in social media. While I was elaborating my model of recommendation system [249], I established interdisciplinary collaborations with social and computer scientists, whose expertise includes computa-




Figure 8.1: Scatterplot of the average Gini index of popularity and average life cycle of the videos of 60 YouTube channels highly active in the French political news debate. The size of the dots corresponds to the total number of views collected by the channel and the color corresponds to the number of published videos (going from dark red for fewer videos, to orange, yellow, green and violet for more videos).

tional social sciences and sociology of new media.

In this context, the PhD thesis of Maria Castaldo, which also included a substantial empirical work of data collection and analysis, has been for me an invaluable opportunity to get a hands-on expertise about social media and, in particular, YouTube. YouTube is indeed a prominent, and still raising, platform in the social media landscape, all the while having been relatively neglected in past research. The analysis of YouTube data has allowed us to uncover fascinating insights into the circadian rhythms of online content consumption, as well as their alterations during the recent COVID-19 pandemics [53]. We have also shed light on a lesser known phenomenon on YouTube, the frequent corrections that the platform performs to view counters in order to combat *fake views* that are due to bots or click farms [51].

From the modelling point of view, during her thesis we have developed a model that describes the competition for the attention of users that takes place on social media [52]. This mathematical model, which is a direct translation in mathematical terms of the sociological model of [152], is able to qualitatively reproduce phenomena of popularity that we observe among YouTube channels. For instance, Figure 8.1 shows that faster dynamics of attention correlate with more unequal popularity of videos.

8.3 Information loops in mobility: the case of navigation apps

In the past decade, there has been a significant increase in the use of navigation apps. This is mainly due to two factors. First, the increase in traffic demand, to which transportation networks are subjected as a result of increasing human socio-economic activities, causes congestion on major roads, leading users of transportation networks to consider using alternative routes to reduce travel time and avoid traffic jams. Second, the increasing number of vehicles equipped with navigation systems and the growing usage of dedicated smartphone applications, such as Waze, Google Maps, and the likes.




Figure 8.2: Control systems perspective: Using *real-time information* implies that navigation apps are introducing a feedback loop on the traffic system.

Navigation apps provide directions for shortest travel time routes based on real-time data. This fact makes apparent their potential to dynamically influence the route decisions made by users of the transportation network. In turn, this observation yields questions on the impact exerted by the use of such applications on road traffic and makes clear the need to understand the underlying feedback mechanisms (Figure 8.2). The relevance of such questions is strengthened by strong empirical evidence showing, over the years, significant changes in traffic patterns and network inefficiencies, which are explained by the new routing habits of app-informed users [39].

The recent scientific literature has already tried to study the effects of navigation apps. The relevant body of research comprises empirical work [39], large-scale simulations [39, 38], and mathematical models [23, 26]. In general, these works focus their attention on two main aspects. First, they highlight how navigation apps can in some cases be responsible for the degradation of the efficiency of the transportation network [39]. This comes from the fact that providing users with shortest route recommendations leads to traffic patterns where no user has an incentive to change route, known as *Wardrop equilibria* or *user equilibria*, which are typically suboptimal with respect to total travel time [289]. Second, they show that the presence of app-informed users tends to spread traffic over the network. Unfortunately, this can induce just a slight alleviation of the pressure on main routes of the network, at the price of a disproportionate increase of congestion on secondary ones, not designed to support the demand they are now subject to [39]. These drawbacks are exacerbated when the fraction of app-informed users increases.

All the mentioned works propose traffic assignment models, where the transportation network is represented by means of a network equipped with origin-destination pairs [23, 161]: users flow through network links from an origin to a destination, choosing one of the possible routes, according to suitable criteria, e.g., travel time. The authors of [39, 263] propose a static model, i.e., in which traffic flows are statically allocated. Static traffic assignment assumes the system to be at equilibrium and is not able to describe the evolution of traffic over time. On the contrary, [111] provides a dynamic model describing the evolution through time of the state of traffic over the network by means of a system of ordinary differential equations (ODEs). However, the complexity of the model prevents the rigorous analysis of its properties. In this work, we aim to go beyond this state of the art and develop a traffic assignment model that is both dynamical in nature and amenable to a full rigorous analysis, so that its equilibria are studied in detail after their stability is established.

The main contribution of the work of my student Tommaso Toso is to provide analytical




Figure 8.3: Illustration of the Grenoble case study. Left: Google Maps caption showing two alternative routes in the Greater Grenoble urban area. Right: Corresponding geometry of roads in the Aimsun simulation environment.

proof that navigation apps can have detrimental effects on traffic. This is done by analysing a novel dynamic traffic assignment model in which a system of ODEs, one for each route of the network, is defined. The considered network geometry consists in an origin-destination pair with two alternative routes: a fraction of the drivers can resort to a navigation system to choose which route to pick. The traffic dynamics on network routes is described by means of a supply and demand mechanism, inspired by [85, 199]: supply and demand functions are very effective to realistically describe the traffic dynamics, also in presence of congestion [76, 177].

In the proposed model, the influence of route recommendations on the demand splitting is captured by state-dependent routing ratios, which define how the traffic demand splits on the two routes. Routing ratios are functions of the route densities that satisfy a natural *monotonicity property*, i.e., the demand directed towards a route decreases if its travel time increases. This property enables a thorough stability analysis, showing that the model admits a globally asymptotically stable equilibrium point. The impact of navigation apps on traffic can thus be studied by studying the properties of the steady-state equilibrium. Our results shown that, at equilibrium, routing recommendations can have negative effects: the network efficiency, measured as a total travel time, can worsen and unsatisfied demand can emerge. Consistently with the literature, we find that the drawbacks of navigation apps become more likely for high penetration rates [39].

The dynamic model we propose is the first featuring state-dependent monotone routing ratios, otherwise studied in [23, 73, 74, 75], together with a supply and demand mechanism [76, 177]. This work-in-progress is presented in a series of papers and preprints [268, 269, 266]. We are currently working towards an extensive analysis of the Grenoble case study shown in Figure 8.3. The analysis is based on a detailed simulation performed on the Aimsun commercial micro-simulator.

Bibliography

- [1] R. P. Abelson. Mathematical models in social psychology. *Advances in experimental social psychology*, 3:1–54, 1967.
- [2] F. Acciani. *Control over unreliable networks: Consensus and platooning*. PhD thesis, University of Twente, 2022.
- [3] F. Acciani, P. Frasca, G. Heijenk, and A. A. Stoorvogel. Achieving robust average consensus over lossy wireless networks. *IEEE Transactions on Control of Network Systems*, 6(1):127–137, March 2019.
- [4] F. Acciani, P. Frasca, G. Heijenk, and A. A. Stoorvogel. Stochastic string stability of vehicle platoons via cooperative adaptive cruise control with lossy communication. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):10912–10922, 2021.
- [5] F. Acciani, P. Frasca, A. Stoorvogel, E. Semsar-Kazerooni, and G. Heijenk. Cooperative adaptive cruise control over unreliable networks: an observer-based approach to increase robustness to packet loss. In *European Control Conference*, pages 1399–1404, Limassol, Cyprus, June 2018.
- [6] D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar. Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research*, 38(1):1–27, 2013.
- [7] G. Aletti and G. Naldi. Opinion dynamics on graphon: The piecewise constant case. *Applied Mathematics Letters*, 133:108227, 2022.
- [8] C. Altafini. Consensus problems on networks with antagonistic interactions. *IEEE Transactions on Automatic Control*, 58(4):935–946, 2013.
- [9] C. Altafini and F. Ceragioli. Signed bounded confidence models for opinion dynamics. *Automatica*, 93:114–125, 2018.
- [10] A. Ammar and D. Shah. Efficient rank aggregation using partial data. *ACM SIGMETRICS Performance Evaluation Review*, 40(1):355–366, 2012.
- [11] J. P. Aubin and A. Cellina. *Differential inclusions*, volume 264 of *Grundlehren der Mathematischen Wissenschaften*. Springer, Berlin, 1984.
- [12] M. Avella-Medina, F. Parise, M. T. Schaub, and S. Segarra. Centrality measures for graphons: Accounting for uncertainty in networks. *IEEE Transactions on Network Science and Engineering*, 7(1):520–537, 2020.
- [13] A. Aydoğdu, M. Caponigro, S. McQuade, B. Piccoli, N. Pouradier Duteil, F. Rossi, and E. Trélat. Interaction network, state space, and control in social dynamics. In N. Bellomo, P. Degond, and E. Tadmor, editors, *Active Particles, Volume 1: Advances in Theory, Models, and Applications*, pages 99–140. Springer International Publishing, 2017.
- [14] A. Aydoğdu, P. Frasca, C. D’Apice, R. Manzo, J. M. Thornton, B. Gachomo, T. Wilson, B. Cheung, U. Tariq, W. Saidel, and B. Piccoli. Modeling birds on wires. *Journal of Theoretical Biology*, 415:102–112, 2017.
- [15] N. Ayi and N. Pouradier Duteil. Mean-field and graph limits for collective dynamics models with time-varying weights. *Journal of Differential Equations*, 299:65–110, 2021.

- [16] A. Bacciotti. Some remarks on generalized solutions of discontinuous differential equations. *International Journal of Pure and Applied Mathematics*, 10(3):257–266, 2003.
- [17] S. Baldi and P. Frasca. Leaderless synchronization of heterogeneous oscillators by adaptively learning the group model. *IEEE Transactions on Automatic Control*, 65(1):412–418, 2019.
- [18] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *PNAS*, 105(4):1232–1237, 2008.
- [19] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Empirical investigation of starling flocks: a benchmark study in collective animal behaviour. *Animal Behaviour*, 76:201–215, July 2008.
- [20] B. Bamieh, M. Jovanovic, P. Mitra, and S. Patterson. Coherence in large-scale networks: Dimension-dependent limitations of local feedback. *IEEE Transactions on Automatic Control*, 57(9):2235–2249, 2012.
- [21] S. Banisch and E. Olbrich. Opinion polarization by learning from social feedback. 2017.
- [22] P. Barooah and J. P. Hespanha. Estimation from relative measurements: Algorithms and scaling laws. *IEEE Control Systems Magazine*, 27(4):57–74, 2007.
- [23] A. M. Bayen, A. Keimer, E. Porter, and M. Spinola. Time-continuous instantaneous and past memory routing on traffic networks: A mathematical analysis on the basis of the link-delay model. *SIAM Journal on Applied Dynamical Systems*, 18(4):2143–2180, 2019.
- [24] C. Bernardo, C. Altafini, A. Proskurnikov, and F. Vasca. Bounded confidence opinion dynamics: A survey. *Automatica*, 159:111302, 2024.
- [25] G. Bianchin, A. Gasparri, P. Frasca, and F. Pasqualetti. The observability radius of networks. *IEEE Transactions on Automatic Control*, 62(6):3006–3013, 2017.
- [26] G. Bianchin and F. Pasqualetti. Routing apps may cause oscillatory congestions in traffic networks. In *IEEE 59th Conference on Decision and Control (CDC)*, pages 253–260, Jeju Island, Korea, 2020.
- [27] M. Bini, P. Frasca, C. Ravazzi, and F. Dabbene. Graph structure-based heuristics for optimal targeting in social networks. *IEEE Transactions on Control of Network Systems*, 9(3):1189–1201, 2022.
- [28] V. Blondel, J. Hendrickx, and J. Tsitsiklis. Continuous-time average-preserving opinion dynamics with opinion-dependent communications. *SIAM Journal on Control and Optimization*, 48(8):5214–5240, 2010.
- [29] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis. On Krause’s multi-agent consensus model with state-dependent connectivity. *IEEE Transactions on Automatic Control*, 54(11):2586–2597, 2009.
- [30] B. Bonnet, N. Pouradier Duteil, and M. Sigalotti. Consensus formation in first-order graphon models with time-varying topologies. *Mathematical Models and Methods in Applied Sciences*, 32(11):2121–2188, 2022.
- [31] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztegombi. Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. *Advances in Mathematics*, 219(6):1801–1851, 2008.
- [32] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztegombi. Convergent sequences of dense graphs II. Multiway cuts and statistical physics. *Annals of Mathematics*, 176(1):151–219, 2012.
- [33] V. S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, 2008.

- [34] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.
- [35] F. Bullo. *Lectures on Network Systems*. Kindle Direct Publishing, 1.7 edition, 2024.
- [36] F. Bullo, R. Carli, and P. Frasca. Gossip coverage control for robotic networks: Dynamical systems on the space of partitions. *SIAM Journal on Control and Optimization*, 50(1):419–447, 2012.
- [37] F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009.
- [38] T. Cabannes, F. Shyu, E. Porter, S. Yao, Y. Wang, M. A. Sangiovanni Vincentelli, S. Hinardi, M. Zhao, and A. M. Bayen. Measuring regret in routing: Assessing the impact of increased app usage. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2589–2594, 2018.
- [39] T. C. P. Cabannes. *The impact of information-aware routing on road traffic. From case studies to game-theoretical analysis and simulations*. PhD thesis, University of California, Berkeley, 2022.
- [40] P. E. Caines and M. Huang. Graphon mean field games and the GMFG equations. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4129–4134. IEEE, 2018.
- [41] C. Canuto, F. Fagnani, and P. Tilli. An Eulerian approach to the analysis of Krause’s consensus models. *SIAM Journal on Control and Optimization*, 50(1):243–265, 2012.
- [42] M. Cao, B. D. Anderson, and A. S. Morse. Sensor network localization with imprecise distances. *Systems & control letters*, 55(11):887–893, 2006.
- [43] T. Carletti, D. Fanelli, A. Guarino, F. Bagnoli, and A. Guazzini. Birth and death in a continuous opinion dynamics model: The consensus case. *The European Physical Journal B*, 64:285–292, 2008.
- [44] R. Carli, G. Como, P. Frasca, and F. Garin. Distributed averaging on digital erasure networks. *Automatica*, 47(1):115–121, 2011.
- [45] R. Carli, G. Como, P. Frasca, and F. Garin. Distributed averaging on digital noisy networks. In *Information Theory and Applications Workshop*, pages 1–9, San Diego, CA, Feb. 2011.
- [46] R. Carli, F. Fagnani, P. Frasca, T. Taylor, and S. Zampieri. Average consensus on networks with transmission noise or quantization. In *European Control Conference*, pages 1852–1857, Kos, Greece, June 2007.
- [47] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Gossip consensus algorithms via quantized communication. *Automatica*, 46(1):70–80, 2010.
- [48] R. Carli, F. Fagnani, A. Speranzon, and S. Zampieri. Communication constraints in the average consensus problem. *Automatica*, 44(3):671–684, 2008.
- [49] M. Castaldo. *Attention dynamics on YouTube: conceptual models, temporal analysis of engagement metrics, fake views*. PhD thesis, Université Grenoble Alpes, 2022.
- [50] M. Castaldo, P. Frasca, and T. Venturini. Online attention dynamics in social media. In *Cyber-Physical-Human Systems: Fundamentals and Applications*, pages 491–510. Wiley Online Library, 2023.
- [51] M. Castaldo, P. Frasca, T. Venturini, and F. Gargiulo. Doing data science with platforms crumbs: an investigation into fake views on YouTube. *arXiv preprint arXiv:2210.01096*, 2022.
- [52] M. Castaldo, T. Venturini, P. Frasca, and F. Gargiulo. Junk news bubbles modelling the rise and fall of attention in online arenas. *New Media & Society*, page 1461444820978640, 2020.

- [53] M. Castaldo, T. Venturini, P. Frasca, and F. Gargiulo. The rhythms of the night: increase in online night activity and emotional resilience during the spring 2020 covid-19 lockdown. *EPJ Data Science*, 10(1):7, 2021.
- [54] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, 2009.
- [55] D. Castro, F. Ruffier, and C. Eloy. Modeling collective behaviors from optic flow and retinal cues. *Physical Review Research*, 6(2):023016, 2024.
- [56] S. Caton and C. Haas. Fairness in machine learning: A survey. *ACM Computing Surveys*, 56(7):1–38, 2024.
- [57] F. Ceragioli. *Discontinuous ordinary differential equations and stabilization*. PhD thesis, Università di Firenze, 2000.
- [58] F. Ceragioli, C. De Persis, and P. Frasca. Quantized average consensus: Discontinuities and hysteresis. In *IFAC Symposium on Nonlinear Control Systems*, pages 78–83, Bologna, Italy, Sept. 2010. Invited paper.
- [59] F. Ceragioli, C. De Persis, and P. Frasca. Discontinuities and hysteresis in quantized average consensus. *Automatica*, 47(9):1916–1928, 2011.
- [60] F. Ceragioli and P. Frasca. Continuous and discontinuous opinion dynamics with bounded confidence. *Nonlinear Analysis: Real World Applications*, 13(3):1239–1251, 2012.
- [61] F. Ceragioli and P. Frasca. Continuous-time consensus dynamics with quantized all-to-all communication. In *European Control Conference*, pages 1120–1125, Linz, Austria, July 2015.
- [62] F. Ceragioli and P. Frasca. Consensus and disagreement: The role of quantized behaviors in opinion dynamics. *SIAM Journal on Control and Optimization*, 56(2):1058–1080, 2018.
- [63] F. Ceragioli and P. Frasca. Discontinuities, generalized solutions, and (dis) agreement in opinion dynamics. In *Control Subject to Computational and Communication Constraints: Current Challenges*, pages 287–309. Springer, 2018.
- [64] F. Ceragioli, P. Frasca, B. Piccoli, and F. Rossi. Generalized solutions to opinion dynamics models with discontinuities. In N. Bellomo and L. Gibelli, editors, *Crowd Dynamics Vol. 3*. Birkhauser, 2021.
- [65] F. Ceragioli, P. Frasca, and W. S. Rossi. Modeling limited attention in opinion dynamics by topological interactions. In *International Conference on Network Games, Control and Optimization*, pages 272–281. Springer, 2021.
- [66] B. Chazelle and C. Wang. Inertial Hegselmann-Krause systems. *IEEE Transactions on Automatic Control*, 62(8):3905–3913, 2016.
- [67] N. R. Chowdhury, I.-C. Morarescu, S. Martin, and S. Srikant. Continuous opinions and discrete actions in social networks: a multi-agent system approach. In *IEEE Conference on Decision and Control*, pages 1739–1744, Las Vegas, NV, USA, Dec. 2016.
- [68] C. Commault, J. van der Woude, and P. Frasca. Functional target controllability of networks: Structural properties and efficient algorithms. *IEEE Transactions on Network Science and Engineering*, 7(3):1521–1530, 2019.
- [69] G. Como and F. Fagnani. From local averaging to emergent global behaviors: The fundamental role of network interconnections. *Systems & Control Letters*, 95:70–76, 2016.
- [70] G. Como, P. Frasca, Y. Mo, F. Parise, C. Ravazzi, and S. Yüksel. Technical committee on networks and communication systems. *IEEE Control Systems Magazine*, 41(6):21–23, 2021.
- [71] G. Como, P. Frasca, F. Parise, C. Ravazzi, K. Savla, and S. Yüksel. Networks and communication systems [technical activities]. *IEEE Control Systems Magazine*, 43(4):20–22, 2023.

- [72] G. Como, P. Frasca, and S. Yuksel. Technical committee on networks and communication systems [technical activities]. *IEEE Control Systems Magazine*, 40(2):19–22, 2020.
- [73] G. Como, E. Lovisari, and K. Savla. Throughput optimality and overload behavior of dynamical flow networks under monotone distributed routing. *IEEE Transactions on Control of Network Systems*, 2(1):57–67, 2015.
- [74] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli. Robust distributed routing in dynamical networks — Part I: Locally responsive policies and weak resilience. *IEEE Transactions on Automatic Control*, 58(2):317–332, 2013.
- [75] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli. Stability analysis of transportation networks with multiscale driver decisions. *SIAM Journal on Control and Optimization*, 51:230–252, 2013.
- [76] S. Coogan and M. Arcak. A compartmental model for traffic networks and its dynamical behavior. *IEEE Transactions on Automatic Control*, 60(10):2698–2703, 2015.
- [77] M. Coraggio, P. De Lellis, and M. di Bernardo. Convergence and synchronization in networks of piecewise-smooth systems via distributed discontinuous coupling. *Automatica*, 129:109596, 2021.
- [78] M. Coraggio, P. DeLellis, and M. di Bernardo. Distributed discontinuous coupling for convergence in heterogeneous networks. *IEEE Control Systems Letters*, 5(3):1037–1042, 2021.
- [79] J. Cortés. Discontinuous dynamical systems – a tutorial on solutions, nonsmooth analysis, and stability. *IEEE Control Systems Magazine*, 28(3):36–73, 2008.
- [80] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.
- [81] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [82] E. Cristiani, P. Frasca, and B. Piccoli. Effects of anisotropic interactions on the structure of animal groups. *Journal of Mathematical Biology*, 62(4):569–588, 2011.
- [83] C. Cros, P.-O. Amblard, C. Prieur, and J.-F. Da Rocha. Fusion of distance measurements between agents with unknown correlations. *IEEE Control Systems Letters*, 7:2143–2148, 2023.
- [84] R. Curtain, O. Iftime, and H. Zwart. A comparison between LQR control for a long string of SISO systems and LQR control of the infinite spatially invariant version. *Automatica*, 46(10):1604–1615, 2010.
- [85] C. F. Daganzo. The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory. *Transportation Research Part B: Methodological*, 28:269–287, 1994.
- [86] C. De Persis and P. Frasca. Robust self-triggered coordination with ternary controllers. *IEEE Transactions on Automatic Control*, 58(12):3024–3038, 2013.
- [87] C. De Persis and P. Frasca. Robust self-triggered coordination with ternary controllers. *IEEE Transactions on Automatic Control*, 58(12):3024–3038, 2013.
- [88] M. de Roo, P. Frasca, and R. Carloni. Optimal event handling by multiple unmanned aerial vehicles. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1230–1236, 2016.
- [89] S. Dean and J. Morgenstern. Preference dynamics under personalized recommendations. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 795–816, 2022.
- [90] R. Defayot, T. Thonet, J.-M. Renders, and M. de Rijke. Generative slate recommendation with reinforcement learning. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 580–588, 2023.

- [91] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(1-4):87–98, 2000.
- [92] J.-F. Delmas, D. Dronnier, and P.-A. Zitt. An infinite-dimensional metapopulation SIS model. *Journal of Differential Equations*, 313:1–53, 2022.
- [93] J.-F. Delmas, P. Frasca, F. Garin, V. C. Tran, A. Velleret, and P.-A. Zitt. Individual based sis models on (not so) dense large random networks. *arXiv preprint arXiv:2302.13385*, 2023.
- [94] F. Dietrich, S. Martin, and M. Jungers. Control via leadership of opinion dynamics with state and time-dependent interactions. *IEEE Transactions on Automatic Control*, 63(4):1200–1207, 2017.
- [95] A. G. Dimakis, S. Kar, J. F. Moura, M. G. Rabbat, and A. Scaglione. Gossip algorithms for distributed signal processing. *Proceedings of the IEEE*, 98(11):1847–1864, 2010.
- [96] J.-M. Dion, C. Commault, and J. van der Woude. Generic properties and control of linear structured systems: a survey. *Automatica*, 39(7):1125–1144, 2003.
- [97] A. Dubovskaya, S. C. Fennell, K. Burke, J. P. Gleeson, and D. O’Kiely. Analysis of mean-field approximation for deffuant opinion dynamics on networks. *SIAM Journal on Applied Mathematics*, 83(2):436–459, 2023.
- [98] R. Dunbar. Neocortex size as a constraint on group size in primates. *J. Human Evolution*, 22(6):469–493, 1992.
- [99] J. W. Durham, R. Carli, P. Frasca, and F. Bullo. Dynamic coverage control with asynchronous one-to-base-station communication. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 5589–5594, Orlando, FL, USA, Dec. 2011.
- [100] J. W. Durham, R. Carli, P. Frasca, and F. Bullo. Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics*, 28(2):364–378, 2012.
- [101] R. M. D’Souza, M. di Bernardo, and Y.-Y. Liu. Controlling complex networks with complex nodes. *Nature Reviews Physics*, 5(4):250–262, 2023.
- [102] M. Egerstedt, S. Martini, M. Cao, K. Camlibel, and A. Bicchi. Interacting with networks: How does structure relate to controllability in single-leader, consensus networks? *IEEE Control Systems Magazine*, 32(4):66–73, 2012.
- [103] M. Eslami, A. Rickman, K. Vaccaro, A. Aleyasen, A. Vuong, K. Karahalios, K. Hamilton, and C. Sandvig. "I always assumed that I wasn’t really that close to [her]": Reasoning about invisible algorithms in news feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 153–162, New York, NY, USA, 2015.
- [104] F. Fagnani and P. Frasca. Broadcast gossip averaging: interference and unbiasedness in large Abelian Cayley networks. *IEEE Journal of Selected Topics in Signal Processing*, 5(4):866–875, 2011.
- [105] F. Fagnani and P. Frasca. *Introduction to Averaging Dynamics over Networks*. Lecture Notes in Control and Information Sciences. Springer Nature, 2017.
- [106] F. Fagnani and S. Zampieri. Asymmetric randomized gossip algorithms for consensus. In *IFAC World Congress*, pages 9052–9056, Seoul, Korea, July 2008.
- [107] F. Fagnani and S. Zampieri. Randomized consensus algorithms over large scale networks. *IEEE Journal on Selected Areas in Communications*, 26(4):634–649, 2008.
- [108] F. Fagnani and S. Zampieri. Average consensus with packet drop communication. *SIAM Journal on Control and Optimization*, 48(1):102–133, 2009.
- [109] A. Feintuch and B. Francis. Infinite chains of kinematic points. *Automatica*, 48(5):901 – 908, 2012.
- [110] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini. The rise of social bots. *Communications of the ACM*, 59(7):96–104, 2016.

- [111] A. Festa and P. Goatin. Modeling the impact of on-line navigation devices in traffic flows. In *IEEE 58th Conference on Decision and Control (CDC)*, pages 323–328, 2019.
- [112] A. F. Filippov. *Differential Equations with Discontinuous Righthand Sides*, volume 18 of *Mathematics and Its Applications*. Kluwer Academic Publishers, 1988.
- [113] A. Firooznia, J. Ploeg, N. Van De Wouw, and H. Zwart. Co-design of controller and communication topology for vehicular platooning. *IEEE Transactions on Intelligent Transportation Systems*, 18(10):2728–2739, 2017.
- [114] M. Franceschelli and P. Frasca. Proportional dynamic consensus in open multi-agent systems. In *IEEE Conference on Decision and Control*, pages 900–905, Dec. 2018.
- [115] M. Franceschelli and P. Frasca. Stability of open multiagent systems and applications to dynamic consensus. *IEEE Transactions on Automatic Control*, 66(5):2326–2331, 2021.
- [116] P. Frasca. Continuous-time quantized consensus: Convergence of Krasovskii solutions. *Systems & Control Letters*, 61(2):273–278, 2012.
- [117] P. Frasca, R. Carli, F. Fagnani, and S. Zampieri. Average consensus on networks with quantized communication. *International Journal of Robust and Nonlinear Control*, 19(16):1787–1816, 2009.
- [118] P. Frasca, F. Garin, B. Gerencser, and J. M. Hendrickx. Optimal one-dimensional coverage by unreliable sensors. *SIAM Journal on Control and Optimization*, 53(5):3120–3140, 2015.
- [119] P. Frasca, F. Garin, and R. Prisant. Opinion dynamics on signed graphs and graphons: Beyond the piece-wise constant case. *arXiv preprint arXiv:2404.08372*, 2024.
- [120] P. Frasca and J. M. Hendrickx. Randomized averaging algorithm: decoupling accuracy and convergence rate. In *Benelux Meeting on Systems and Control*, Heijen, Netherlands, Mar. 2012.
- [121] P. Frasca and J. M. Hendrickx. Large network consensus is robust to packet losses and interferences. In *European Control Conference*, pages 1782–1787, Zürich, Switzerland, July 2013.
- [122] P. Frasca and J. M. Hendrickx. On the mean square error of randomized averaging algorithms. *Automatica*, 49(8):2496–2501, 2013.
- [123] P. Frasca, H. Ishii, C. Ravazzi, and R. Tempo. Distributed randomized algorithms for opinion formation, centrality computation and power systems estimation. *European Journal of Control*, 24:2–13, July 2015.
- [124] P. Frasca, C. Ravazzi, R. Tempo, and H. Ishii. Gossips and prejudices: Ergodic randomized dynamics in social networks. In *IFAC Workshop on Estimation and Control of Networked Systems*, pages 212–219, Koblenz, Germany, Sept. 2013.
- [125] P. Frasca and F. Rossi. Caratheodory solutions and their associated graphs in opinion dynamics with topological interactions. *IFAC-PapersOnLine*, 55(30):436–441, 2022.
- [126] P. Frasca, S. Tarbouriech, and L. Zaccarian. A hybrid model of opinion dynamics with limited confidence. In *IFAC Symposium on Nonlinear Control Systems*, pages 351–355, Monterey, CA, USA, Aug. 2016.
- [127] P. Frasca, S. Tarbouriech, and L. Zaccarian. Hybrid models of opinion dynamics with opinion-dependent connectivity. *Automatica*, 100:153–161, 2019.
- [128] J. R. P. French. A formal theory of social power. *Psychological Review*, 63:181–94, 1956.
- [129] N. E. Friedkin. *A structural theory of social influence*. Cambridge University Press, 2006.
- [130] N. E. Friedkin. The problem of social control and coordination of complex systems in sociology: A look at the community cleavage problem. *IEEE Control Systems*, 35(3):40–51, 2015.

- [131] N. E. Friedkin and E. C. Johnsen. Social influence networks and opinion change. In E. J. Lawler and M. W. Macy, editors, *Advances in Group Processes*, volume 16, pages 1–29. JAI Press, 1999.
- [132] N. E. Friedkin and E. C. Johnsen. *Social Influence Network Theory: A Sociological Examination of Small Group Dynamics*. Cambridge University Press, 2011.
- [133] S. Galam. Contrarian deterministic effects on opinion dynamics: “the hung elections scenario”. *Physica A: Statistical Mechanics and its Applications*, 333(Supplement C):453–460, 2004.
- [134] S. Galam. *Sociophysics: a physicist’s modeling of psycho-political phenomena*. Springer Science & Business Media, 2012.
- [135] S. Gao and P. Caines. Graphon control of large-scale networks of linear systems. *IEEE Transactions on Automatic Control*, 65(10):4090–4105, 2019.
- [136] S. Gao and P. E. Caines. The control of arbitrary size networks of linear systems via graphon limits: An initial investigation. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1052–1057. IEEE, 2017.
- [137] F. Gargiulo, M. Castaldo, T. Venturini, and P. Frasca. Distribution of labor, productivity and innovation in collaborative science. *Applied Network Science*, 7(1):19, 2022.
- [138] Y. Ge, S. Zhao, H. Zhou, C. Pei, F. Sun, W. Ou, and Y. Zhang. Understanding echo chambers in e-commerce recommender systems. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2261–2270, 2020.
- [139] V. Giammarino, S. Baldi, P. Frasca, and M. L. Delle Monache. Traffic flow on a ring with a single autonomous vehicle: An interconnected stability perspective. *IEEE Transactions on Intelligent Transportation Systems*, 22(8):4998–5008, 2020.
- [140] C. M. Gisolo, M. L. Delle Monache, F. Ferrante, and P. Frasca. Nonlinear analysis of stability and safety of optimal velocity model vehicle groups on ring roads. *IEEE Transactions on Intelligent Transportation Systems*, 23(11):20628–20635, 2022.
- [141] S. Gracy, F. Garin, and A. Kibangou. Structural and strongly structural input and state observability of linear network systems. *IEEE Transactions on Control of Network Systems*, 2018.
- [142] S. Grauwin and P. Jensen. Opinion group formation and dynamics: Structures that last from nonlasting entities. *Physical Review E*, 85(6):066113, 2012.
- [143] E. Hargreaves, C. Agosti, D. Menasché, G. Neglia, A. Reiffers-Masson, and E. Altman. Biases in the Facebook news feed: a case study on the Italian elections. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 806–812, 2018.
- [144] R. S. V. Haro. *Contributions to open multi-agent systems: consensus, optimization and epidemics*. PhD thesis, Université Paris-Saclay, 2022.
- [145] N. Hayashi. Distributed subgradient method in open multiagent systems. *IEEE Transactions on Automatic Control*, 68(10):6192–6199, 2023.
- [146] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3):1–33, 2002.
- [147] R. Hegselmann and U. Krause. Opinion dynamics under the influence of radical groups, charismatic leaders, and other constant signals: A simple unifying model. *Networks and Heterogeneous Media*, 10(3):477–509, 2015.
- [148] J. Hendrickx, A. Olshevsky, and V. Saligrama. Minimax rate for learning from pairwise comparisons in the btl model. In *International Conference on Machine Learning*, pages 4193–4202. PMLR, 2020.

- [149] J. M. Hendrickx and S. Martin. Open multi-agent systems: Gossiping with deterministic arrivals and departures. In *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1094–1101, 2016.
- [150] J. M. Hendrickx and S. Martin. Open multi-agent systems: Gossiping with random arrivals and departures. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 763–768, 2017.
- [151] J. M. Hendrickx and M. G. Rabbat. Stability of decentralized gradient descent in open multi-agent systems. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 4885–4890, 2020.
- [152] S. Hilgartner and C. L. Bosk. The rise and fall of social problems: A public arenas model. *American journal of Sociology*, 94(1):53–78, 1988.
- [153] M. Hirst. Towards a political economy of fake news. *The Political Economy of Communication*, 5(2), 2017.
- [154] A. Howard, M. J. Mataric, and G. S. Sukhatme. Cooperative relative localization for mobile robot teams: An ego-centric approach. In *Naval Research Lab. Workshop on Multi-Robot Systems*, pages 65–76, 2003.
- [155] Y.-G. Hsieh, F. Iutzeler, J. Malick, and P. Mertikopoulos. Optimization in open networks via dual averaging. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 514–520, 2021.
- [156] T. D. Huynh, N. R. Jennings, and N. R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, Sep 2006.
- [157] P. E. Jabin and S. Motsch. Clustering and asymptotic behavior in opinion formation. *Journal of Differential Equations*, 257(11):4165–4187, 2014.
- [158] M. O. Jackson. *Social and economic networks*. Princeton University Press, 2010.
- [159] R. Jiang, S. Chiappa, T. Lattimore, A. György, and P. Kohli. Degenerate feedback loops in recommender systems. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 383–390, 2019.
- [160] X. Jiang, L.-H. Lim, Y. Yao, and Y. Ye. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.
- [161] A. Keimer and A. M. Bayen. Routing on traffic networks incorporating past memory up to real-time information on the network state. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):151–172, 2019.
- [162] S. S. Kia, B. Van Scoy, J. Cortes, R. A. Freeman, K. M. Lynch, and S. Martinez. Tutorial on dynamic average consensus: The problem, its applications, and the algorithms. *IEEE Control Systems Magazine*, 39(3):40–72, 2019.
- [163] A. D. I. Kramer, J. E. Guillory, and J. T. Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *PNAS*, 111(24):8788–8790, 2014.
- [164] J. Krause and G. D. Ruxton. *Living in groups*. Oxford Series in Ecology and Evolution. Oxford University Press, 2002.
- [165] C. Kuehn. *Moment Closure—A Brief Review*, pages 253–271. Springer International Publishing, Cham, 2016.
- [166] C. Kuehn and S. Throm. Power network dynamics on graphons. *SIAM Journal on Applied Mathematics*, 79(4):1271–1292, 2019.
- [167] S. L. Lim and P. J. Bentley. Opinion amplification causes extreme polarization in social networks. *Scientific Reports*, 12(1):18131, 2022.
- [168] C.-T. Lin. Structural controllability. *IEEE Transactions on Automatic Control*, 19(3):201–208, 1974.

- [169] G. Lindmark and C. Altafini. Minimum energy control for complex networks. *Scientific reports*, 8(1):3188, 2018.
- [170] D. Liu, S. Baldi, V. Jain, W. Yu, and P. Frasca. Cyclic communication in adaptive strategies to platooning: The case of synchronized merging. *IEEE Transactions on Intelligent Vehicles*, 6(3):490–500, 2020.
- [171] D. Liu, S. Mair, K. Yang, S. Baldi, P. Frasca, and M. Althoff. Resilience in platoons of cooperative heterogeneous vehicles: Self-organization strategies and provably-correct design. *IEEE Transactions on Intelligent Vehicles*, 9(1):2262–2275, 2024.
- [172] L. Liu, S. Dean, E. Rolf, M. Simchowitz, and M. Hardt. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, pages 3150–3158. PMLR, 2018.
- [173] Y. Liu, J. Slotine, and A. Barabasi. Controllability of complex networks. *Nature*, 473(7346), 2011.
- [174] P. Lorenz-Spreen, B. M. Mønsted, P. Hövel, and S. Lehmann. Accelerating dynamics of collective attention. *Nature communications*, 10(1):1759, 2019.
- [175] L. Lovász. *Large networks and graph limits*. American Mathematical Society, 2012.
- [176] L. Lovász and B. Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96(6):933–957, 2006.
- [177] E. Lovisari, G. Como, and K. Savla. Stability of monotone dynamical flow networks. In *53rd IEEE Conference on Decision and Control*, pages 2384–2389, Los Angeles, California, 2014.
- [178] E. Lovisari, F. Garin, and S. Zampieri. Resistance-based performance analysis of the consensus algorithm over geometric graphs. *SIAM Journal on Control and Optimization*, 51(5):3918–3945, 2013.
- [179] T. Lykouris, V. Syrgkanis, and E. Tardos. Learning and efficiency in games with dynamic population. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '16*, pages 120–129, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics.
- [180] G. C. Maffettone, A. Boldini, M. Di Bernardo, and M. Porfiri. Continuification control of large-scale multiagent systems in a ring. *IEEE Control Systems Letters*, 7:841–846, 2022.
- [181] G. C. Maffettone, M. Porfiri, and M. di Bernardo. Continuification control of large-scale multiagent systems under limited sensing and structural perturbations. *IEEE Control Systems Letters*, 2023.
- [182] N. Martin. *Network partitioning algorithms with scale-free objective*. PhD thesis, Université Grenoble Alpes [2020-....], 2020.
- [183] N. Martin, P. Frasca, and C. Canudas-de Wit. Large-scale network reduction towards scale-free structure. *IEEE Transactions on Network Science and Engineering*, 6(4):711–723, 2018.
- [184] N. Martin, P. Frasca, and C. Canudas-De-Wit. Mergetocure: a new strategy to allocate cure in an epidemic over a grid-like network using a scale-free abstraction. *IFAC-PapersOnLine*, 51(23):34–39, 2018.
- [185] N. Martin, P. Frasca, and C. Canudas-de Wit. Subgraph detection for average detectability of lti systems. *IEEE Transactions on Network Science and Engineering*, 2020.
- [186] N. Martin, P. Frasca, T. Ishizaki, J.-I. Imura, and C. Canudas-de Wit. The price of connectedness in graph partitioning problems. In *2019 18th European control conference (ECC)*, pages 2313–2318. IEEE, 2019.
- [187] A. C. Martins. Continuous opinions and discrete actions in opinion dynamics problems. *International Journal of Modern Physics C*, 19(04):617–624, 2008.

- [188] A. C. Martins. Trust in the CODA model: Opinion dynamics and the reliability of other agents. *Physics Letters A*, 377(37):2333–2339, 2013.
- [189] S. Mauras, V. Cohen-Addad, G. Duboc, M. Dupré la Tour, P. Frasca, C. Mathieu, L. Opatowski, and L. Viennot. Mitigating covid-19 outbreaks in workplaces and schools by hybrid telecommuting. *PLoS Computational Biology*, 17(8):e1009264, 2021.
- [190] G. S. Medvedev. The nonlinear heat equation on dense graphs and graph limits. *SIAM Journal on Mathematical Analysis*, 46(4):2743–2766, 2014.
- [191] G. S. Medvedev. The nonlinear heat equation on w-random graphs. *Archive for Rational Mechanics and Analysis*, 212:781–803, 2014.
- [192] A. Mironchenko. Live systems of varying dimension: Modeling and stability. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 3956–3961, 2023.
- [193] A. Mirtabatabaei and F. Bullo. Opinion dynamics in heterogeneous networks: Convergence conjectures and theorems. *SIAM Journal on Control and Optimization*, 50(5):2763–2785, 2012.
- [194] M. Mobilia. Does a single zealot affect an infinite group of voters? *Physical Review Letters*, 91(2):028701, 2003.
- [195] C. Monnoyer de Galland and J. M. Hendrickx. Fundamental performance limitations for average consensus in open multi-agent systems. *IEEE Transactions on Automatic Control*, 68(2):646–659, 2023.
- [196] C. Monnoyer de Galland, S. Martin, and J. M. Hendrickx. Modelling gossip interactions in open multi-agent systems. *arXiv preprint arXiv:2009.02970*, 2020.
- [197] C. Monnoyer de Galland, R. Vizuete, J. M. Hendrickx, P. Frasca, and E. Panteley. Random coordinate descent algorithm for open multi-agent systems with complete topology and homogeneous agents. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 1701–1708. IEEE, 2021.
- [198] C. Monnoyer de Galland, R. Vizuete, J. M. Hendrickx, E. Panteley, and P. Frasca. Random coordinate descent for resource allocation in open multi-agent systems. *IEEE Transactions on Automatic Control*, 2024.
- [199] I.-C. Morărescu and C. Canudas de Wit. Highway traffic model-based density estimation. In *Proceedings of the American Control Conference*, pages 2012–2017, San Francisco, California, 2011.
- [200] S. Motsch and E. Tadmor. Heterophilious dynamics enhances consensus. *SIAM Review*, 56(4):577–621, 2014.
- [201] S. V. Nagpal, G. G. Nair, S. H. Strogatz, and F. Parise. Synchronization in random networks of identical phase oscillators: A graphon approach. *arXiv preprint arXiv:2403.13998*, 2024.
- [202] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis. On distributed averaging algorithms and quantization effects. *IEEE Transactions on Automatic Control*, 54(11):2506–2517, 2009.
- [203] M. U. B. Niazi, C. Canudas-de Wit, and A. Y. Kibangou. Thermal monitoring of buildings by aggregated temperature estimation. *IFAC-PapersOnLine*, 53(2):4132–4137, 2020.
- [204] M. U. B. Niazi, X. Cheng, C. Canudas-de Wit, and J. M. Scherpen. Clustering-based average state observer design for large-scale network systems. *Automatica*, 151:110914, 2023.
- [205] D. Nikitin. *Scalable large-scale control of network aggregates*. PhD thesis, Université Grenoble Alpes [2020-....], 2021.
- [206] D. Nikitin, C. Canudas-de Wit, and P. Frasca. Boundary control for stabilization of large-scale networks through the continuation method. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 4792–4797. IEEE, 2021.

- [207] D. Nikitin, C. Canudas-de Wit, and P. Frasca. Control of average and deviation in large-scale linear networks. *IEEE Transactions on Automatic Control*, 67(4):1639–1654, 2021.
- [208] D. Nikitin, C. Canudas-de Wit, and P. Frasca. A continuation method for large-scale modeling and control: from odes to pde, a round trip. *IEEE Transactions on Automatic Control*, 67(10):5118–5133, 2022.
- [209] D. Nikitin, C. Canudas-de Wit, P. Frasca, and U. Ebels. Synchronization of spin-torque oscillators via continuation method. *IEEE Transactions on Automatic Control*, 68(11):6621–6635, 2023.
- [210] A. Olshevsky. Minimal controllability problems. *IEEE Transactions on Control of Network Systems*, 1(3):249–258, Sept 2014.
- [211] G. Orosz. Connected cruise control: modelling, delay effects, and nonlinear behaviour. *Vehicle System Dynamics*, 54(8):1147–1176, 2016.
- [212] B. Osting, C. Brune, and S. J. Osher. Optimal data collection for improved rankings expose well-connected graphs. *Journal of Machine Learning Research*, 15:2981–3012, 2014.
- [213] N. Pagan, J. Baumann, E. Elokda, G. De Pasquale, S. Bolognani, and A. Hannák. A classification of feedback loops and their relation to biases in automated decision-making systems. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–14, 2023.
- [214] E. Panteley and A. Loria. Synchronization and dynamic consensus of heterogeneous networked systems. *IEEE Transactions on Automatic Control*, 62(8):3758–3773, 2017.
- [215] F. Parise and A. Ozdaglar. Graphon games: A statistical framework for network games and interventions. volume 91, pages 191–225. Wiley Online Library, 2023.
- [216] E. Pariser. *The Filter Bubble: What The Internet Is Hiding From You*. Penguin, 2011.
- [217] S. E. Parsegov, A. V. Proskurnikov, R. Tempo, and N. E. Friedkin. Novel multidimensional models of opinion dynamics in social networks. *IEEE Transactions on Automatic Control*, 62(5):2270–2285, 2017.
- [218] F. Pasqualetti, S. Zampieri, and F. Bullo. Controllability metrics, limitations and algorithms for complex networks. *IEEE Transactions on Control of Network Systems*, 1(1):40–52, 2014.
- [219] R. Patel, P. Frasca, and F. Bullo. Centroidal area-constrained partitioning for robotic networks. *ASME Journal on Dynamic Systems, Measurement, and Control*, 136(3):031024, 2014.
- [220] R. Patel, P. Frasca, J. W. Durham, R. Carli, and F. Bullo. Dynamic partitioning and coverage control with asynchronous one-to-base-station communication. *IEEE Transactions on Control of Network Systems*, 3(1):24–33, March 2016.
- [221] A. Pentland. *Social Physics: How Good Ideas Spread—The Lessons from a New Science*. Penguin, 2014.
- [222] S. Pequito, S. Kar, and A. P. Aguiar. A framework for structural input/output and control configuration selection in large-scale systems. *IEEE Transactions on Automatic Control*, 61(2):303–318, 2016.
- [223] N. Perra and L. E. Rocha. Modelling opinion dynamics in the age of algorithmic personalisation. *Scientific reports*, 9(1):7261, 2019.
- [224] J. Petit, R. Lambiotte, and T. Carletti. Random walks on dense graphs and graphons. *SIAM Journal on Applied Mathematics*, 81(6):2323–2345, 2021.
- [225] I. Pinyol and J. Sabater-Mir. Computational trust and reputation models for open multi-agent systems: a review. *Artificial Intelligence Review*, 40(1):1–25, Jun 2013.

- [226] B. Polyak. *Introduction to Optimization*. Optimization Software, Inc., New York, NY, 1987.
- [227] R. Prisant, L. Cataldo, F. Ceragioli, and P. Frasca. Disagreement, limit cycles and Zeno solutions in continuous opinion dynamics with binary actions. In *2024 European Control Conference (ECC)*, Stockholm, Sweden, June 2024.
- [228] A. V. Proskurnikov, A. S. Matveev, and M. Cao. Opinion dynamics in social networks with hostile camps: Consensus vs. polarization. *IEEE Transactions on Automatic Control*, 61(6):1524–1536, 2016.
- [229] A. V. Proskurnikov and R. Tempo. A tutorial on modeling and analysis of dynamic social networks. Part I. *Annual Reviews in Control*, 43:65–79, 2017.
- [230] A. V. Proskurnikov and R. Tempo. A tutorial on modeling and analysis of dynamic social networks. Part II. *Annual Reviews in Control*, 45:166–190, 2018.
- [231] W. B. Qin, M. M. Gomez, and G. Orosz. Stability and frequency response under stochastic communication delays with applications to connected cruise control design. *IEEE Transactions on Intelligent Transportation Systems*, 18(2):388–403, 2017.
- [232] D. Quevedo, P. Frasca, S. Yuksel, and G. Como. Technical committee on networks and communication systems [technical activities]. *IEEE Control Systems*, 38(4):18–19, 2018.
- [233] D. Quevedo, P. Frasca, D. Zelazo, G. Como, V. Saligrama, and Y. Mo. Technical committee on networks and communication systems [technical activities]. *IEEE Control Systems Magazine*, 36(4):13–15, 2016.
- [234] C. Ravazzi, N. Chan, and P. Frasca. Distributed estimation from relative measurements of heterogeneous and uncertain quality. *IEEE Transactions on Signal and Information Processing over Networks*, 5(2):203–217, 2019.
- [235] C. Ravazzi, P. Frasca, H. Ishii, and R. Tempo. A distributed randomized algorithm for relative localization in sensor networks. In *European Control Conference*, pages 1776–1781, Zürich, Switzerland, July 2013.
- [236] C. Ravazzi, P. Frasca, R. Tempo, and H. Ishii. Almost sure convergence of a randomized algorithm for relative localization in sensor networks. In *IEEE Conference on Decision and Control*, Florence, Italy, Dec. 2013. to appear.
- [237] C. Ravazzi, P. Frasca, R. Tempo, and H. Ishii. Ergodic randomized algorithms and dynamics over networks. *IEEE Transactions on Control of Network Systems*, 2(1):78–87, 2015.
- [238] M. Rodriguez-Vega, C. Canudas-de Wit, and H. Fourati. Average density estimation for urban traffic networks: Application to the grenoble network. *Transportation Research Part B: Methodological*, 154:21–43, 2021.
- [239] W. S. Rossi and P. Frasca. An index for the “local” influence in social networks. In *2016 European Control Conference (ECC)*, pages 525–530, 2016.
- [240] W. S. Rossi and P. Frasca. Mean-field analysis of the convergence time of message-passing computation of harmonic influence in social networks. *IFAC-PapersOnLine*, 50(1):2409–2414, 2017. 20th IFAC World Congress.
- [241] W. S. Rossi and P. Frasca. Asynchronous opinion dynamics on the k -nearest-neighbors graph. In *IEEE Conference on Decision and Control*, pages 3648–3653, Miami, FL, USA, Dec. 2018.
- [242] W. S. Rossi and P. Frasca. Effects of network communities and topology changes in message-passing computation of harmonic influence in social networks. *arXiv preprint arXiv:1804.07093*, 2018.
- [243] W. S. Rossi and P. Frasca. On the convergence of message passing computation of harmonic influence in social networks. *IEEE Transactions on Network Science and Engineering*, 6(2):116–129, 2019.

- [244] W. S. Rossi and P. Frasca. Opinion dynamics with topological gossiping: Asynchronous updates under limited attention. *IEEE Control Systems Letters*, 4(3):566–571, 2020.
- [245] W. S. Rossi, P. Frasca, and F. Fagnani. Transient and limit performance of distributed relative localization. In *IEEE Conference on Decision and Control*, pages 2744–2748, Maui, HI, USA, Dec. 2012.
- [246] W. S. Rossi, P. Frasca, and F. Fagnani. Limited benefit of cooperation in distributed relative localization. In *IEEE Conference on Decision and Control*, pages 5427–5431, Florence, Italy, Dec. 2013.
- [247] W. S. Rossi, P. Frasca, and F. Fagnani. Average resistance of toroidal graphs. *SIAM Journal on Control and Optimization*, 53(4):2541–2557, 2015.
- [248] W. S. Rossi, P. Frasca, and F. Fagnani. Distributed estimation from relative and absolute measurements. *IEEE Transactions on Automatic Control*, 62(12):6385–6391, 2017.
- [249] W. S. Rossi, J. W. Polderman, and P. Frasca. The closed loop between opinion formation and personalized recommendations. *IEEE Transactions on Control of Network Systems*, 9(3):1092–1103, 2022.
- [250] G. S. Seyboth, D. V. Dimarogonas, K. H. Johansson, P. Frasca, and F. Allgöwer. On robust synchronization of heterogeneous linear multi-agent systems with static couplings. *Automatica*, 53:392–399, 2015.
- [251] D. Shah and J. Shin. Dynamics in congestion games. *Performance Evaluation Review*, 38(1):107–118, 2010.
- [252] N. B. Shah and M. J. Wainwright. Simple, robust and optimal ranking from pairwise comparisons. *Journal of machine learning research*, 18(199):1–38, 2018.
- [253] G. Shi, C. Altafini, and J. S. Baras. Dynamics over signed networks. *SIAM Review*, 61(2):229–257, 2019.
- [254] G. Shi, A. Proutiere, M. Johansson, J. S. Baras, and K. H. Johansson. The evolution of beliefs over signed social networks. *Operations Research*, 64(3):585–604, 2016.
- [255] J. Shin, L. Jian, K. Driscoll, and F. Bar. The diffusion of misinformation on social media: Temporal pattern, message, and source. *Computers in Human Behavior*, 83:278–287, 2018.
- [256] B. Sprenger, G. De Pasquale, R. Soloperto, J. Lygeros, and F. Dörfler. Control strategies for recommendation systems in social networks. *arXiv preprint arXiv:2403.06152*, 2024.
- [257] R. Stern, S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, R. Haulcy, H. Pohlmann, F. Wu, B. Piccoli, B. Seibold, J. Sprinkle, and D. Work. Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transportation Research Part C: Emerging Technologies*, 89:205–221, 2018.
- [258] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S.-i. Tadaki, and S. Yukawa. Traffic jams without bottlenecks experimental evidence for the physical mechanism of the formation of a jam. *New Journal of Physics*, 10((3)):033001, 2008.
- [259] Y. Sun, A. Cuesta-Infante, and K. Veeramachaneni. The backfire effects of fairness constraints. In *ICML 2022 Workshop on Responsible Decision Making in Dynamic Environments*, Baltimore, Maryland, USA, 2022.
- [260] D. Swaroop and J. K. Hedrick. String stability of interconnected systems. *IEEE Transactions on Automatic Control*, 41(3):349–357, 1996.
- [261] F. A. Taha, K. Rokade, and F. Parise. Gradient dynamics in linear quadratic network games with time-varying connectivity and population fluctuation. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 1991–1996, 2023.
- [262] A. Tahbaz-Salehi and A. Jadbabaie. A necessary and sufficient condition for consensus over random networks. *IEEE Transactions on Automatic Control*, 53(3):791–795, 2008.

- [263] J. Thai, N. Laurent-Brouty, and A. M. Bayen. Negative externalities of GPS-enabled routing applications: A game theoretical approach. In *IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pages 595–601, 2016.
- [264] Y. Tian and L. Wang. Dynamics of opinion formation, social power evolution, and naïve learning in social networks. *Annual Reviews in Control*, 55:182–193, 2023.
- [265] J. Török, G. Iñiguez, T. Yasserli, M. San Miguel, K. Kaski, and J. Kertész. Opinions, conflicts, and consensus: modeling social dynamics in a collaborative environment. *Physical review letters*, 110(8):088701, 2013.
- [266] T. Toso, A. Y. Kibangou, and P. Frasca. Impact on traffic of delayed information in navigation systems. *IEEE Control Systems Letters*, 7:1500–1505, 2023.
- [267] T. Toso, A. Y. Kibangou, and P. Frasca. Modeling the impact of route recommendations in road traffic. In *22nd World Conference of the International Federation of Automatic Control*, 2023.
- [268] T. Toso, A. Y. Kibangou, and P. Frasca. Modeling the impact of route recommendations in road traffic. *IFAC-PapersOnLine*, 56(2):4179–4185, 2023. 22nd IFAC World Congress.
- [269] T. Toso, A. Y. Kibangou, and P. Frasca. Potential degradation of transportation network efficiency due to route recommendations. *arXiv preprint arXiv:2303.17923*, 2023.
- [270] L. Tumash, C. Canudas-de Wit, and M. L. Delle Monache. Multi-directional continuous traffic model for large-scale urban networks. *Transportation research part B: methodological*, 158:374–402, 2022.
- [271] D. Urbig. Attitude dynamics with limited verbalisation capabilities. *Journal of Artificial Societies and Social Simulation*, 6(1):1–23, 2003.
- [272] V. S. Varma, I.-C. Morărescu, and M. Ayouni. Analysis of opinion dynamics under binary exogenous and endogenous signals. *Nonlinear Analysis: Hybrid Systems*, 38:100910, 2020.
- [273] V. S. Varma, I. C. Morarescu, and D. Nescic. Open multi-agent systems with discrete states and stochastic interactions. *IEEE Control Systems Letters*, 2(3):375–380, July 2018.
- [274] L. Vassio, F. Fagnani, P. Frasca, and A. Ozdaglar. Message passing optimization of harmonic influence centrality. *IEEE Transactions on Control of Network Systems*, 1(1):109–120, 2014.
- [275] T. Venturini. From fake to junk news, the data politics of online virality. In D. Bigo, E. Isin, and E. Ruppert, editors, *Data Politics: Worlds, Subjects, Rights*. Routledge, London, 2019.
- [276] E. I. Verriest. Pseudo-continuous multi-dimensional multi-mode systems: Behavior, structure and optimal control. *Discrete Event Dynamic Systems*, 22:27–59, 2012.
- [277] A. Vespignani. Predicting the behavior of techno-social systems. *Science*, 325(5939):425–428, 2009.
- [278] N. K. Vishnoi. $Lx = b$. *Laplacian Solvers and Their Algorithmic Applications*, volume 8 of *Foundations and Trends® in Theoretical Computer Science*. Now Publishers, Inc., 2013.
- [279] R. Vizuete, C. M. de Galland, J. M. Hendrickx, P. Frasca, and E. Panteley. Resource allocation in open multi-agent systems: an online optimization analysis. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 5185–5191. IEEE, 2022.
- [280] R. Vizuete, P. Frasca, and F. Garin. Graphon-based sensitivity analysis of SIS epidemics. *IEEE Control Systems Letters*, 4(3):542–547, 2020.
- [281] R. Vizuete, P. Frasca, and E. Panteley. On the influence of noise in randomized consensus algorithms. *IEEE Control Systems Letters*, 5(3):1025–1030, 2020.
- [282] R. Vizuete, P. Frasca, and E. Panteley. Gradient descent for resource allocation with packet loss. *IFAC-PapersOnLine*, 55(13):109–114, 2022.

- [283] R. Vizuete, P. Frasca, and E. Panteley. On Lyapunov functions for open Hegselmann-Krause dynamics. *arXiv preprint arXiv:2303.07074*, 2023.
- [284] R. Vizuete, P. Frasca, and E. Panteley. SIS epidemics on open networks: A replacement-based approximation. *arXiv preprint arXiv:2403.16727*, 2024.
- [285] R. Vizuete, F. Garin, and P. Frasca. The Laplacian spectrum of large graphs sampled from graphons. *IEEE Transactions on Network Science and Engineering*, 8(2):1711–1721, 2021.
- [286] R. Vizuete, C. Monnoyer de Galland, P. Frasca, E. Panteley, and J. M. Hendrickx. Trends and questions in open multi-agent systems. In *Hybrid and Networked Dynamical Systems: Modeling, Analysis and Control*, pages 219–252. Springer, 2024.
- [287] U. von Luxburg, M. Belkin, and O. Bousquet. Consistency of spectral clustering. *The Annals of Statistics*, pages 555–586, 2008.
- [288] C. Wang, Q. Li, E. Weinan, and B. Chazelle. Noisy Hegselmann-Krause systems: Phase transition and the 2R-conjecture. *Journal of Statistical Physics*, 166(5):1209–1225, 2017.
- [289] J. G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers, Part II*, 1:325–378, 1952.
- [290] J. Wei, X. Yi, H. Sandberg, and K. Johansson. Nonlinear consensus protocols with applications to quantized systems. In *IFAC World Congress*, pages 16010–16015, Toulouse, France, July 2017.
- [291] P. Wieland, R. Sepulchre, and F. Allgöwer. An internal model principle is necessary and sufficient for linear output synchronization. *Automatica*, 47(5):1068–1074, 2011.
- [292] M. Xue, Y. Tang, W. Ren, and F. Qian. Stability of multi-dimensional switched systems with an application to open multi-agent systems. *Automatica*, 146:110644, 2022.
- [293] Y. Yang, D. V. Dimarogonas, and X. Hu. Opinion consensus of modified Hegselmann-Krause models. *Automatica*, 50(2):622–627, 2014.
- [294] Y. Yi, T. Castiglia, and S. Patterson. Shifting opinions in a social network through leader selection. *IEEE Transactions on Control of Network Systems*, 8(3):1116–1127, 2021.
- [295] J. Zhang and Y. Hong. Opinion evolution analysis for short-range and long-range Deffuant-Weisbuch models. *Physica A*, 392(21):5289–5297, 2013.
- [296] Y. Zhang, E. Levina, and J. Zhu. Estimating network edge probabilities by neighbourhood smoothing. *Biometrika*, 104(4):771–783, 2017.
- [297] M. Zhu and S. Martínez. Discrete-time dynamic average consensus. *Automatica*, 46(2):322–329, 2010.
- [298] S. Öncü, J. Ploeg, N. van de Wouw, and H. Nijmeijer. Cooperative adaptive cruise control: Network-aware analysis of string stability. *IEEE Transactions on Intelligent Transportation Systems*, 15(4):1527–1537, 2014.

Appendix A

Curriculum Vitae

A.1 Positions and employment

Current position

CNRS Researcher (October 2016 – today)

Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, GIPSA-Lab, 38000 Grenoble, France

Previous positions

Assistant Professor (June 2013 – September 2016)

Hybrid Systems group

Department of Applied Mathematics

Faculty of Electrical Engineering, Mathematics and Computer Science

University of Twente

Enschede, The Netherlands

Post-doc researcher (February 2010 – May 2013)

Dipartimento di Scienze Matematiche, Politecnico di Torino

Working with Fabio Fagnani

Post-doc researcher (January 2009 – February 2010)

Dipartimento di Ingegneria dell'Informazione e Matematica Applicata, Università di Salerno & Istituto per le Applicazioni del Calcolo, C.N.R., Rome, Italy

Working with Benedetto Piccoli

Ph.D. student (January 2006 – December 2008)

Dipartimento di Matematica, Politecnico di Torino, Torino, Italy

Courtesy appointments

Research Associate (March 2013 – October 2016)

IEIIT-CNR Institute, National Research Council, Torino

Adjoint faculty member (November 2016 – October 2020)

Department of Applied Mathematics, University of Twente

A.2 Education

Ph.D., Matematica per le Scienze dell’Ingegneria. (March 2009)

Politecnico di Torino, Torino, Italy

Thesis: *Distributed optimization with communication constraints*

Advisors: Fabio Fagnani, Sandro Zampieri

Laurea Specialistica (M.S.), Ingegneria Matematica, *cum laude*. (October 2005)

Politecnico di Torino

Thesis: *Modelli cinetici alle velocità discrete per il traffico veicolare* (in italian)

Advisor: Nicola Bellomo

Laurea (B.S.), Matematica per le Scienze dell’Ingegneria, *cum laude*. (October 2003)

Politecnico di Torino

Thesis: *Il problema di Stokes. Spazi funzionali e formulazioni variazionali* (in italian)

Advisor: Claudio Canuto

Academic certificates

BKO–University Teaching Qualification (September 2016)

Teaching and supervision skills certificate for university lecturers in the Netherlands

ASN–National Scientific Habilitation as Associate Professor (October 2014)

09/G1 – Automatica (Required to become Associate Professor in Italy)

Qualification aux fonctions de Maître de Conférences (2010)

Section 26 (Required to become Associate Professor in France)

A.3 Awards

Prix de thèse Impact 2022 de la Fondation CentraleSupélec

Deuxième Prix, catégorie Impact Science (as advisor of Renato Vizuete)

2022 Outstanding Student Paper Prize, granted by the Networks and Communication Systems Technical Committee of the IEEE Control Systems Society for the paper “C. Monnoyer, R. Vizuete, J.M. Hendrickx, P. Frasca, and E. Panteley. Random coordinate descent algorithm for Open Multi-Agent Systems with complete topology and homogeneous agents. IEEE Conference on Decision and Control, December 2021, pp. 1701–1708” (as advisor of Renato Vizuete)

2013 SIAG/CST Best SICON Paper Prize, granted by the SIAM Activity Group on Control and Systems Theory for the paper “F. Bullo, R. Carli, and P. Frasca. Gossip coverage control for robotic networks: dynamical systems on the space of partitions. SIAM Journal on Control and Optimization, 50 (1), pp. 419-447, 2012”

A.4 Invited faculty positions & academic visits

In this section I list some relevant visits at academic institutions, including some Invited Professor positions.

University of Cagliari, Cagliari, Italy. *Invited Professor*. April–May 2017

Institut National des Sciences Appliquées & LAAS CNRS, Toulouse, France. *Invited Professor*. August–September 2016

Université Paul Sabatier & LAAS CNRS, Toulouse, France. *Invited Professor*. June–July 2016

University of California at Riverside, Riverside (CA), USA. Host: Fabio Pasqualetti. December 2014

University of California at Santa Barbara, Santa Barbara (CA), USA. Host: Francesco Bullo. September–October 2012, June 2011, and February–July, 2008

Université catholique de Louvain, Louvain-la-Neuve, Belgium. Host: Jean-Charles Delvenne. April 2011

Rutgers University of New Jersey (Camden), Camden (NJ), USA. Host: Benedetto Piccoli. May 2010

Lunds Tekniska Högskola, Lund, Sweden. Host: Anders Rantzer. February–April 2010

A.5 Supervised students and postdocs

In this section, I provide a full list of the students and young researchers that have supervised. For each student, I report some key information about funding and joint publications (the latter are referenced per the lists in Appendix B). A description of the research of selected students of mine of their subsequent career is available in Section 1.3.

Post-docs

3. **Aurélien Velleret** January 2022–December 2022
 “Epidemics on finite and infinite networks”
 Jointly advised with F. Garin, J.-F. Delmas and P.-A. Zitt
 Supported by a MODCOV19 grant
 Joint publications: [P3]
2. **Andrea Mattioni** October 2021–September 2022
 “Reinforcement learning control of complex systems”
 Jointly advised with Christophe Prieur
 Supported by MIAI Grenoble Artificial Intelligence Institute
 Joint publications: [J57]
1. **Wilbert Samuel Rossi** January 2015–December 2018
 “Network dynamics”
 University of Twente. Jointly advised with Jan Willem Polderman
 Supported by University of Twente
 Joint publications (as post-doc): [J52] [J37][J33] [BC3] [C38] [C31] [C27]

PhD students

7. **Raoul Prisant** October 2023–ongoing
 EEATS doctoral school, Grenoble. Jointly advised (50%) with Federica Garin
 Supported by ANR grant COCOON

6. **Tommaso Toso** October 2021–ongoing
 EEATS doctoral school, Grenoble. Jointly advised (50%) with Alain Kibangu
 Supported by IRGA (*Initiative de Recherche Grenoble Alpes*) grant ON-ROUTE
 Joint publications: [J56][C49]
5. **Maria Castaldo** November 2019–November 2022
 “Attention dynamics on YouTube: conceptual models, temporal analysis of engagement metrics, fake views”
 EEATS doctoral school, Grenoble. Jointly advised (50%) with T. Venturini
 Supported by CNRS MITI 80 PRIME grant “DOOM” (Disorders of Online Media)
 Joint publications: [J51] [J48] [J42] [BC5]
4. **Renato Sebastian Vizuite Haro** September 2019–September 2022
 “Contributions to open multi-agent systems: consensus, optimization and epidemics”
 Paris Saclay University. Jointly advised (50%) with Elena Panteley
 Supported by ANR project “HANDY– Hybrid and Networked Dynamical Systems”
 Joint publications: [J60] [J41] [BC6][C50] [C45] [C46] [C48]
3. **Denis Nikitin** September 2018–September 2021
 “Scalable large-scale control of network aggregates”
 EEATS doctoral school, Grenoble. Jointly advised (50%) with C. Canudas De Wit
 Supported by ERC project “Scale-freeBack”
 Joint publications: [J58] [J54] [J49] [C44][C43] [C41]
2. **Nicolas Martin** December 2016–February 2020
 “Network partitioning algorithms with scale-free objective”
 EEATS doctoral school, Grenoble. Jointly advised (50%) with C. Canudas De Wit
 Supported by ERC project “Scale-freeBack”
 Joint publications: [J39] [J34] [C40] [C37] [C33]
1. **Francesco Acciani** September 2014–June 2022
 “Control over unreliable networks: Consensus and platooning”
 University of Twente. Jointly advised (50%) with Geert Heijenk and Anton Stoorvogel
 Supported by University of Twente
 Joint publications: [J50] [J30] [C34] [C32] [C28]

Master theses

17. **Raoul Prisant** October 2022–September 2023
 “Continuous Opinion and Discrete Actions in social dynamics: analysis of a quantized model”
 Politecnico di Torino. Jointly advised with Francesca Ceragioli
 Joint publications: [C51]
16. **Antoine Houssard** May 2021–October 2021
 “Toxicity on Twitch: theory and empirical work”
 Sorbonne University. Co-advised with Floriana Gargiulo and Tommaso Venturini
 Supported by 80 PRIME DOOM. A. Houssard is a PhD student in Paris.
15. **Luca Cataldo** March 2021–October 2021
 “Social dynamics with quantized states: simulations and analytical results”

Politecnico di Torino. Jointly advised with Francesca Ceragioli
 Joint publications: [C51]

14. **Xiaochen Li** February 2021–July 2021
 “The effects of Covid-19 on mobility”
 INSA Centre Val de Loire. Co-advised with Maria Laura Delle Monache
 Supported by Inria Grenoble
13. **Ali Salame** February 2021–June 2021
 “Stability and Control of vehicle platoons with connected Autonomous Vehicle”
 MiSCIT master, UGA. Co-advised with M. L. Delle Monache and F. Garin
 Supported by GIPSA-lab
12. **Cristina Magnetti Gisolo** September 2020–February 2021
 “Nonlinear analysis of stability and safety of Optimal Velocity Model vehicle groups on ring roads”
 Politecnico di Torino. Jointly advised with M. L. Delle Monache and F. Ferrante
 Supported by ANR HANDY
 Joint publications: [J55]
11. **Massimo Bini** November 2019–February 2020
 “Optimal sparse control of dynamics over networks”
 Politecnico di Torino. Jointly advised with Chiara Ravazzi and Fabrizio Dabbene
 Supported by ANR HANDY
 Joint publications: [J53]
 M. Bini went to the University of Tuebingen for a PhD in machine learning.
10. **Edwige De Colnet** May 2020–September 2020
 “Observing filter bubbles in online shopping”
 INSA Rennes. Jointly advised with Tommaso Venturini and Floriana Gargiulo
 Supported by Grenoble Data Institute, in collaboration with CDiscount, Bordeaux
9. **Renato Vizuet** February–June 2019
 “Graphons: from data models to control systems for large-scale networks”
 GIPSA-lab, Grenoble. Co-advised with F. Garin
 Joint publications: [J44] [J36]
8. **Tienyi Wei** February–June 2019
 “Stabilizing traffic via autonomous vehicles”
 MS, GIPSA-lab, Grenoble. Co-advised with M.L Delle Monache
7. **Alexandre Olikier** December 2017–June 2018
 “Open multi-agent systems with fixed size and possibly not complete topologies”
 Université catholique de Louvain. Jointly advised with Julien Hendrickx
6. **Nelson Chan** July 2015–March 2016
 “On the analysis and design of algorithms for robust estimation from relative measurements”
 University of Twente
 Joint publications: [J31].
 N. Chan went to University of Groningen as a Ph.D. student. He is now Assistant Professor at the University College, University of Groningen, NL
5. **Martijn de Roo** September 2014–July 2015
 “Optimal event handling by multiple UAVs”

University of Twente. Jointly advised with Raffaella Carloni, UT
 Joint publications: [C26]

4. **Floris Broekema** September 2014–May 2015
 “Robust sensor coverage in a two-dimensional area”
 University of Twente
3. **Bram de Witte** September 2014–April 2015
 “Network games, information spread and endogenous security” (MS thesis), University of Twente
 Jointly advised with Bastiaan Overvest, Centraal Planbureau (CPB), the Hague
 Joint publications: [J40]
2. **Luca Vassio** November 29, 2012
 “Harmonic influence centrality on networks: a message passing optimization approach”
 Politecnico di Torino. Jointly advised with F. Fagnani and A. Ozdaglar
 Joint publications: [J17] [C24]
 L. Vassio went to DIMEAS, Politecnico di Torino and Avio S.p.a. as a researcher.
 He is now Assistant Professor at Politecnico di Torino
1. **Wilbert S. Rossi** July 9, 2011
 “Localizzazione relativa su grandi reti” (Relative localization on large networks)
 Politecnico di Torino, *cum laude*. Jointly advised with F. Fagnani and P. Colaneri
 Joint publications: [J26] [J21] [C22] [C15]
 W.S. Rossi went to Politecnico di Torino as a Ph.D. student. He then joined me
 as a Postdoc in Twente and later became Assistant Professor at the University
 College, University of Groningen, NL

A.6 Editorial and referee activities

Since 2008, I have regularly acted as referee for journal and conference papers. I have received the **2015 Outstanding Reviewer Award** from the IEEE Transactions on Automatic Control.

Since 2013, I have been Associate Editor in the editorial boards of several conferences and journals of the control system community.

Below, I list my main responsibilities in editorial boards, award committees and reviewing funding proposals.

Journals

- Automatica, Associate Editor (2021–today)
- IEEE Control Systems Letters, Associate Editor (2017–2021)
- Asian Journal of Control (Wiley), Associate Editor (2017–2022)
- International Journal of Robust and Nonlinear Control (Wiley), Subject Editor (2014–2018)

Conferences (International Program Committees)

- IFAC World Congress: 2020, 2023

A.7. LEADERSHIP ROLES, SERVICE TO THE COMMUNITY AND ORGANIZATION ACTIVITIES⁹³

- 2021 IFAC Symposium on Robust Control Design (ROCOND 2021)
- 55th IEEE Conference on Decision and Control, Associate Editor at Large, 2016
- IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys): 2016, 2018, 2020
- International Symposium on Mathematical Theory of Networks and Systems (MTNS): 2014, 2016, 2018
- IEEE Robotics and Automation Society, CASE Conference Editorial Board, Associate Editor (2016–2017)
IEEE Intl. Conference on Automation Science and Engineering (CASE), 2017
- European Control Association (EUCA), Conference Editorial Board, Associate Editor (2013–2016)
2014, 2015, 2016 European Control Conference
- IEEE Control System Society, Conference Editorial Board, Associate Editor (2013–2016)
2014, 2015, 2016, 2017 American Control Conference
2013, 2014, 2015, 2016 IEEE Conference on Decision and Control

Award Committees:

- National French Best PhD Thesis “GdR MACS” (2017, 2018, 2024)
- Outstanding Student Paper Prize, IEEE CSS TC on Networks and Communications (2023)

Referee of funding proposals:

French National Research Agency ANR (2014, 2015, 2016), European Research Council (2017, 2018, 2022, 2024), ETHZ Research Commission (2018), Italian Ministry of University and Scientific Research–MIUR (2018, 2021), Dutch Research Council NWO (2019), Israel Science Foundation (2020)

A.7 Leadership roles, service to the community and organization activities

Local service

2020–: Team leader of “Dynamics and Control of Networks (DANCE)”, joint research team between *GIPSA-Lab* and *Inria Centre at the University Grenoble Alpes*.

The DANCE team, for which I coordinate the research activities, currently includes five tenured members (including myself) and about ten non-tenured members (PhD students, post-docs, research scientists, and engineers). The tenured members are independent PIs who are Associate Professors at the University of Grenoble or Researchers at CNRS or INRIA. The research of the DANCE team focuses on control methods for large-scale socio-technical systems, with a particular emphasis on mobility applications.

Service to national community

Member of the steering committee of the GdR MACS (body coordinating national activities in Automatic Control in France), 2019–today

Participation to hiring committees

Concours CR2 (Tenured Researcher), INRIA Saclay, Paris, Spring 2017

Participation to PhD committees and PhD theses reviews

16. Vera Sosnovik. *Detection and analysis of issue and political ads*. Université Grenoble Alpes, PhD advisor: Oana Goga. September 4, 2023
15. Olivier Lindamoulage de Silva. *On the efficiency of decentralized epidemic management and competitive viral marketing*. PhD Advisors: Irinel-Constantin Morarescu, Samson Lasaulce. Université de Lorraine, September 28, 2023
14. Jonathan Adams. *Mathematical modelling of person-to-person opinion exchange: understanding and quantifying polarisation*. Queensland University of Technology, Brisbane, Australia. PhD advisor: Gentry White. *Referee*, 2023
13. Carmela Bernardo. *Convergence analysis of heterogeneous opinion dynamics with bounded confidence and stubbornness*. University of Sannio, Benevento, Italy. Ph.D. advisor: Francesco Vasca, *Referee*, 2022
12. Zohreh Sanai. *Coordination of Open Multi-Agent Systems*. University of Cagliari, Italy. PhD advisors: Carla Seatzu and Mauro Franceschelli. January 2022
11. Diego Deplano. *Coordination of multi-agent systems: stability via nonlinear Perron-Frobenius theory and consensus for desynchronization and dynamic estimation*. University of Cagliari, Italy. PhD advisors: Alessandro Giua and Mauro Franceschelli. *Referee*, March 2021
10. Francesco Sasso. *Distributed Bayesian methods for estimation and learning over networks*. University of Salento, Lecce, Italy. Ph.D. advisors: Giuseppe Notarstefano and Angelo Coluccia, *Referee*, 2020
9. Wenjing Yang. *Influence Maximization in Social Networks*. Université Aix-Marseille, France. PhD advisors: Alessandro Giua and Leonardo Brenner. November 2019
8. Zhiyang Ju. *Persistent Communication Connectivity of Multi-agent Systems*. University of Melbourne, Australia. PhD advisors: Dragan Nesić and Iman Shames. *Referee*, February 2019
7. Pierre-Yves Chevalier. *Inhomogeneous Products of Stochastic Matrices with Application to Consensus Systems*. Université catholique de Louvain, Louvain-la-Neuve, Belgium. Ph.D. advisors: Julien Hendrickx and Raphael Jungers, June 2018
6. Domenico Tangredi. *Consensus in Heterogeneous Opinion Dynamics Networks*. University of Sannio, Benevento, Italy. Ph.D. advisor: Francesco Vasca, *Referee*, 2018
5. Florian Dietrich. *Analyse et contrôle de systèmes de dynamiques d'opinions*. CRAN, Université de Lorraine, Nancy, France. Ph.D. advisors: Marc Jungers and Samuel Martin, *Examinateur*, November 22, 2017

4. Laura Dal Col, *On distributed control analysis and design for multi-agent systems subject to limited information*. Institut National des Sciences Appliquées, Toulouse, France. Ph.D. advisors: Luca Zaccarian and Sophie Tarbouriech, *Examineur*, October 25, 2016
3. Georg S. Seyboth, *On distributed and cooperative control design for networks of dynamical systems*. University of Stuttgart, Germany. Ph.D. advisor: Frank Allgöwer, January 15, 2016
2. Francesco Ferrante, *On quantization and sporadic measurements in control systems: Stability, stabilization, and observer design*. Université Paul Sabatier, Toulouse, France. Ph.D. advisors: Frédéric Gouaisbaut and Sophie Tarbouriech, *Examineur*, October 21, 2015
1. Matin Jafarian, *Coordination with binary controllers: Formation control and disturbance rejection*. University of Groningen, the Netherlands. Ph.D. advisor: Claudio De Persis, June 25, 2015

Participation to Licentiate and Master committees (abroad)

2. Gustav Lindmark. *Methods and algorithms for control input placement in complex networks*. Linköping University, Sweden. Advisor: Claudio Altafini. August 2018
1. Charles Monnoyer. *Efficiency bounds in computations in open multi-agent systems*. Université catholique de Louvain, Louvain-la-Neuve, Belgium. Advisor: Julien Hendrickx. June 2018

Organization of scientific events

- Open Invited Track on “Control for Socio-Technical Network Systems”, IFAC World Congress, Yokohama, Japan, July 2023 (with G. Como, K. Savla, F. Parise)
- Co-organizer of ANR HANDY workshop on “Hybrid and Networked Dynamical Systems”, Toulouse, France, June 2022
- Co-organizer (with M. Fiacchini) of the 42-nd Grenoble Summer School of Automatic Control on the topic “Data and Learning for Control”, September 6-10, 2021
- Organizer of the Automatic Control Seminar Series, GIPSA-lab, 2017–2021.
- Open Invited Track on “Social Systems: Dynamics, Games and Control on Networks”, IFAC World Congress, Berlin, Germany, July 2020 (with G. Como)
- GdR MACS Workshop “CRISE-MACS: Pendant et au-delà de la crise sanitaire COVID-19”, July 7, 2020 (Online).
- International workshop on “Disorders of Online Media” (with T. Venturini as PI of the DOOM project), Paris, November 16-17, 2019
- Co-organizer of the international ERC Scale-free Back workshop on “Analysis and Control of Large-Scale Complex Networks”, Grenoble, September 2018
- Invited session on “Model reduction and control in large-scale networks”. 2018 European Control Conference. Cyprus, June 2018 (with C. Canudas)

- Co-organizer of the international ERC Scale-free Back workshop on “Modelling reduction tools for large-scale complex networks”, Grenoble, September 2017
- Open Invited session on “Dynamics and control in social networks”, IFAC World Congress, July 2017 (with G. Como)
- Workshop on “Dynamics and Control in Social Networks”, IEEE CDC, Las Vegas, Nevada, December 11, 2016
- Invited session on “Dynamics in social networks: Opinions, games and optimization”, IEEE Conference on Decision and Control, December 2014 (with C. Ravazzi, H. Ishii, R. Tempo)
- Local organizer and co-chair (with F. Dabbene), 2014 Workshop on Uncertain Dynamical Systems (WUDS 2014), Amsterdam, Netherlands, August 2014
- Member of the National Organizing Committee, 21st International Symposium on Mathematical Theory of Networks and Systems (MTNS 2014), Groningen, Netherlands, July 2014
- International workshop on “Algorithms and dynamics over networks”, Dipartimento di Scienze Matematiche, Politecnico di Torino, February 8–9, 2012 (with G. Calafiore, F. Fagnani, and S. Zampieri)
- International workshop on “Algorithms and dynamics over networks”, Dipartimento di Matematica, Politecnico di Torino, July 14–15, 2010 (with G. Calafiore, F. Ceragioli, and F. Fagnani)
- Invited session on “Crowd and swarm dynamics: interactions, self-organization, mathematics, applications”, 10th biannual meeting of SIMAI (Società Italiana di Matematica Applicata ed Industriale), Cagliari, Italy, June 21–25, 2010 (with A. Tosin)

Membership of professional societies

IEEE, *Member*, 2013–2018. *Senior Member*, 2018–today
Control Systems Society, 2013–today

SIAM, *Member*, 2023–today

SIMAI (Italian Society for Industrial and Applied Mathematics), 2010–2020

Membership of activity groups

CSS Technical Committee “Networks and Communications Systems”, 2013–today
Working Group on Control of Network Systems, *chair*, 2015–today

CSS Technical Committee “Smart Cities”, 2024–today

IFAC Technical Committee 1.5 on Networked Systems, 2013–today

IFAC Technical Committee 2.5 on Robust Control, 2014–today

IFAC Technical Committee 9.2 on Social Impact of Automation, 2018–today

Dutch 4TU Federation, ICT Next Generation (ICTng), Core team member, 2015–2016

A.8 Funding and research grants

Responsibility of externally funded research activities

COCOON (Continuous Methods for the Control of Large Networks).

ANR PRME (2022–2027) PI: P. Frasca. Amount: 408k €

Abstract: COCOON is a 4-year research project funded by ANR (the French national science foundation). The COCOON project advocates a scalable approach to large networks that is based on continuous network models instead of the usual (discrete) graphs. Towards this broad objective, this proposal aims at concurrently developing and cross-fertilising two promising methods to define continuous dynamics that approximate large-network dynamics: (1) Using graph limit objects such as graphons; (2) Defining analog approximations through a continuation process that replaces a large systems of ordinary differential equations with a single partial differential equation. These methods can be beneficial in a multitude of potential applications: the project will address three distinct applications with potentially high societal impact: epidemic models, electro-mobility networks and, with a bigger thrust, multimodal mobility networks.

HANDY (Hybrid and Networked Dynamical Systems).

ANR PRC (2019–2023). PI: Luca Zaccarian. Co-PI, local leader: P. Frasca. Amount: 105k€ for the Grenoble unit

Abstract: Networked dynamical systems are ubiquitous in current and emerging technologies. From energy grids, fleets of connected autonomous vehicles to online social networks, the same scenario arises in each case: dynamical units interact locally to achieve a global behavior. When considering a networked system as a whole, very often continuous-time dynamics are affected by instantaneous changes, called jumps, leading to so-called hybrid dynamical systems. Hybrid phenomena thus play an essential role in these control applications, and call upon the development of novel adapted tools for stability and performance analysis and control design. In this context, the aim of HANDY project is to provide methodological control-oriented tools for realistic networked models, which account for hybrid phenomena. The project brings together researchers from LAAS in Toulouse, CRAN in Nancy, GIPSA in Grenoble and LSS in Gif-sur-Yvette, with expertise in various domains of automatic control, ranging from geometric control and optimization, switched systems, hybrid dynamics, nonlinear control, and multi-agent systems. See also: <http://projects.laas.fr/handy>

DOOM (Systems-theory for the Disorders Of Online Media).

80 PRIME grant from CNRS MITI (2019–2022). PI: P. Frasca. Amount: 20k € + 1 PhD salary

Abstract: Online social media have a key role in contemporary society and the debates that take place on them are known to shape political and societal trends. For this reason, pathological phenomena like the formation of “filter bubbles” and the viral propagation of “fake news” are observed with concern. The scientific assumption of this proposal is that these information disorders are direct consequences of the inherent nature of these communication media, and more specifically of the collective dynamics of attention thereby. In order to capture these dynamics, this proposal advocates the mathematical modelling of the interplay between the medium (algorithmic component) and the users (human component). The resulting dynamics shall be explored by a system-theoretic approach, using notions such as feedback and stability. This quantitative and rigorous approach will not only unlock fundamental insights but also deliver suggestions on

suitable policies to manage the media. See also: <https://cis.cnrs.fr/doom>. The follow-up project **ADOBE (Attention Dynamics in Online Media and Beyond)**, in collaboration with F. Gargiulo and T. Venturini, has also been funded by the MITI (18 k€, 2022–2023).

ON-ROUTE (Online Routing Recommendations in Multimodal Transportation: Impact, Potential Shortcomings, and their Mitigation).

IRGA grant from Grenoble IDEX (2021–2024). PI: A. Kibangu, co-PI: P. Frasca. Amount: 5 k€ + 1 PhD salary

Abstract: Millions of users rely on Online Routing Applications (ORAs) like Waze, Google Maps, and TomTom for navigation guidance, impacting traffic congestion and transportation modes. ORAs function as social feedback systems, gathering and influencing user choices in dynamic traffic networks. Our project shifts focus from service providers to transportation managers. Online routing is in general formulated as a multicriteria optimization problem, which is solved by the ORA to satisfy user needs: instead, the transportation network manager aims at optimizing some overall measure of the efficiency of the network. This project therefore aims at analyzing the effect of ORAs on transportation networks and at finding mitigation strategies against their adverse effects, thanks to their regulation or to suitable traffic control actions, such as variable speed limits or intelligent traffic lights.

C2S2 (Control of Cyber-Social Systems)

IRS grant from Grenoble IDEX (2017–2018). PI: P. Frasca. Amount: 8 k€

Abstract: The project concentrates on cyber-social systems, that is, complex systems with interacting social and technological components. A strong motivation for this novel research direction comes from the need for innovative tools for the management of vehicular traffic. In this application, state-of-the-art approaches concentrate on hard control actions, like traffic lights: instead, future management methods should exploit soft control actions aimed at controlling the traffic demand, that is, the aggregated behaviors of the drivers. Within this perspective, this project explores a selection of original scientific challenges related to controlling the distribution of the system’s state.

MOB (Models of Bubbles in Online Social Networks)

CNRS PEPS S2IH 2018. Amount: 10 k€

Abstract: This exploratory project focuses on the effects of online recommendation systems on social dynamics, which may entail the formation of “filter bubbles” that distort the experience of the users. The project will develop a mathematical model to demonstrate these effects and propose designs for their mitigation. The research will be conducted by a blend of tools from dynamical systems, network theory, complex systems, and control systems.

Participation to externally funded research activities

- FeedingBias ANR PRC project. PI: Gilles Bastin (2023–2026)
- “Algorithms & Society” MIAI Chair (AI and Sociology). Core faculty member. PI: Gilles Bastin (2019–2023)
- “AI and dynamical systems: new paradigms for control and robots” MIAI Chair (AI and Robotics). Core faculty member. PI: Christophe Prieur (2019–2023)

- “Contextual Recommendations in Action - Bridging AI and Real-Life Economics” MIAI Chair. Associated faculty member. PI: Sihem Amer-Yahia (2019–2023)
- Scale-FreeBack ERC Advanced Grant. PI: Carlos Canudas de Wit. (October 2016–2021)
- COOPS Joint International Lab, between CNR–IEIIT and Osaka University/Tokyo Institute of Technology/Mitsubishi Electric Corporation. Topic: Cooperative Control of Energy Management Systems. Project leader: R. Tempo. (Dec. 2014 – December 2017)
- Institute for Collaborative Biotechnologies (ICB), Santa Barbara, USA. “Opinion Dynamics in Social Networks: Peer-to-Peer Interactions, Information Assimilation, and Strategic Manipulation”. (External collaborator) Project leader: F. Bullo. (Nov. 2011 – Dec. 2012)

A.9 Outreach and non-academic collaborations

Industrial collaborations

- CDiscount, Bordeaux, France. Topic: Ethics & algorithms (2019–2020).
- SELEX Sistemi Integrati (Finmeccanica group), Rome, Italy. Topic: simulation and radar detection of stochastic signals (2009–2010)

Press coverage

- Interview by Marie Lyan, La Tribune – Auvergne-Rhône-Alpes, about our Covid-19 research published in PLOS Computational Biology (30/08/2021).
- Stress, travail, médias... Comment avons-nous vécu le confinement ?. Article on “Le Journal du CNRS” on online behaviors during COVID-19 lock-down (09/06/2020)
- Press release by INS2I institute about project DOOM, 2019
- Online technology magazine Tech Xplore article about recommender systems, September 2018

A.10 Teaching

These are graduate and undergraduate courses that I have designed and thought. My book with Fabio Fagnani originated from the lecture notes of our courses at Politecnico di Torino. Notice that since I became a CNRS researcher, I have had no teaching duties, but I have been teaching one graduate course since 2021 and I intend to expand my teaching activities in the future.

- **Topics in Smart Transportation.** Master Autonomous Robotic Systems (MARS), Institut National Polytechnique de Grenoble. 8 h in fall 2021, 2022, 2023
- **Introduction to network coordination,** MiSCIT: Master in Systems, Control & IT, Grenoble, France. 15 h, winter 2015.

- **Network dynamics** (with G. Como and M. Cao), DISC graduate school, Utrecht. 8 h, fall 2014.
- **Signals and systems**. Undergraduate course, Creative Technology program—Smart Technology, University of Twente. 3 ECTS in fall 2014, 2015
- **Signals and systems**. Undergraduate course, Creative Technology program—New Media, University of Twente. 3 ECTS in fall 2014, 2015
- **Hybrid dynamical systems**. Master in Applied Mathematics, University of Twente. 5 ECTS in spring 2014, 2015, 2016
- **Dynamics over networks** (with F. Fagnani). Master in Mathematical Engineering, Politecnico di Torino. 20 h, April–June 2012 and 2013
- **Control of/over networks** (with F. Fagnani). Master in Automatica and Control Technologies, Politecnico di Torino. 20 h, January–March 2011 and 2012

A.11 Invited presentations

54. “Graphons and dynamics on large graphs, from a control systems perspective”, Symposium on Collective Models for Networked Particle Systems, University of Pavia, April 16, 2024
53. “Navigation systems in traffic networks: Route recommendations and performance degradation”, Focus Period Symposium on Network Dynamics and Control, Linköping, September 20–22, 2023
52. “Distributed algorithms over graphs: Consensus, gossiping, relative estimation, distributed optimization”, Peyresq Summer School, Peyresq, France, June 25–July 1, 2023
51. “Potential deterioration in transportation network efficiency due to route recommendations”, Games, Learning, and Networks. Workshop on Games on networks, Singapore, 3–6 April 2023
50. “The closed loop between opinion formation and personalised recommendations”, Laboratoire Jacques-Louis Lions, Paris Sorbonne, November 25, 2022
49. “The closed loop between opinion formation and personalised recommendations”, Workshop on “Stochastic modeling and complex networks: application to the dynamics of opinions in social networks”, Avignon, October 2022
48. “Graphons: A tool for the analysis of dynamical systems on large networks”, Scuola Superiore Meridionale, Naples, Italy, November 25, 2021
47. “The closed loop between opinion formation and personalised recommendations”, Technion, Haifa, Israel, January 4, 2021
46. “Traffic flow on a ring with a single autonomous vehicle: An interconnected stability perspective”, Southeast University, Nanjing, China, September 4, 2020
45. “The closed loop between opinion formation and personalised recommendations”, Peking University, Beijing, China, July 21, 2020

44. "The closed loop between opinion formation and personalised recommendations", Workshop "Network Dynamics in the Social, Economic, and Financial Sciences", Turin, Italy, November 5-8, 2019
43. "Non-smooth opinion dynamics", Workshop of the European Network for Nonsmooth Dynamical Systems, Grenoble, September 18, 2019
42. "The closed loop between opinion formation and personalised recommendations", Workshop "Reti sociali e comportamenti emergenti", Napoli, February 4, 2019
41. "The harmonic influence in social networks and its distributed computation by message passing", IXXI, ENS Lyon, July 3, 2018
40. "Randomization and quantization in opinion dynamics", IRSTEA, Clermont-Ferrand, March 14, 2018
39. "Message-passing computation of the harmonic influence in social networks", L2S, Paris-Saclay, November 21, 2017
38. "Harmonic influence in social networks and identification of influencers by message passing", WUDS'17 workshop, Banyuls-sur-mer, July 6, 2017
37. "Non-smooth dynamical systems in opinion dynamics", University of Twente, Enschede, NL, June 15, 2017
36. "The observability radius of network systems", University of Cagliari, Cagliari, Italy, May 4, 2017
35. "Non-smooth and hybrid systems in opinion dynamics", IEEE CDC satellite workshop on Dynamics and Control in Social Networks, Las Vegas, Nevada, December 11, 2016
34. "Non-smooth and hybrid systems in opinion dynamics", ANR Workshop "Control subject to computational and communication constraints" (CO4), Toulouse, France, October 26-28, 2016
33. "The observability radius of network systems", LAAS-CNRS, Toulouse, France, September 29, 2016
32. "Non-smooth systems in opinion dynamics", LAAS MAC workshop, Banyuls-sur-Mer, France, June 23, 2016
31. "The harmonic influence in social networks and its distributed computation", LAAS-CNRS, Toulouse, France, June 17, 2016
30. "The local influence in social networks and its distributed estimation", Université catholique de Louvain, Louvain-la-neuve, Belgium, June 7, 2016
29. "The observability radius of network systems", COOPS workshop, Osaka, Japan, December 12, 2015
28. "Distributed estimation from relative and absolute measurements", Kyoto University, Japan, December 11, 2015
27. "Distributed estimation from relative and absolute measurements", Osaka University, Japan, December 9, 2015

26. “Ergodic opinion dynamics from prejudices and gossiping”. Workshop on Mathematical Models in Social Dynamics, Politecnico di Torino, Torino, Italy, October 1-2, 2015
25. “Distributed estimation from relative and absolute measurements”, LAAS Toulouse, France, June 23, 2015
24. “Distributed estimation from relative measurements: Fundamental limitations, algorithms, and application to power systems” (**keynote**), SICE International Symposium on Control Systems (ISCS), Tokyo, Japan, March 4–7, 2015
23. “Distributed estimation from relative measurements”, TU Delft, the Netherlands, February 18, 2015
22. “Message passing optimization of harmonic influence centrality”, SCONES workshop, Boston University, Boston, USA, October 27–28, 2014
21. “Ergodic dynamics in social networks”, WUDS’14 workshop, Amsterdam, August 22, 2014
20. “On the most influential node in a social network: competing leaders in opinion dynamics”, CWTS-UT mini-workshop on Network Analysis, University of Twente, June 18, 2014
19. “Opinion dynamics in social networks: modelling, analysis, and control”, TU Eindhoven, May 28, 2014 and Zilverling colloquium, University of Twente, June 10, 2014
18. “Estimation from relative measurements: fundamental limitations and distributed algorithms in sensor networks”, TU Delft, The Netherlands, Novembre 15, 2013
17. “Topics in social networks: opinion dynamics and control”, IEIIT-CNR, Torino, Italy, May 23, 2013
16. “Robust self-triggered coordination by ternary controllers”, Gipsa-Lab, Grenoble, France, November 15, 2012
15. “The wisdom of randomly interacting crowds”, University of California at Santa Barbara, September 26 and University of California at San Diego, September 28, 2012
14. “Robust self-triggered coordination by ternary controllers”, University of California at Los Angeles, September 25, 2012
13. “On the mean square error of randomized averaging algorithms”, Universität Stuttgart, Germany, May 15, 2012
12. “Coordination with little communication: Averaging, animal groups, automated deployment”, Rijksuniversiteit Groningen, The Netherlands, February 21, 2012
11. “Coverage control via gossip communication”, Rijksuniversiteit Groningen, The Netherlands, August 24, 2011
10. “Discontinuous feedback in consensus-seeking systems: Bounded confidence, quantization, hysteresis”, University of California, Santa Barbara, USA, June 17, 2011
9. “Estimation of Gaussian signals via hexagonal sensor networks: Wildfire and pollutant detection”. CASHMA FIRB review meeting, Università di Salerno, Salerno, Italy, March 1, 2011

8. “Gossip coverage control on graphs”, Université catholique de Louvain, Louvain-la-Neuve, Belgium, April 5, 2011
7. “Broadcast gossip averaging algorithms: interference and asymptotical error”, Kungliga Tekniska Högskola, Stockholm, Sweden, April 16, 2010
6. “Distributed averaging on digital noisy networks”, Lunds Tekniska Högskola, Lund, Sweden, April 9, 2010
5. “Broadcast gossip averaging algorithms: interference and asymptotical error”, Lunds Tekniska Högskola, Lund, Sweden, March 2, 2010
4. “Average consensus with communication constraints”. Meeting on mathematical control theory: controllability, optimization, stability (organized with PRIN 2006 funding). Dipartimento di Matematica, Politecnico di Torino, Torino, Italy, April 20–21, 2009
3. “The role of communication constraints in optimization problems over networks”. IAC-CNR, Roma, Italy, March 2009
2. “Gossip coverage control”. Dipartimento di Ingegneria dell’Informazione, Università di Padova, Padova, Italy, Sep. 2008
1. “Quantized gossip consensus”, Arizona State University, Tempe (AZ), USA, April 15, 2008

Conference presentations (contributed)

20. “Attention dynamics and disorders on YouTube”. CCS/Italy 2023, Naples, Italy. October 9–11, 2023
19. “SIS models on (not so) dense large random networks”, SIAM Conference on Applications of Dynamical Systems (DS23), May 14–18, Portland, Oregon, 2023
18. “Modeling limited attention in opinion dynamics by topological interactions”. NET-GCOOP 2020, September 23, 2021
17. “A network reduction method inducing scale-free degree distributions”. European Control Conference, Limassol, Cyprus, June 15, 2018
16. “Achieving robust average consensus dynamics over wireless networks”. European Control Conference, Aalborg, Denmark, June 29, 2016
15. “An index for the local influence in social networks”. European Control Conference, Aalborg, Denmark, June 29, 2016
14. “Harmonic influence in large-scale networks: Analysis, optimization, and applications to opinion dynamics and distributed control”, Mini-course at the International Symposium on Mathematical Theory of Networks and Systems, Groningen, Netherlands, July 2014 (with Giacomo Como and Fabio Fagnani)
13. “A message passing algorithm for the evaluation of social influence”. European Control Conference 2014, Strasbourg, France, June 2014
12. “Self-triggered rendezvous of gossiping second-order agents”. IEEE Conference on Decision and Control, Florence, Italy, Dec. 2013

11. “Gossips and Prejudices: Ergodic Randomized Dynamics in Social Networks”. IFAC Workshop on Distributed Estimation and Control in Networked Systems, Koblenz, Germany, September 25–26, 2013 (poster)
10. “On optimal coverage with unreliable sensors”. IFAC Workshop on Distributed Estimation and Control in Networked Systems, Koblenz, Germany, September 25–26, 2013 (poster)
9. “An ergodic randomized algorithm for relative localization”. European Control Conference, Zürich, Switzerland, July 18, 2013
8. “Robust self-triggered coordination with ternary controllers”. 32nd Benelux Meeting on Systems and Control, Houffalize, Belgium, March 26–28, 2013
7. “Self-triggered coordination with ternary controllers”. 3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems, Santa Barbara, CA, September 14, 2012 (poster)
6. “Continuous-time discontinuous equations in bounded confidence opinion dynamics”. IFAC World Congress, Milan, Italy, Aug. 2011
5. “The asymptotical error of broadcast gossip averaging algorithms”. IFAC World Congress, Milan, Italy, Aug. 2011
4. “Discontinuous feedback in consensus-seeking systems: Bounded confidence, quantization, hysteresis”, 2011 Santa Barbara Control Workshop, Santa Barbara, USA, June 24, 2011 (poster)
3. “Mathematical issues of distributed complex systems”. M. Anile SIMAI workshop, Roma, Italy, Oct. 9, 2009
2. “Average consensus on networks with transmission noise or quantization”. European Control Conference, Kos, Greece, July 2007
1. “Agreement on networks”. V Séminaire Transalpin de Physique, Torino, Italy, Feb. 2007 (poster)

Appendix B

Publications

This appendix is devoted to presenting the publications that I have authored based on my research. I will commence by offering an overview of my publication practices and output in Section B.1, inclusive of references to collaborations and activity as advisor. Following this, in Section B.2, I will present a selection of papers organized by topic. Finally, Section B.3 contains the full list of my publications, categorized by type (books, book chapters, journal papers, conference papers).

B.1 Overview

Summary of figures: Since 2007, I have (co)authored 1 book, 1 edited book, 6 book chapters, 60 articles in international journals, and 51 full-length articles in the proceedings of international conferences. My publications span over 25 journals and the proceedings of numerous conferences. Among the journals, the IEEE Transactions on Automatic Control (10 articles) and the IEEE Transactions on Control of Network Systems (7 articles) stand out as my most frequent venues. Among the conferences, most are conferences in control systems that are organized by IEEE or IFAC: the most frequent is the IEEE Conference on Decision and Control with 14 articles.

With over 120 co-authors, all but one of my publications are collaborations with colleagues or students. The vast majority of my journal publications (all but 7) have not involved my Ph.D. advisors. Over time, joint publications with my students have increasingly constituted a substantial portion of my work, with their contributions becoming the predominant majority since 2021.

Main academic collaborations (after my PhD):

Simone Baldi, Delft Technical University and Southeast University: coordination of heterogeneous multi-vehicle systems, 2014–2024

Francesco Bullo, University of California, Santa Barbara: robotic networks, 2008–2016

Carlos Canudas-de-Wit, CNRS, Grenoble: large-scale systems, 2016–2023

Francesca Ceragioli, Politecnico di Torino: non-smooth systems, 2009–ongoing

Claudio De Persis, University of Groningen: quantized and hybrid systems, 2010–2013

Federica Garin, INRIA, GIPSA-Lab, Grenoble: network systems, 2009–ongoing

Julien Hendrickx, Université catholique de Louvain: multi-agent systems, 2011–2024

Roberto Tempo, IEIIT-CNR, Torino: randomized dynamics, 2012–2015.

B.2 Selected papers on key topics

Theory of Control, Multi-agent and Network Systems:

- D. Nikitin, C. Canudas de Wit, and P. Frasca. A continuation method for large-scale modeling and control: from ODEs to PDE, a round trip. *IEEE Transactions on Automatic Control*, 67 (10): 5118–5133, 2022
- M. Franceschelli and P. Frasca. Stability of Open Multi-Agent Systems and applications to dynamic consensus. *IEEE Transactions on Automatic Control*, 66 (5): 2326–2331, 2021
- R. Vizuete, F. Garin, and P. Frasca. The Laplacian spectrum of large graphs sampled from graphons. *IEEE Transactions on Network Science and Engineering*, 8 (2): 1711–1721, 2021
- S. Baldi and P. Frasca. Leaderless synchronization of heterogeneous oscillators by adaptively learning the group model. *IEEE Transactions on Automatic Control*, 65 (1): 412–418, 2020
- Chiara Ravazzi, Paolo Frasca, Roberto Tempo, and Hideaki Ishii. Ergodic randomized algorithms and dynamics over networks. *IEEE Transactions on Control of Network Systems*, 2 (1): 78–87, 2015
- Paolo Frasca and Julien M. Hendrickx. On the mean square error of randomized averaging algorithms. *Automatica*, 49 (8): 2496–2501, 2013
- Francesco Bullo, Ruggero Carli, and Paolo Frasca. Gossip coverage control for robotic networks: dynamical systems on the space of partitions. *SIAM Journal on Control and Optimization*, 50 (1): 419–447, 2012
- Francesca Ceragioli, Claudio De Persis, and Paolo Frasca. Discontinuities and hysteresis in quantized average consensus. *Automatica*, 47(9): 1916–1928, 2011

Opinion Dynamics and Computational Social Sciences:

- W.S. Rossi, J.W. Polderman, and P. Frasca. The closed loop between opinion formation and personalised recommendations, *IEEE Transactions on Control of Network Systems*, 9 (3): 1092–1103, 2022
- M. Castaldo, T. Venturini, P. Frasca, and F. Gargiulo. Junk News Bubbles: Modelling the rise and fall of attention in online arenas. *New Media & Society*, February 2021
- M. Castaldo, F. Gargiulo, T. Venturini, and P. Frasca. The Rhythms of the Night: increase in online night activity and emotional resilience during the Covid-19 lockdown. *EPJ Data Science*, 10:7, 2021
- W.S. Rossi and P. Frasca. Opinion dynamics with topological gossiping: asynchronous updates under limited attention. *IEEE Control Systems Letters*, 4 (3), 566–571, 2020
- Francesca Ceragioli and Paolo Frasca. Consensus and disagreement: the role of quantized behaviours in opinion dynamics. *SIAM Journal on Control and Optimization*, 56(2): 1058–1080, 2018

- Paolo Frasca, Chiara Ravazzi, Roberto Tempo, and Hideaki Ishii. Gossips and prejudices: ergodic randomized dynamics in social networks. *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Koblenz, Germany, Sep. 2013, pp. 212–219

Intelligent Transportation Systems:

- Francesco Acciani, Paolo Frasca, Geert Heijenk, and Anton A. Stoorvogel. Stochastic string stability of vehicle platoons via cooperative adaptive cruise control with lossy communication. *IEEE Transactions on Intelligent Transportation Systems*, 23 (8): 10912–10922, 2022
- D. Liu, S. Baldi, V. Jain, W. Yu, and P. Frasca. Cyclic communication in adaptive strategies to platooning: the case of synchronized merging. *IEEE Transactions on Intelligent Vehicles*, 6 (3): 490–500, 2021
- V. Giammarino, S. Baldi, P. Frasca and M.L. Delle Monache. Traffic flow on a ring with a single autonomous vehicle: an interconnected stability perspective, *IEEE Transactions on Intelligent Transportation Systems*, 22 (8): 4998–5008, 2021

Cooperative Robotics

- M. de Roo, P. Frasca, and R. Carloni. Optimal event handling by multiple Unmanned Aerial Vehicles. *IEEE Int. Conference on Robotics and Automation (ICRA)*, May 2016, pp. 1230–1236
- J.W. Durham, R. Carli, P. Frasca, and F. Bullo. Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics*, 28 (2): 364–378, 2012

Biology (Epidemics and Animal Groups):

- S. Mauras, V. Cohen-Addad, G. Duboc, M. Dupré la Tour, P. Frasca, C. Mathieu, L. Opatowski, L. Viennot. Mitigating COVID-19 outbreaks in workplaces and schools by hybrid telecommuting. *PLOS Computational Biology*, 17 (8): e1009264, August 2021
- A. Aydogdu, P. Frasca, et al.. Modeling birds on wires. *Journal of Theoretical Biology*, 415:102–112, 2017
- E. Cristiani, P. Frasca, and B. Piccoli. Effects of anisotropic interactions on the structure of animal groups. *Journal of Mathematical Biology*, 62 (4): 569–588, 2011

B.3 Full list of publications, per type

B.3.1 Books

- [B1] Fabio Fagnani and Paolo Frasca. *Introduction to Averaging Dynamics over Networks*. Vol. 472 in Lecture Notes in Control and Information Sciences (LNCIS) Series, Springer, 2017. ISBN: 978-3-319-68021-7.

This self-contained textbook provides a concise introduction to the theory of averaging dynamics and consensus. With approximately one hundred exercises, it thoroughly covers the fundamental concepts while exploring the main applications of averaging in

modeling, estimation, and control challenges across diverse domains, including robotic, sensor, and social networks.

B.3.2 Edited books

- [E1] Romain Postoyan, Paolo Frasca, Elena Panteley, Luca Zaccarian. *Hybrid and Networked Dynamical Systems*. Lecture Notes in Control and Information Sciences (LNCIS) Series, volume 493, Springer, 2024. ISBN: 978-3-031-49554-0.

B.3.3 Chapters in books

- [BC6] Renato Vizuete, Charles Monnoyer de Galland, Paolo Frasca, Elena Panteley, Julien Hendrickx. Trends and Questions in Open Multi-agent Systems. In: R. Postoyan, P. Frasca, E. Panteley, L. Zaccarian (eds) *Hybrid and Networked Dynamical Systems*. Lecture Notes in Control and Information Sciences, vol 493. Springer, Cham, 2024
- [BC5] Maria Castaldo, Paolo Frasca, and Tommaso Venturini. Online Attention Dynamics in Social Media. A. M. Annaswamy, P. P. Khargonekar, F. Lamnabhi-Lagarrigue, S. K. Spurgeon (eds.), *Cyber-Physical-Human Systems: Fundamentals and Applications*, Wiley, 2023, pp. 491-510
- [BC4] Francesca Ceragioli, Paolo Frasca, Benedetto Piccoli, Francesco Rossi. Generalized solutions to opinion dynamics models with discontinuities. N. Bellomo, L. Gibelli (eds.), *Crowd Dynamics, Volume 3: Modeling and Social Applications in the Time of COVID-19*, Modeling and Simulation in Science, Engineering and Technology, Springer, 2021
- [BC3] Francesca Ceragioli, Paolo Frasca and Wilbert Samuel Rossi. Modeling limited attention in opinion dynamics by topological interactions. In *Network Games, Control and Optimization. NETGCOOP 2021*, Lasaulce S., Mertikopoulos P., Orda A. (eds). *Communications in Computer and Information Science*, vol 1354. Springer, Cham, 2021, pp. 272–281
- [BC2] Francesca Ceragioli and Paolo Frasca. Discontinuities, generalized solutions, and (dis)agreement in opinion dynamics. In *Control Subject to Computational and Communications Constraints*, S. Tarbouriech, A. Girard and L. Hetel, Eds., Springer, LNCIS Series, 2018, pp. 287-309.
- [BC1] Claudio De Persis and Paolo Frasca. Hybrid coordination of flow networks. In *Hybrid Systems with Constraints*, J. Daafouz and S. Tarbouriech and M. Sigalotti, Eds., Wiley-ISTE, 2013

B.3.4 Journal papers

- [J60] Charles Monnoyer de Galland de Carnières, Renato Vizuete, Julien M. Hendrickx, Elena Panteley, and Paolo Frasca. Random Coordinate Descent for Resource Allocation in Open Multi-Agent Systems. *IEEE Transactions on Automatic Control*, Early Access, 2024 <https://doi.org/10.1109/TAC.2024.3394349>
- [J59] Di Liu, Sebastian Mair, Kang Yang, Simone Baldi, Paolo Frasca, and Matthias Althoff. Resilience in platoons of cooperative heterogeneous vehicles: self-organization strategies and provably-correct design. *IEEE Transactions on Intelligent Vehicles*, 9 (1): 2262–2275, 2024

- [J58] Denis Nikitin, Carlos Canudas-de-Wit, Paolo Frasca, and Ursula Ebels. Synchronization of Spin-Torque Oscillators via Continuation Method. *IEEE Transactions on Automatic Control*, 68 (11): 6621–6635, 2023
- [J57] Andrea Mattioni, Samuele Zoboli, Bojan Mavkov, Daniele Astolfi, Vincent Andrieu, Emmanuel Witrant, Paolo Frasca, Christophe Prieur. Enhancing deep reinforcement learning with integral action to control Tokamak safety factor, *Fusion Engineering and Design*, Volume 196, 114008, 2023
- [J56] Tommaso Toso, Alain Y. Kibangou and Paolo Frasca, Impact on traffic of delayed information in navigation systems, *IEEE Control Systems Letters*, 7: 1500–1505, 2023.
- [J55] Cristina Magnetti Gisolo, Maria Laura Delle Monache, Francesco Ferrante, and Paolo Frasca. Nonlinear analysis of stability and safety of Optimal Velocity Model vehicle groups on ring roads. *IEEE Transactions on Intelligent Transportation Systems*, 23 (11): 1524–9050, 2022
- [J54] Denis Nikitin, Carlos Canudas-de-Wit, and Paolo Frasca. A continuation method for large-scale modeling and control: from ODEs to PDE, a round trip. *IEEE Transactions on Automatic Control*, 67 (10): 5118–5133, 2022
- [J53] Massimo Bini, Paolo Frasca, Chiara Ravazzi, and Fabrizio Dabbene. Graph structure-based heuristics for optimal targeting in social networks. *IEEE Transaction on Control of Network Systems*, 9 (3): 1189–1201, 2022
- [J52] Wilbert Samuel Rossi, Jan Willem Polderman, and Paolo Frasca. The closed loop between opinion formation and personalised recommendations, *IEEE Transactions on Control of Network Systems*, 9 (3): 1092–1103, 2022
- [J51] Maria Castaldo, Tommaso Venturini, Paolo Frasca, and Floriana Gargiulo. Junk News Bubbles: Modelling the rise and fall of attention in online arenas. *New Media & Society*, 24 (9), 2027–2045, 2022
- [J50] Francesco Acciani, Paolo Frasca, Geert Heijenk, and Anton A. Stoorvogel. Stochastic string stability of vehicle platoons via cooperative adaptive cruise control with lossy communication. *IEEE Transactions on Intelligent Transportation Systems*, 23 (8): 10912–10922, 2022
- [J49] Denis Nikitin, Carlos Canudas-de-Wit, and Paolo Frasca. Control of average and deviation in large-scale linear networks. *IEEE Transactions on Automatic Control*, 67 (4): 1639–1654, 2022
- [J48] Floriana Gargiulo, Maria Castaldo, Tommaso Venturini and Paolo Frasca. Distribution of labor, productivity and innovation in collaborative science. *Applied Network Science*, 7:19, March 2022
- [J47] Di Liu, Simone Baldi, Vishrut Jain, Wenwu Yu, and Paolo Frasca. Cyclic communication in adaptive strategies to platooning: the case of synchronized merging. *IEEE Transactions on Intelligent Vehicles*, 6 (3): 490–500, 2021
- [J46] Simon Mauras, Vincent Cohen-Addad, Guillaume Duboc, Max Dupre la Tour, Paolo Frasca, Claire Mathieu, Lulla Opatowski, Laurent Viennot. Mitigating COVID-19 outbreaks in workplaces and schools by hybrid telecommuting. *PLOS Computational Biology*, 17 (8): e1009264, August 2021

- [J45] Vittorio Giammarino, Simone Baldi, Paolo Frasca and Maria Laura Delle Monache. Traffic flow on a ring with a single autonomous vehicle: an interconnected stability perspective, *IEEE Transactions on Intelligent Transportation Systems*, 22 (8): 4998–5008, 2021
- [J44] Renato Vizuete, Federica Garin, and Paolo Frasca. The Laplacian spectrum of large graphs sampled from graphons. *IEEE Transactions on Network Science and Engineering*, 8 (2): 1711–1721, 2021
- [J43] Mauro Franceschelli and Paolo Frasca. Stability of Open Multi-Agent Systems and applications to dynamic consensus. *IEEE Transactions on Automatic Control*, 66 (5): 1558–2523, 2021
- [J42] Maria Castaldo, Tommaso Venturini, Paolo Frasca, and Floriana Gargiulo. The Rhythms of the Night: increase in online night activity and emotional resilience during the spring 2020 Covid-19 lockdown. *EPJ Data Science* 10: 7, 2021.
- [J41] Renato Vizuete, Paolo Frasca, and Elena Panteley. On the influence of noise in randomized consensus algorithms. *IEEE Control Systems Letters*, 5 (3): 1025–1030, 2021
- [J40] Bram de Witte, Paolo Frasca, Bastiaan Overvest, and Judith Timmer. Protecting shared information in networks: a network security game with strategic attacks. *International Journal of Robust and Nonlinear Control*, 30 (11): 4255–4277, 2020
- [J39] Nicolas Martin, Paolo Frasca, and Carlos Canudas-de-Wit. Subgraph detection for average detectability of LTI systems. *IEEE Transactions on Network Science and Engineering*, 7 (4): 2787–2798, 2020
- [J38] Christian Commault, Jacob van der Woude, and Paolo Frasca. Functional target controllability of networks: structural properties and efficient algorithms. *IEEE Transactions on Network Science and Engineering*, 7 (3): 1521–1530, 2020
- [J37] Wilbert Samuel Rossi and Paolo Frasca. Opinion dynamics with topological gossiping: asynchronous updates under limited attention. *IEEE Control Systems Letters*, 4 (3), 566–571, 2020
- [J36] Renato Vizuete, Paolo Frasca, and Federica Garin. Graphon-based sensitivity analysis of SIS epidemics. *IEEE Control Systems Letters*, 4 (3), 542–547, 2020
- [J35] Simone Baldi and Paolo Frasca. Leaderless synchronization of heterogenous oscillators by adaptively learning the group model. *IEEE Transactions on Automatic Control*, 65 (1): 412–418, 2020
- [J34] Nicolas Martin, Paolo Frasca, and Carlos Canudas-de-Wit. Large-scale network reduction towards scale-free structure. *IEEE Transactions on Network Science and Engineering*, 6 (4): 711–723, 2019
- [J33] Wilbert Samuel Rossi and Paolo Frasca. On the convergence of message-passing computation of harmonic influence in social networks. *IEEE Transactions on Network Science and Engineering*, 6 (2): 116–129, 2019
- [J32] Simone Baldi, Shuai Yuan and Paolo Frasca. Output synchronization of unknown heterogeneous agents via distributed model reference adaptation. *IEEE Transactions on Control of Network Systems*, 6 (2): 515–525, 2019

- [J31] Chiara Ravazzi, Nelson P. K. Chan, and Paolo Frasca. Distributed estimation from relative measurements of heterogeneous and uncertain quality. *IEEE Transactions on Signal and Information Processing over Networks*, 5 (2): 203–217, 2019
- [J30] Francesco Acciani, Anton Stoorvogel, Geert Heijenk, and Paolo Frasca. Achieving robust average consensus over lossy wireless networks. *IEEE Transactions on Control of Network Systems*, 6 (1): 127–137, 2019
- [J29] Simone Baldi and Paolo Frasca. Adaptive synchronization of unknown heterogeneous agents: An adaptive virtual model reference approach. *Journal of The Franklin Institute*, 356 (2): 935–955, 2019
- [J28] Paolo Frasca, Sophie Tarbouriech and Luca Zaccarian. Hybrid models of opinion dynamics with limited confidence. *Automatica*, 100: 153–161, February 2019
- [J27] Francesca Ceragioli and Paolo Frasca. Consensus and disagreement: the role of quantized behaviours in opinion dynamics. *SIAM Journal on Control and Optimization*, 56(2): 1058–1080, 2018
- [J26] Wilbert Samuel Rossi, Paolo Frasca, and Fabio Fagnani. Distributed estimation from relative and absolute measurements. *IEEE Transactions on Automatic Control*, 62 (12): 6385–6391, 2017
- [J25] Gianluca Bianchin, Paolo Frasca, Andrea Gasparri, and Fabio Pasqualetti. The observability radius of networks. *IEEE Transactions on Automatic Control*, 62 (6): 3006–3013, 2017
- [J24] Aylin Aydoğdu, Paolo Frasca, Ciro D’Apice, Rosanna Manzo, Julianne M. Thornton, Bryan Gachomo, Tevin Wilson, Brian Cheung, Umair Tariq, William Saidel, Benedetto Piccoli. Modeling birds on wires. *Journal of Theoretical Biology*, 415: 102–112, 2017
- [J23] Rushabh Patel, Paolo Frasca, Joseph W. Durham, Ruggero Carli, and Francesco Bullo. Dynamic partitioning and coverage control with asynchronous one-to-base-station communication. *IEEE Transactions on Control of Network Systems*, 3 (1): 24–33, 2016
- [J22] Paolo Frasca, Federica Garin, Balázs Gerencsér, and Julien M. Hendrickx. Optimal one-dimensional coverage by unreliable sensors. *SIAM Journal on Control and Optimization*, 53 (5): 3120–3140, 2015
- [J21] Wilbert Samuel Rossi, Paolo Frasca, and Fabio Fagnani. Average resistance of toroidal graphs. *SIAM Journal on Control and Optimization*, 53 (4): 2541–2557, 2015
- [J20] Paolo Frasca, Hideaki Ishii, Chiara Ravazzi, and Roberto Tempo. Distributed randomized algorithms for opinion formation, centrality computation and power systems estimation: A tutorial overview. *European Journal of Control*, 24: 2–13, July 2015
- [J19] Chiara Ravazzi, Paolo Frasca, Roberto Tempo, and Hideaki Ishii. Ergodic randomized algorithms and dynamics over networks. *IEEE Transactions on Control of Network Systems*, 2 (1): 78–87, 2015
- [J18] Georg S. Seyboth, Dimos V. Dimarogonas, Karl H. Johansson, Paolo Frasca, and Frank Allgöwer. On robust synchronization of heterogeneous linear multi-agent systems with static couplings. *Automatica*, 53: 392–399, March 2015

- [J17] Luca Vassio, Fabio Fagnani, Paolo Frasca, and Asuman Ozdaglar. Message passing optimization of Harmonic Influence Centrality. *IEEE Transactions on Control of Network Systems*, 1 (1): 109–120, 2014
- [J16] Rushab Patel, Paolo Frasca, and Francesco Bullo. Centroidal area-constrained partitioning for robotic networks. *ASME Journal of Dynamic Systems, Measurement and Control*, 136 (3): 031024, 2014
- [J15] Claudio De Persis and Paolo Frasca. Robust self-triggered coordination with ternary controllers. *IEEE Transactions on Automatic Control*, 58 (12): 3024–3038, 2013
- [J14] Paolo Frasca and Julien M. Hendrickx. On the mean square error of randomized averaging algorithms. *Automatica*, 49 (8): 2496–2501, 2013
- [J13] Joseph W. Durham, Ruggero Carli, Paolo Frasca, and Francesco Bullo. Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics*, 28 (2): 364–378, 2012
- [J12] Francesco Bullo, Ruggero Carli, and Paolo Frasca. Gossip coverage control for robotic networks: dynamical systems on the space of partitions. *SIAM Journal on Control and Optimization*, 50 (1): 419–447, 2012
- [J11] Francesca Ceragioli and Paolo Frasca. Continuous and discontinuous opinion dynamics with bounded confidence. *Nonlinear Analysis: Real World Applications*, 13 (3): 1239–1251, 2012
- [J10] Paolo Frasca. Continuous-time quantized consensus: Convergence of Krasovskiy solutions. *Systems and Control Letters*, 61(2): 273–278, 2012
- [J9] Andrea Tosin and Paolo Frasca. Existence and approximation of probability measure solutions to models of collective behaviors. *Networks and Heterogeneous Media*, 6 (3): 561–596, 2011
- [J8] Francesca Ceragioli, Claudio De Persis, and Paolo Frasca. Discontinuities and hysteresis in quantized average consensus. *Automatica*, 47(9): 1916–1928, 2011
- [J7] Fabio Fagnani and Paolo Frasca. Broadcast gossip averaging: interference and unbiasedness in large Abelian Cayley networks. *IEEE Journal of Selected Topics in Signal Processing*, 5 (4): 866–875, 2011
- [J6] Emiliano Cristiani, Paolo Frasca, and Benedetto Piccoli. Effects of anisotropic interactions on the structure of animal groups. *Journal of Mathematical Biology*, 62 (4): 569–588, 2011
- [J5] Ruggero Carli, Giacomo Como, Paolo Frasca, and Federica Garin. Distributed averaging on digital erasure networks. *Automatica*, 47 (3): 115–121, 2011
- [J4] Ruggero Carli, Fabio Fagnani, Paolo Frasca, and Sandro Zampieri. Gossip consensus algorithms via quantized communication. *Automatica*, 46 (1): 70–80, 2010
- [J3] Paolo Frasca, Paolo Mason, and Benedetto Piccoli. Detection of Gaussian signals via hexagonal sensor networks. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1 (1–2): 39–55, 2009
- [J2] Paolo Frasca, Ruggero Carli, Fabio Fagnani, and Sandro Zampieri. Average consensus on networks with quantized communication. *International Journal of Robust and NonLinear Control*, 19 (16): 1787–1816, 2009

- [J1] Vincenzo Coscia, Marcello Delitala, and Paolo Frasca. On the mathematical theory of vehicular traffic flow II: Discrete velocity kinetic models. *International Journal of Non-Linear Mechanics*, 42 (3): 411–421, 2007

B.3.5 Conference proceedings papers

These are full-length peer-reviewed papers that appeared in the proceedings of international conferences.

- [C51] Raoul Prisant, Luca Cataldo, Francesca Ceragioli, Paolo Frasca. Disagreement, limit cycles and Zeno solutions in continuous opinion dynamics with binary actions. *2024 European Control Conference (ECC)*, Stockholm, Sweden, June 2024
- [C50] Renato Vizuete, Paolo Frasca, Elena V. Panteley. SIS epidemics on open networks: a replacement-based approximation. *2024 European Control Conference (ECC)*, Stockholm, Sweden, June 2024
- [C49] Tommaso Toso, Alain Kibangou, Paolo Frasca. Modeling the impact of route recommendations in road traffic. *IFAC World Congress*, Yokohama, Japan, July 2023
- [C48] Renato Vizuete, Charles Monnoyer de Galland de Carnières, Julien M. Hendrickx, Elena Panteley, Paolo Frasca. Resource allocation in open multi-agent systems: an online optimization analysis. *IEEE Conference on Decision and Control*, Cancun, Mexico, December 2022, pp. 5185-5191
- [C47] Paolo Frasca and Francesco Rossi. Carathéodory solutions to opinion dynamics with topological interactions and their associated graphs. *International Symposium on Mathematical Theory of Networks and Systems*, IFAC-PapersOnLine 55 (30), 436-441. *IFAC Conference on Networked Systems*, Bayreuth, Germany, September 2022.
- [C46] Renato Vizuete, Paolo Frasca and Elena Panteley, "Gradient descent for resource allocation with packet loss," IFAC-PapersOnLine 55 (13), 109-114. *IFAC Conference on Networked Systems*, Zurich, CH, July 2022.
- [C45] Charles Monnoyer de Galland, Renato Vizuete, Julien M. Hendrickx, Paolo Frasca, and Elena Panteley. Random coordinate descent algorithm for Open Multi-Agent Systems with complete topology and homogeneous agents. *IEEE Conference on Decision and Control*, December 2021, pp. 1701–1708
- [C44] Denis Nikitin, Carlos Canudas-de-Wit, and Paolo Frasca. Boundary Control for Stabilization of Large-Scale Networks through the Continuation Method. *IEEE Conference on Decision and Control*, December 2021, pp. 4786–4791
- [C43] Denis Nikitin, Carlos Canudas-de-Wit, and Paolo Frasca. Shape-based nonlinear model reduction for 1D conservation laws. *IFAC World Congress*, July 2020
- [C42] Vittorio Giammarino, Maolong Lv, Simone Baldi, Paolo Frasca and Maria Laura Delle Monache. On a Weaker Notion of String/Ring Stability for Mixed Traffic. *IEEE Conference on Decision and Control*, Nice, France, December 2019, pp. 335–340
- [C41] Denis Nikitin, Carlos Canudas-de-Wit, and Paolo Frasca. Boundary Output Regulation for Large Scale-Free Positive Network Systems. *IEEE Conference on Decision and Control*, Nice, France, December 2019, pp. 5050-5055

- [C40] Nicolas Martin, Paolo Frasca, Takayuki Ishizaki, Jun-Ichi Imura, Carlos Canudas-de-Wit. The price of connectedness in graph partitioning problems. *2019 European Control Conference (ECC)*, Naples, Italy, June 2019, pp. 2313–2318
- [C39] Mauro Franceschelli and Paolo Frasca. Proportional dynamic consensus in Open Multi-Agent Systems. *IEEE Conference on Decision and Control*, Miami, FL, USA, December 2018, pp. 900–905
- [C38] Wilbert Samuel Rossi and Paolo Frasca. Asynchronous opinion dynamics on the k -nearest-neighbors graph. *IEEE Conference on Decision and Control*, Miami, FL, USA, December 2018, pp. 3648–3653
- [C37] Nicolas Martin, Paolo Frasca, and Carlos Canudas-de-Wit. MergeToCure: a new strategy to allocate cure in an epidemic over a grid-like network using a scale-free abstraction. *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Groningen, August 2018, pp. 34–39
- [C36] Simone Baldi, Muhammad Ridho Rosa, Paolo Frasca, and Elias B. Kosmatopoulos. Platooning merging maneuvers in the presence of parametric uncertainty. *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Groningen, August 2018, pp. 148–153
- [C35] Simone Baldi, Muhammad Ridho Rosa, and Paolo Frasca. Adaptive state-feedback synchronization with distributed input: the cyclic case. *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Groningen, August 2018, pp. 1–6
- [C34] Francesco Acciani, Paolo Frasca, Anton Stoorvogel, Elham Semsar-Kazerooni, and Geert Heijenk. Cooperative adaptive cruise control over unreliable networks: an observer-based approach to increase robustness to packet loss. *2018 European Control Conference (ECC)*, Limassol, Cyprus, June 12–15, 2018, pp. 1399–1404
- [C33] Nicolas Martin, Paolo Frasca, and Carlos Canudas-De-Wit. Large-scale network reduction towards scale-free structure. *2018 European Control Conference (ECC)*, Limassol, Cyprus, June 12–15, 2018, pp. 2236–2241
- [C32] Francesco Acciani, Paolo Frasca, Anton A. Stoorvogel, Geert Heijenk. Using a linear gain to accelerate average consensus over unreliable networks. *IEEE Conference on Decision and Control*, Melbourne, Australia, December 2017, pp. 3569–3574
- [C31] Wilbert Samuel Rossi and Paolo Frasca. Mean-field analysis of the convergence time of message-passing computation of harmonic influence in social networks. *IFAC World Congress*, Toulouse, France, July 2017, pp. 2461–2466
- [C30] Paolo Frasca, Sophie Tarbouriech, and Luca Zaccarian. A hybrid model of opinion dynamics with limited confidence. *IFAC Symposium on Nonlinear Control Systems*, Monterey, CA, USA, August 2016, pp. 351–355
- [C29] Gianluca Bianchin, Paolo Frasca, Andrea Gasparri, and Fabio Pasqualetti. The observability radius of network systems. *American Control Conference*, Boston, MA, USA, July 2016, pp. 185–190
- [C28] Francesco Acciani, Geert Heijenk, and Paolo Frasca. Achieving robust average consensus dynamics over wireless networks. *European Control Conference*, Aalborg, Denmark, June 2016, pp. 555–560

- [C27] Wilbert Samuel Rossi and Paolo Frasca. An index for the local influence in social networks. *European Control Conference*, Aalborg, Denmark, June 2016, pp. 525–530
- [C26] Martijn de Roo, Paolo Frasca, and Raffaella Carloni. Optimal event handling by multiple unmanned aerial vehicles. *IEEE International Conference on Robotics and Automation*, Stockholm, Sweden, May 2016, pp. 1230–1236
- [C25] Francesca Ceragioli and Paolo Frasca. Continuous-time consensus dynamics with quantized all-to-all communication, *European Control Conference*, Linz, Austria, July 2015, pp. 1120–1125
- [C24] Luca Vassio, Fabio Fagnani, Paolo Frasca, and Asuman Ozdaglar. A message passing algorithm for the evaluation of social influence. *European Control Conference*, Strasbourg, France, June 2014, pp. 190–195
- [C23] Claudio De Persis, Paolo Frasca, and Julien M. Hendrickx. Self-triggered rendezvous of gossiping second-order agents. (Invited paper) *IEEE Conference on Decision and Control*, Florence, Italy, Dec. 2013, pp. 7403–7408
- [C22] Wilbert Samuel Rossi, Paolo Frasca, and Fabio Fagnani. Limited benefit of cooperation in distributed relative localization. *IEEE Conference on Decision and Control*, Florence, Italy, Dec. 2013, pp. 5427–5431
- [C21] Chiara Ravazzi, Paolo Frasca, Roberto Tempo and Hideaki Ishii. Almost sure convergence of a randomized algorithm for relative localization in sensor networks. (Invited paper) *IEEE Conference on Decision and Control*, Florence, Italy, Dec. 2013, pp. 4778–4783
- [C20] Rushab Patel, Paolo Frasca, and Francesco Bullo. Centroidal area-constrained partitioning for robotic networks. *ASME Dynamic Systems and Control Conference*, Stanford, CA, USA, Oct. 2013
- [C19] Paolo Frasca, Chiara Ravazzi, Roberto Tempo, and Hideaki Ishii. Gossips and prejudices: ergodic randomized dynamics in social networks. *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Koblenz, Germany, Sep. 2013, pp. 212–219
- [C18] Paolo Frasca and Federica Garin. On optimal coverage with unreliable sensors. *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Koblenz, Germany, Sep. 2013, pp. 38–42
- [C17] Chiara Ravazzi, Paolo Frasca, Hideaki Ishii, and Roberto Tempo. A distributed randomized algorithm for relative localization in sensor networks. *European Control Conference*, Zürich, Switzerland, July 2013, pp. 1776–1781
- [C16] Paolo Frasca and Julien M. Hendrickx. Large network consensus is robust to packet losses and interferences. *European Control Conference*, Zürich, Switzerland, July 2013, pp. 1782–1787
- [C15] Wilbert Samuel Rossi, Paolo Frasca, and Fabio Fagnani. Transient and limit performance of distributed relative localization. *IEEE Conference on Decision and Control*, Maui, HI, USA, Dec. 2012, pp. 2744–2748
- [C14] Claudio De Persis and Paolo Frasca. Self-triggered coordination with ternary controllers. *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Santa Barbara, CA, USA, Sep. 2012, pp. 43–48

- [C13] Joseph W. Durham, Ruggero Carli, Paolo Frasca, and Francesco Bullo. Dynamic partitioning and coverage control with asynchronous one-to-base-station communication. (Invited paper) *IEEE Conference on Decision and Control and European Control Conference*, Orlando, FL, USA, Dec. 2011, pp. 5589–5594
- [C12] Francesca Ceragioli and Paolo Frasca. Continuous-time discontinuous equations in bounded confidence opinion dynamics. *IFAC World Congress*, Milan, Italy, Aug. 2011, pp. 1986–1990
- [C11] Fabio Fagnani and Paolo Frasca. The asymptotical error of broadcast gossip averaging algorithms. *IFAC World Congress*, Milan, Italy, Aug. 2011, pp. 10027–10031
- [C10] Ruggero Carli, Giacomo Como, Paolo Frasca, and Federica Garin. Distributed averaging on digital noisy networks. (Invited paper) *Information Theory and Applications Workshop*, San Diego, CA, USA, Feb. 2011, pp. 1–9
- [C9] Francesca Ceragioli, Claudio De Persis, and Paolo Frasca. Quantized average consensus: Discontinuities and hysteresis. (Invited paper) *IFAC Symposium on Nonlinear Control Systems*, Bologna, Italy, Sep. 2010, pp. 78–83
- [C8] Joseph W. Durham, Ruggero Carli, Paolo Frasca, and Francesco Bullo. Discrete partitioning and coverage control with gossip communication. *ASME Dynamic Systems and Control Conference*, Hollywood, CA, Oct. 2009, pp. 225–232
- [C7] Ruggero Carli, Giacomo Como, Paolo Frasca, and Federica Garin. Average consensus on digital noisy networks. *IFAC Workshop on Estimation and Control of Networked Systems*, Venice, Italy, Sep. 2009, pp. 36–41
- [C6] Paolo Frasca, Ruggero Carli, Francesco Bullo. Multiagent coverage algorithms with gossip communication: Control systems on the space of partitions. *American Control Conference*, St. Louis, MO, USA, June 2009, pp. 2228–2235
- [C5] Paolo Frasca, Ruggero Carli, Fabio Fagnani, and Sandro Zampieri. Average consensus by gossip algorithms with quantized communication. (Invited paper) *IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2008, pp. 4831–4836
- [C4] Ruggero Carli, Fabio Fagnani, Paolo Frasca, and Sandro Zampieri. The quantization error in the average consensus problem. *Mediterranean Conference on Control and Automation*, Ajaccio, France, June 2008, pp. 1592–1597
- [C3] Ruggero Carli, Fabio Fagnani, Paolo Frasca, and Sandro Zampieri. A probabilistic analysis of the average consensus algorithm with quantized communication. *IFAC World Congress*, Seoul, Korea, Aug. 2008, pp. 8062–8067
- [C2] Ruggero Carli, Fabio Fagnani, Paolo Frasca, Thomas Taylor, and Sandro Zampieri. Average consensus on networks with transmission noise or quantization. *European Control Conference*, Kos, Greece, July 2007, pp. 1852–1857
- [C1] Ruggero Carli, Fabio Fagnani, Paolo Frasca, and Sandro Zampieri. Efficient quantized techniques for consensus algorithms. *International Workshop on Networked Control System: Tolerant to Faults*, Nancy, France, June 2007

B.3.6 Technical reports

These are technical reports that are not subsumed by refereed publications. Some contain results that have been presented at conferences without proceedings, others contain variations or extensions of refereed publications.

- [T4] Wilbert Samuel Rossi and Paolo Frasca. Effects of network communities and topology changes in message-passing computation of harmonic influence in social networks. (Extended abstract) *IFAC Conference on Cyber-Physical & Human-Systems*, Miami, FL, USA, December 2018. Pages 170–173. Available at <https://arxiv.org/abs/1804.07093>
- [T3] Wilbert Samuel Rossi, Paolo Frasca, and Fabio Fagnani. Note on “Average resistance of toroidal graphs” by Rossi, Frasca and Fagnani. <https://arxiv.org/abs/1702.00293>, 2017
- [T2] Francesca Ceragioli and Paolo Frasca. Continuous-time discontinuous equations in bounded-confidence opinion dynamics. <https://arxiv.org/abs/1010.0066>, 2010
- [T1] Ruggero Carli, Fabio Fagnani, Paolo Frasca, and Sandro Zampieri. Efficient quantization for average consensus. <https://arxiv.org/abs/0903.1337>, 2009

B.3.7 Preprints

These are documents that are in various stages of the peer-review and publication process.

- [P5] M. Castaldo, P. Frasca, T. Venturini, and F. Gargiulo (2022). Doing data science with platforms crumbs: an investigation into fake views on YouTube. <https://arxiv.org/abs/2210.01096>, 2022.
- [P4] Paolo Frasca, Federica Garin, and Raoul Prisant. Opinion dynamics on signed graphs and graphons: Beyond the piece-wise constant case. <https://arxiv.org/abs/2404.08372>, 2024.
- [P3] J. F. Delmas, P. Frasca, F. Garin, V. C. Tran, A. Velleret, and P. A. Zitt. Individual based SIS models on (not so) dense large random networks. <https://arxiv.org/abs/2302.13385>, 2023
- [P2] Tommaso Toso, Alain Y Kibangou, and Paolo Frasca. Potential degradation of transportation network efficiency due to route recommendations. <https://arxiv.org/abs/2303.17923>, 2023.
- [P1] Renato Vizuite, Paolo Frasca, and Elena Panteley. On Lyapunov functions for open Hegselmann-Krause dynamics. <https://arxiv.org/abs/2303.07074>, 2023